



ПРИМЛ. ЕНО: 15 03 2018			
Рад. јед.	б р о ј	Арх. шифра	Прилог
0801	353/1		

## Научном већу Института за физику у Београду

### Предлог за Студентску награду Института за физику у Београду

Поштовани,

Велико ми је задовољство да предложим др **Владимира Лончара** за Студентску награду Института за физику у Београду за докторску тезу под називом "Hybrid Parallel Algorithms for Solving Nonlinear Schrödinger Equation", коју је одбранио 17. октобра 2017. године на Природно-математичком факултету Универзитета у Новом Саду.

Истраживачки рад др Владимира Лончара је у области примене рачунара у физици, а посебно развоја паралелних и хибридних нумеричких алгоритама за решавање нелинеарних парцијалних диференцијалних једначина и њихова примена на проучавање ултрахладних бозонских система са дипол-дипол интеракцијом. Једначине овог типа се јављају не само у физици ултрахладних атома, него и у нелинеарној оптици и другим областима. Колега Лончар се бавио развојем паралелних алгоритама за решавање једног облика нелинеарне парцијалне диференцијалне једначином познате као Грос-Питаевски једначина (ГПЈ), која се јавља приликом проучавања физичког феномена познатог као Бозе-Ајнштајн кондензација, са нагласком на облик једначине, који поред нелинеарног члана садржи и интегрално језгро конволуционог типа. Решавање ГПЈ нумеричким поступком на рачунару је могуће, али су претходно били доступни само серијски алгоритми. Због великог броја рачунских операција које ови алгоритми подразумевају, њихове серијске имплементације су ограничене на симулације са веома грубом дискретизационом шемом, те се јавља потреба за убрзањем алгоритама и њихових имплементација путем паралелизације. Др Лончар је развио више паралелних алгоритама за различите хардверске архитектуре, од класичних Intel процесора, до Nvidia графичких картица, који се могу извршавати како на једном рачунару, тако и на рачунарском кластеру. У оквиру овог истраживачка правца кандидат је објавио неколико публикација у међународним часописима, односно у зборницима радова са међународних конференција.

Употреба графичких процесора као алтернатива централним процесорима рачунара за убрзање нумеричких поступака се често примењује у научним круговима, јер се поједини нумерички поступци (попут Фуријеове трансформације) много ефикасније могу извршити на графичком процесору. У овом правцу, Владимир се бавио употребом графичких процесора за решавање ГПЈ и развио паралелни алгоритам и пратећу имплементацију за Nvidia графичке процесоре. Алгоритам се ослања на више нумеричких метода, међу којима су Кренк-Николсон нумерички метод, Фуријеова трансформација и нумеричка интеграција путем Симпсоновог правила које су имплементирани као језгра која се извршавају на графичком процесору путем CUDA технологије. Посебан нагласак је стављен на оптимизацију употребе меморије на графичком процесору, које у највећем броју случајева има мање него основне RAM меморије рачунара. У тестовима перформанси, имплементација је показала значајно убрзање у односу на серијску верзију. Овај резултат детаљније је описан у раду:



- CUDA programs for solving the time-dependent dipolar Gross-Pitaevskii equation in an anisotropic trap  
V. Lončar, A. Balaž, A. Bogojević, S. Škrbić, P. Muruganandam, and S. K. Adhikari  
Comput. Phys. Commun. 200, 406 (2016). [M21a, IF(2016)=3.936]

Поред алгоритама за графичке процесоре, кандидат је развијао и паралелне алгоритме за модерне вишејезгарне процесоре као и системе са дистрибуираном меморијом односно рачунарске кластере. У развијеним алгоритмима подаци су подељени између процесирајућих елемената (нити или процеса) тако да ниједан процесирајући елемент нема приступ целом скупу података, већ извршава рачунске операције само на задатом подскупу података и по потреби врши размену података са другим процесирајућим елементима. На овај начин је омогућено извршавање симулација на много већем броју паралелних рачунарских ресурса. Са једне стране, ове имплементације омогућавају брже извршавање постојећих симулација, док са друге стране такође омогућавају и симулације са много финијим дискретизационим шемама које су неопходне за испитивање појединих физичких феномена и које до сад нису могле бити коришћене због ограничења доступне меморије једног рачунара. Алгоритме прати и детаљан преглед перформанси њихових имплементација, док је у дисертацији кандидата такође разјен и модел перформанси. Ови резултати представљени су у следећим радовима:

- OpenMP, OpenMP/MPI, and CUDA/MPI C programs for solving the time-dependent dipolar Gross-Pitaevskii equation  
V. Lončar, L. E. Young-S., S. Škrbić, P. Muruganandam, S. Adhikari, and A. Balaž  
Comput. Phys. Commun. 209, 190 (2016). [M21a, IF(2016)=3.936]
- OpenMP GNU and Intel Fortran programs for solving the time-dependent Gross-Pitaevskii equation  
L. E. Young-S., P. Muruganandam, S. Adhikari, V. Lončar, D. Vudragović, and A. Balaž  
Comput. Phys. Commun. 220, 503 (2017). [M21a, IF(2016)=3.936]

Кандидат је развијао и алгоритме за хетерогене рачунарске системе, који комбинују алгоритме за вишејезгарне процесоре и алгоритме за графичке процесоре. У оквиру овог приступа развијено је више хетерогених нумеричких метода, међу којим се издваја хетерогени алгоритам за брзу Фуријеову трансформацију, који је применљив и на много ширем низу математичких проблема. Да би се минимизовало време извршавања хетерогеног алгоритма за решавање ГПЈ, које зависи од релативних перформанси централног и графичког процесора рачунарског система који се користи, развијен је и метод оптимизације параметара алгоритма. Више различитих метода оптимизације је евалуирано, међу којима је хеуристички метод заснован на генетском алгоритму показао најбоље перформансе, и који је применљив на све хетерогене алгоритме. Хибридни алгоритми за решавање ГПЈ на хетерогеним системима детаљно су описани у докторској дисертацији кандидата, као и у следећем поглављу у тематском зборнику радова:

- Efficient Numerical Tools for Solving the Nonlinear Schroedinger Equation  
V. Lončar, I. Vasić, and A. Balaž  
Scientific Computing: Studies and Applications, Ed. C. Erling, pp. 63-157  
ISBN: 978-1-53612-564-1, NOVA Science Publishers, 2017. [M13]



Осим развоја алгоритама, кандидат се такође бавио и визуелизацијом података добијених из нумеричких симулација. За ове потребе развио је механизме преко којих се подаци лако визуелизују током извршавања симулације, и преко којих се може управљати симулацијом, тзв. *in situ* визуелизација. Развијени механизми значајно олакшавају анализу и обраду података добијених из симулације.

Користећи развијене алгоритме и механизме за обраду података кандидат је проучавао утицај дипол-дипол интеракције на особине Бозе-Ајнштајн кондензата хрома и диспрозијума, као и формирање квантних дроплета приликом нагле промене контактне интеракције у кондензатима са јаком дипол-дипол интеракцијом.

Сви ови резултати су приказани детаљно у оквиру докторске тезе под називом "Hybrid Parallel Algorithms for Solving Nonlinear Schrödinger Equation", која је написана на енглеском језику и има 10 поглавља.

У првом поглављу је дат увод у тезу и објашњен значај проучавања алгоритама за решавање парцијалних диференцијалних једначина, а посебно нелинеарне Шредингерове једначине, као и кратак преглед области науке у којима се она примењује. Такође, побројани су главни научни доприноси тезе и описана је њена структура. У другом поглављу је дат кратак преглед аналитичких и нумеричких метода за решавање нелинеарне Шредингерове једначине, односно Грос-Питаевски једначине са диполним чланом. Представљена је бездимензина форма ове једначине у три, две и једној просторној димензији, а након тога и семи-имплицитни Кренк-Николсон алгоритам подељеног корака, који представља основу за све друге алгоритме развијене у оквиру тезе.

Треће поглавље, које представља оригиналан научни допринос, даје опис паралелног нумеричког алгоритма за решавање диполне Грос-Питаевски једначине, као и имплементације овог алгоритма за системе са дељеном меморијом, уз коришћење OpenMP парадигме. У четвртном поглављу представљен је одговарајући паралелни алгоритам и његова имплементација за системе са хардверским (графичким) акцелераторима у CUDA окружењу, што такође представља оригиналан научни допринос. Пето поглавље комбинује претходна два алгоритма и описује нови, хибридни алгоритам за решавање нелинеарне Шредингерове једначине са диполним чланом. Хибридни алгоритам је намењен рачунарским системима са вишејезгарним централним процесорима и графичким картицама, и представља оригиналан научни допринос.

У шестом поглављу су описана проширења алгоритама из претходних поглавља која су намењена системима са дистрибуираном меморијом уз коришћење MPI парадигме. Ово укључује следеће алгоритме и њихове имплементације, што такође представља оригиналан научни допринос:

- имплементација која ради само на централним процесорима, заснована на алгоритму за дељену меморију, односно његовој OpenMP паралелизацији из трећег поглавља (OpenMP/MPI),
- имплементација која ради само на графичким процесорима, а заснована је на имплементацији из четвртог поглавља (CUDA/MPI),
- хибридна имплементација заснована на MPI проширењу алгоритма из петог поглавља (Hybrid/MPI).

Седмо поглавље се бави визуелизацијом велике количине података коју генеришу програми развијени у оквиру тезе и описује ефикасан приступ решавању овог проблема коришћењем VisIt софтверског окружења.

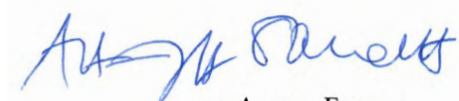
Осмо поглавље, које представља оригиналан научни допринос тезе, проучава перформансе и комплексност свих развијених алгоритама и њихових имплементација. У овом поглављу се теоријски моделира убрзање и ефикасност паралелних имплементација, а ови модели се верификују поређењем са детаљним резултатима мерења перформанси свих програма. Посебан допринос овог поглавља је примена генетског алгорита на оптимизацију параметара хибридних алгоритама, што је неопходно за максимализацију њихових перформанси. Поред могућности да се на овај начин оцене потребни рачунарски ресурси за решавање одговарајућег проблема (нелинеарне Шредингерове једначине у потребној резолуцији), развијена методологија за тестирање и модели су од ширег значаја, јер се могу користити и на другим рачунарским архитектурама, као и за друге врсте хибридних алгоритама.

Девето поглавље демонстрира како се програми развијени у оквиру ове тезе могу користити у истраживању у физици кондензоване материје, односно у области ултрахладних атома, на примеру проучавања формирања квантних вортекса у Бозе-Ајнштајн кондензатима кроз које се креће препрека (ласерски сноп). У овом поглављу се прво верификују развијени програми на примеру за који постоје експериментални резултати, а након тога се испитује зависност критичне брзине препреке за појаву вортекса од јачине дипол-дипол интеракције. Ови резултати представљају зачетак научног доприноса тезе из области физике и на овој теми кандидат планира да ради у будућем периоду. Десето поглавље сумира резултате остварене у докторској тези и даје преглед будућих праваца истраживања.

Владимир Лончар је до сада објавио један рад категорије M13 (поглавље у тематском зборнику), три рада категорије M21a, два саопштења категорије M33 и три саопштења категорије M34. Према бази Web of Science, ови радови су цитирани укупно 45 пута, а 43 пута без аутоцитата, уз Хиршов индекс 3. Један од радова категорије M21a је у Web of Science био означен као Highly Cited Paper (top 1%) за период од објављивања у марту 2016. године до октобра 2017. године.

**Имајући све наведено у виду, са задовољством предлажем др Владимира Лончара за Студентску награду Института за физику у Београду за најбољу докторску тезу одбрањену током 2017. године.**

У Београду, 15. 03. 2018. године



др Антун Балаж  
научни саветник  
Институт за физику у Београду





## Биографија: др Владимир Лончар

Владимир Лончар је рођен 28. октобра 1985. године у Новом Саду. Основне студије на Природно-математичком факултету Универзитета у Новом Саду, смер дипломирани информатичар - пословна информатика, уписао је 2004. године, а завршио 2009. године. Мастер студије на истом факултету, на смеру информациони системи, завршио је 2011. године. Школске 2011/2012 године је уписао докторске студије информатике на Департману за математику и информатику Природно-математичког факултета Универзитета у Новом Саду, а докторску дисертацију под називом "Hybrid Parallel Algorithms for Solving Nonlinear Schrödinger Equation" одбранио је 17. октобра 2017. године под руководством др Антуна Балажа и проф. др Срђана Шкрбића.

Др Владимир Лончар је од 2012. до краја 2014. године активно учествовао у развоју информационог система Природно-математичког факултета у Новом Саду, где је био и запослен. Од фебруара 2015. је запослен у Институту за физику у Београду, у оквиру Лабораторије за примену рачунара у науци Националног центра изузетних вредности за изучавање комплексних система, на пројекту основних истраживања ОН171017 "Моделирање и нумеричке симулације сложених вишечестичних система".

Главна тема истраживања др Владимира Лончара су паралелни и хибридни нумерички алгоритми за решавање нелинеарних парцијалних диференцијалних једначина и њихова примена на проучавање ултрахладних бозонских система са дипол-дипол интеракцијом. До сада је објавио један рад категорије M13, три рада категорије M21a, два саопштења категорије M33 и три саопштења категорије M34. Према бази Web of Science, ови радови су цитирани укупно 45 пута, а 43 пута без аутоцитата, уз Хиршов индекс 3. Један од радова категорије M21a је у Web of Science био означен као Highly Cited Paper (top 1%) за период од објављивања у марту 2016. године до октобра 2017. године.

# Списак публикација др Владимира Лончара

## Поглавље у истакнутој монографији међународног значаја (M13)

1. **V. Lončar**, I. Vasić, and A. Balaž,  
Efficient numerical tools for solving the nonlinear Schrödinger equation,  
in *Scientific Computing: Studies and Applications*, Ed. C. Erling, pp. 63–157  
NOVA Science Publishers (2017), ISBN: 978-1-53612-564-1

## Радови у међународним часописима изузетних вредности (M21a)

1. **V. Lončar**, A. Balaž, A. Bogojević, S. Škrbić, P. Muruganandam, and S. K. Adhikari,  
CUDA programs for solving the time-dependent dipolar Gross-Pitaevskii equation in an anisotropic trap, *Comput. Phys. Commun.* **200**, 406 (2016) [ИФ(2016)=3.936]  
DOI: 10.1016/j.cpc.2015.11.014
2. **V. Lončar**, L. E. Young-S., S. Škrbić, P. Muruganandam, S. K. Adhikari, and A. Balaž,  
OpenMP, OpenMP/MPI, and CUDA/MPI C programs for solving the time-dependent dipolar Gross-Pitaevskii equation, *Comput. Phys. Commun.* **209**, 190 (2016) [ИФ(2016)=3.936]  
DOI: 10.1016/j.cpc.2016.07.029
3. L. E. Young-S., P. Muruganandam, S. K. Adhikari, **V. Lončar**, D. Vudragović and A. Balaž,  
OpenMP GNU and Intel Fortran programs for solving the time-dependent Gross-Pitaevskii equation, *Comput. Phys. Commun.* **220**, 503 (2017) [ИФ(2016)=3.936]  
DOI: 10.1016/j.cpc.2017.07.013

## Саопштења са међународних скупова штампана у целини (M33)

1. **V. Lončar**, S. Škrbić, and A. Balaž,  
Parallelization of minimum spanning tree algorithms using distributed memory architectures,  
in *Transactions on Engineering Technologies*, Eds. G.-C. Yang, S-I. Ao, L. Gelman, pp. 543–554  
Springer (2014), DOI: 10.1007/978-94-017-8832-8\_39
2. **V. Lončar**, S. Škrbić,  
Parallel implementation of minimum spanning tree algorithms using MPI,  
in *IEEE 13th International Symposium on Computational Intelligence and Informatics (CINTI)*,  
pp. 35–38 (2012), DOI: 10.1109/CINTI.2012.6496797

## Саопштења са међународних скупова штампана у изводу (M34)

1. **V. Lončar**, A. Balaž, A. Pelster,  
Trapped Bose-Einstein Condensates with Strong Disorder,  
V International School and Conference on Photonics - Photonica 2015, Belgrade, Serbia (2015)
2. **V. Lončar**, D. Vudragović, A. Balaž, A. Pelster,  
Rosensweig instability due to three-body interaction or quantum fluctuations?  
DPG 2016 conference, Q17.2, Hannover, Germany (2016)
3. **V. Lončar**, D. Vudragović, S. K. Adhikari, A. Balaž,  
Parallel solvers for dipolar Gross-Pitaevskii equation,  
VI International School and Conference on Photonics - Photonica 2017, Belgrade, Serbia (2017)

## Одбрањена докторска дисертација (M70)

1. **V. Lončar**,  
Hybrid parallel algorithms for solving nonlinear Schrödinger equation,  
Univerzitet u Novom Sadu, Prirodno-matematički fakultet (2017)

Web of Science InCites Journal Citation Reports Essential Science Indicators EndNote Publons Sign In Help English

# Web of Science

Clarivate Analytics

Search Search Results My Tools Search History Marked List

**Citation report for 5 results from Web of Science Core Collection between** 1996 and 2018 **Go**

You searched for: **AUTHOR: (loncar v\*) ...More**

This report reflects citations to source items indexed within Web of Science Core Collection. Perform a Cited Reference Search to include citations to items not indexed within Web of Science Core Collection.

Export Data: Save to Text File

**Total Publications**

**5**

1998 2017

**h-index**

**3**

Average citations per item

**9**

**Sum of Times Cited**

**45**

Without self citations


**43**

**Citing articles**

**32**

Without self citations

**30**



Sort by: Times Cited Date More

Page 1 of 1

Use the checkboxes to remove individual items from this Citation Report

or restrict to items published between 1996 and 2018 **Go**

- 1. **CUDA programs for solving the time-dependent dipolar Gross-Pitaevskii equation in an anisotropic trap**  
By: Loncar, Vladimir; Balaz, Antun; Boojevic, Aleksandar; et al.

2014	2015	2016	2017	2018	Total	Average Citations per Year
0	0	12	27	6	45	15.00
0	0	12	14	3	29	9.67

COMPUTER PHYSICS COMMUNICATIONS Volume: 200 Pages: 406-410  
Published: MAR 2016

2. **OpenMP, OpenMP/MPI, and CUDA/MPI C programs for solving the time-dependent dipolar Gross-Pitaevskii equation**

By: Loncar, Vladimir; Young-S, Luis E.; Skrbic, Srdjan; et al.  
COMPUTER PHYSICS COMMUNICATIONS Volume: 209 Pages: 190-196  
Published: DEC 2016

3. **OpenMP GNU and Intel Fortran programs for solving the time-dependent Gross-Pitaevskii equation**

By: Young-S, Luis E.; Muruganandam, Paulsamy; Adhikari, Sadhan K.; et al.  
COMPUTER PHYSICS COMMUNICATIONS Volume: 220 Pages: 503-506  
Published: NOV 2017

4. **Distributed Memory Parallel Algorithms for Minimum Spanning Trees**

By: Loncar, Vladimir; Skrbic, Srdjan; Balaz, Antun  
Edited by: Ao, Si; Gelman, L; Hukins, DWL; et al.  
Conference: World Congress on Engineering (WCE 2013) Location: Imperial Coll London, London, ENGLAND Date: JUL 03-05, 2013  
Sponsor(s): IAENG, Soc Artificial Intelligence; IAENG, Soc Bioinformat; IAENG, Soc Comp Sci; IAENG, Soc Data Mining; IAENG, Soc Elect Engn; IAENG, Soc Imaging Engn; IAENG, Soc Ind Engn; IAENG, Soc Informat Syst Engn; IAENG, Soc Internet Comp & Web Serv; IAENG, Soc Mech Engn; IAENG, Soc Operat Res; IAENG, Soc Sci Comp; IAENG, Soc Software Engn; IAENG, Soc Software Engn; IAENG, Soc Wireless Networks  
WORLD CONGRESS ON ENGINEERING - WCE 2013, VOL II Book Series: Lecture Notes in Engineering and Computer Science Pages: 1271-+  
Published: 2013

5. **Parallel implementation of minimum spanning tree algorithms using MPI**

By: Loncar, Vladimir; Skrbic, Srdjan  
Edited by: Szakal, A  
Conference: 13th IEEE International Symposium on Computational Intelligence and Informatics (CINTI) Location: Budapest, HUNGARY Date: NOV 20-22, 2012  
Sponsor(s): IEEE  
13TH IEEE INTERNATIONAL SYMPOSIUM ON COMPUTATIONAL INTELLIGENCE AND INFORMATICS (CINTI 2012) Book Series: International Symposium on Computational Intelligence and Informatics Pages: 35-38  
Published: 2012

0	0	0	10	2	12	4.00	
0	0	0	2	1	3	1.50	
0	0	0	1	0	1	0.17	
0	0	0	0	0	0	0.00	

Select Page

Sort by: **Times Cited** Date  Page 1 of 1

5 records matched your query of the 40,229,813 in the data limits you selected.



Копије објављених радова и других  
публикација

Computer Science, Technology and Applications

# Scientific Computing



Caj Erling  
Editor

## Studies and Applications

NOVA

**COMPUTER SCIENCE, TECHNOLOGY AND APPLICATIONS**

**SCIENTIFIC COMPUTING**  
**STUDIES AND APPLICATIONS**

No part of this digital document may be reproduced, stored in a retrieval system or transmitted in any form or by any means. The publisher has taken reasonable care in the preparation of this digital document, but makes no expressed or implied warranty of any kind and assumes no responsibility for any errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of information contained herein. This digital document is sold with the clear understanding that the publisher is not engaged in rendering legal, medical or any other professional services.

# **COMPUTER SCIENCE, TECHNOLOGY AND APPLICATIONS**

Additional books in this series can be found on Nova's website under the Series tab.

Additional e-books in this series can be found on Nova's website under the eBooks tab.



**COMPUTER SCIENCE, TECHNOLOGY AND APPLICATIONS**

**SCIENTIFIC COMPUTING  
STUDIES AND APPLICATIONS**

**CAJ ERLING  
EDITOR**



Copyright © 2017 by Nova Science Publishers, Inc.

**All rights reserved.** No part of this book may be reproduced, stored in a retrieval system or transmitted in any form or by any means: electronic, electrostatic, magnetic, tape, mechanical photocopying, recording or otherwise without the written permission of the Publisher.

We have partnered with Copyright Clearance Center to make it easy for you to obtain permissions to reuse content from this publication. Simply navigate to this publication's page on Nova's website and locate the "Get Permission" button below the title description. This button is linked directly to the title's permission page on copyright.com. Alternatively, you can visit copyright.com and search by title, ISBN, or ISSN.

For further questions about using the service on copyright.com, please contact:

Copyright Clearance Center

Phone: +1-(978) 750-8400

Fax: +1-(978) 750-4470

E-mail: [info@copyright.com](mailto:info@copyright.com).

### **NOTICE TO THE READER**

The Publisher has taken reasonable care in the preparation of this book, but makes no expressed or implied warranty of any kind and assumes no responsibility for any errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of information contained in this book. The Publisher shall not be liable for any special, consequential, or exemplary damages resulting, in whole or in part, from the readers' use of, or reliance upon, this material. Any parts of this book based on government reports are so indicated and copyright is claimed for those parts to the extent applicable to compilations of such works.

Independent verification should be sought for any data, advice or recommendations contained in this book. In addition, no responsibility is assumed by the publisher for any injury and/or damage to persons or property arising from any methods, products, instructions, ideas or otherwise contained in this publication.

This publication is designed to provide accurate and authoritative information with regard to the subject matter covered herein. It is sold with the clear understanding that the Publisher is not engaged in rendering legal or any other professional services. If legal or any other expert assistance is required, the services of a competent person should be sought. FROM A DECLARATION OF PARTICIPANTS JOINTLY ADOPTED BY A COMMITTEE OF THE AMERICAN BAR ASSOCIATION AND A COMMITTEE OF PUBLISHERS.

Additional color graphics may be available in the e-book version of this book.

### **Library of Congress Cataloging-in-Publication Data**

ISBN 978-1-53612-584-9 (eBook)

*Published by Nova Science Publishers, Inc. † New York*

# CONTENTS

<b>Preface</b>		<b>vii</b>
<b>Chapter 1</b>	Exploring the Potential for Scientific Computing on the Java Platform of Scala and Groovy with the Scalalab and Groovylab Environments <i>Stergios Papadimitriou and Lefteris Moussiades</i>	<b>1</b>
<b>Chapter 2</b>	The Structure and Dynamics of Meetup Social Networks <i>Marija Mitrović Dankulov and Jelena Smiljanić</i>	<b>33</b>
<b>Chapter 3</b>	Efficient Numerical Tools for Solving the Nonlinear Schrödinger Equation <i>Vladimir Lončar, Ivana Vasić and Antun Balaž</i>	<b>63</b>
<b>Index</b>		<b>159</b>

*Chapter 3*

# EFFICIENT NUMERICAL TOOLS FOR SOLVING THE NONLINEAR SCHRÖDINGER EQUATION

*Vladimir Lončar, Ivana Vasić and Antun Balazž\**

Scientific Computing Laboratory,  
Center for the Study of Complex Systems,  
Institute of Physics Belgrade, University of Belgrade,  
Belgrade, Serbia

## Abstract

The nonlinear Schrödinger equation or Gross-Pitaevskii equation plays an important role in several areas of physical sciences. It is used in various forms to describe prominent physical phenomena, such as propagation of light in a nonlinear medium and the dynamics of weakly interacting Bose-Einstein condensates studied in experiments with ultracold quantum gases. Large-scale numerical simulations for realistically large physical systems based on this equation have provided valuable insights into the properties of many interesting phenomena. Furthermore, numerical solutions of this equation are now routinely used for interpretation of experimental data and are indispensable for day-to-day operation of many laboratories, as well as for the design of new experimental setups.

In this chapter we review an efficient numerical algorithm for solving this nonlinear partial differential equation in three spatial coordinates

---

\*Corresponding Author Email: antun.balaz@ipb.ac.rs.



and either real or imaginary time, based on the Crank-Nicolson split-step semi-implicit method. Integration over imaginary time provides a direct access to stationary solutions of the full time-dependent problem. General spatially anisotropic solutions in three spatial variables are considered, as well as dimensionally-reduced problems such as radially or spherically symmetric cases.

Stability and efficiency of the presented method is ensured in the following way. The discretization error in both time and spatial step is of the second order. Moreover, integration over the time variable is split into two steps that are dealt with in a different, but optimized way. In the first step, only local part of the equation is taken into account and integration is performed exactly with respect to it. In this way the nonlinear term is considered in a straightforward, but optimized way that allows for a treatment of strong nonlinearities. In the second step, spatial derivatives are considered separately. Semi-implicit nature of the algorithm also contributes to its stability.

This method is numerically implemented using the standard C programming language. Further computational optimization is achieved by employing the state-of-the-art parallelization techniques for different computing platforms: OpenMP (for shared memory multiprocessing), OpenMP/MPI (for distributed memory systems, where each compute node uses OpenMP to maximize performance), CUDA (for single GPGPUs), CUDA/MPI (for distributed memory systems, where each compute node has GPGPU installed), as well as their hybrid combinations. The scalability of the algorithm with the number of processors and compute nodes is discussed.

As applications and extensions of the above developed algorithms, we present three relevant examples from the field of ultracold atoms. The first application illustrates how the non-equilibrium dynamics of multicomponent Bose-Einstein condensates can be numerically studied, including the excitation dynamics of a condensate with a coupling of the effective spin and angular momentum in the second application. The third application demonstrates how formation of vortices can be studied in a condensate featuring contact, as well as long-range dipole-dipole interaction.

**PACS:** 02.60.Lj, 02.60.Jh, 02.60.Cb, 03.75.-b

**Keywords:** Bose-Einstein condensate, Gross-Pitaevskii equation, Dipole-dipole interaction, OpenMP, MPI, CUDA, Crank-Nicolson method

## 1. Introduction

One of the important breakthroughs in physics of the twentieth century was the development of quantum mechanics. The key ingredient of this theory is the Schrödinger equation: a linear equation that governs the time evolution of a quantum system. In one of its common representations, Schrödinger equation is a partial differential equation.

Full quantum description of an interacting many-body system easily becomes computationally demanding due to exponentially increasing underlying vector space with the system size. For this reason, approximate descriptions that capture main physical phenomena are highly sought-after. The nonlinear Schrödinger (NLS) equation represents an effective, mean-field description of a quantum system. It is of great relevance in the field of cold atoms, in particular for a description and understanding of a weakly interacting bosonic systems at low temperatures in the regime of Bose–Einstein condensation [1]. In this context, the equation is widely known as Gross–Pitaevskii equation [2, 3] and it has provided valuable theoretical understanding of experimental results. Another field of research where the NLS equation plays a prominent role is the field of nonlinear optics that was initiated by the discovery of lasers in 1960 [4]. As analytical solutions of the NLS equation are limited to just a few special cases, the NLS equation was investigated numerically starting in the seventies using several different numerical methods [5]. The development of present-day computer architectures promotes numerical simulations and enables studies of more complex physical problems for experimentally relevant system sizes.

In this chapter we first introduce the NLS equation and derive an algorithm for obtaining its solutions based on a split-step method. In the next step we present an optimized implementation of this algorithm that heavily exploits parallelization techniques available on present-day state-of-the-art computer architectures. In the last sections of this chapter we illustrate applications of the described method for three examples that are currently in a research focus of cold-atom community.

## 2. Algorithm and Computational Implementation

### 2.1. Introduction

The time–dependent nonlinear Schrödinger equation in a dimensionless form is given by

$$i \frac{\partial \psi(\mathbf{r}, t)}{\partial t} = \left( -\frac{1}{2} \Delta + V(\mathbf{r}) + g |\psi(\mathbf{r}, t)|^2 \right) \psi(\mathbf{r}, t), \quad (1)$$

where  $\Delta$  is the three-dimensional Laplace operator  $\Delta \equiv \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}$ . The usual Schrödinger equation describing the time–evolution of a single–particle wave function  $\psi(\mathbf{r}, t)$  is recovered by eliminating the nonlinear term by setting  $g = 0$ . The wave function depends on four variables: three spatial coordinates  $\mathbf{r} = (x, y, z)$  and time  $t$ . The real–valued function  $V(\mathbf{r})$  corresponds to an external potential, and before mentioned Laplacian stems from the kinetic energy. Initial conditions are appropriately set by fixing  $\psi(\mathbf{r}, t = 0) = \psi_{\text{initial}}(\mathbf{r})$  and then the time evolution for  $t > 0$  is studied. The equation (1) conserves the normalization of the function  $\psi(\mathbf{r}, t)$

$$\int d\mathbf{r} |\psi(\mathbf{r}, t)|^2 = 1, \quad (2)$$

which in physical terms corresponds to the conservation of the particle number, as well as the total energy of the system,

$$E = \int d\mathbf{r} \psi^*(\mathbf{r}, t) \left( -\frac{1}{2} \Delta + V(\mathbf{r}) + \frac{g}{2} |\psi(\mathbf{r}, t)|^2 \right) \psi(\mathbf{r}, t). \quad (3)$$

In the context of cold atoms, the complex–valued function  $\psi(\mathbf{r}, t)$  stands for a condensate wave function – an effective single particle wave function macroscopically occupied by bosonic particles. The constant  $g$  is set by weak atomic interactions that we approximate by contact interactions: repulsive interactions yield  $g > 0$ , while attractive interactions are described by  $g < 0$ . The phase of a system described by the nonlinear Schrödinger equation is called a Bose–Einstein condensate (BEC). Macroscopic occupation of a single particle state (a finite condensate fraction) is a consequence of bosonic statistics (Bose–Einstein statistics) that becomes of key importance at very low temperatures. The phenomenon of Bose–Einstein condensation was predicted in early days of quantum mechanics in 1924 when the Bose–Einstein statistics was derived for the

first time. After several decades of experimental efforts, first experimental realization of a BEC with nanokelvin-cold atoms was reported in 1995 [6, 7, 8]. This success marked the onset of intense research activity in the field of cold atoms where the description based on the nonlinear Schrödinger equation played an important role [9, 10, 11, 12, 1, 13, 14]. Formally, the equation (1) can be derived by using the time-dependent variational principle. The time-dependent spatial density distribution of a condensate is given by

$$n(\mathbf{r}, t) = |\psi(\mathbf{r}, t)|^2. \quad (4)$$

Equation (1) governs the time evolution of a weakly interacting BEC. Another relevant question concerns finding a stationary solution  $\psi(\mathbf{r})$  that minimizes the energy

$$E_0 = \int d\mathbf{r} \psi^*(\mathbf{r}) \left( -\frac{1}{2}\Delta + V(\mathbf{r}) + \frac{g}{2}|\psi(\mathbf{r})|^2 \right) \psi(\mathbf{r}), \quad (5)$$

for a given potential  $V(\mathbf{r})$  and interaction constant  $g$ . For a moment we are considering a stationary problem, so  $\psi$  is a function of three spatial variables only. To address this question we will use a variant of a well established gradient descent method [15]. To this end, we introduce a label  $n = 0, 1, \dots$ , where  $\psi(\mathbf{r})_0$  is our initial guess for the stationary solution. In the next step, the function  $\psi$  is incrementally changed as

$$\psi(\mathbf{r})_{n+1} - \psi(\mathbf{r})_n = -\epsilon \frac{\delta E_0(\psi)}{\delta \psi^*}, \quad (6)$$

where  $\epsilon$  is an optimally chosen update step. The reasoning behind equation (6) is a general idea of the gradient descent method: we move through the space of functions  $\psi(\mathbf{r})$  along the direction in which the energy functional exhibits the fastest decay. For a stationary solution  $\psi_0(\mathbf{r})$  it holds true  $\frac{\delta E_0(\psi)}{\delta \psi^*} = 0$  and it corresponds to an extremum (minimum) of  $E_0$ . The last equation can be rewritten as

$$\frac{\partial \psi(\mathbf{r}, \tau)}{\partial \tau} = - \left( -\frac{1}{2}\Delta + V(\mathbf{r}) + g|\psi(\mathbf{r}, \tau)|^2 \right) \psi(\mathbf{r}, \tau). \quad (7)$$

where  $\tau = n\epsilon$  is a label that keeps track of propagation through the space of functions  $\psi(\mathbf{r})$ . We notice that equation (7) can be derived from equation (1) by a formal replacement  $t \rightarrow i\tau$ . For this reason, the variable  $\tau$  is named imaginary

time. In contrast to the energy conservation (3), imaginary time propagation by its construction ensures that with a right choice of initial function, we will reach a minimum of the energy functional (5) at long enough  $\tau$ . Moreover, imaginary time propagation as given in (7) does not preserve normalization (2), so this condition should be provided additionally, either by introducing an appropriate Lagrange multiplier or by re-normalizing  $\psi(\mathbf{r}, \tau)$  in each propagation step (6). Finally, we note that the solution  $\psi_0(\mathbf{r})$  which minimizes  $E_0$  evolves in time according to the nonlinear equation (1) as

$$\psi(\mathbf{r}, t) = e^{-i\mu t} \psi_0(\mathbf{r}), \quad (8)$$

where  $\mu$  is a chemical potential

$$\mu = \int d\mathbf{r} \psi_0^*(\mathbf{r}) \left( -\frac{1}{2}\Delta + V(\mathbf{r}) + g|\psi_0(\mathbf{r})|^2 \right) \psi_0(\mathbf{r}). \quad (9)$$

More recent experimental progress in the field of cold atoms has enabled study of more complex and fascinating physical systems whose theoretical understanding necessitates generalizations and extensions of the basic NLS equation (1). One of the topics that has attracted lot of attention recently is realization of a BEC of atoms with strong magnetic dipole moment [16, 17, 18]. To describe the system, we define  $z$ -axis as the polarization axis of the dipole moments. In this case the dipole-dipole interaction is proportional to [19]

$$V_{\text{dd}}(\mathbf{R}) = \frac{1 - 3 \cos^2 \theta}{|\mathbf{R}|^3}, \quad (10)$$

where vector  $\mathbf{R} = \mathbf{r} - \mathbf{r}'$  stands for the relative position of the two dipoles, and  $\theta$  is the angle between the vector  $\mathbf{R}$  and the  $z$ -axis. The dipolar NLS equation is given by

$$i \frac{\partial \psi(\mathbf{r}, t)}{\partial t} = \left( -\frac{1}{2}\Delta + V(\mathbf{r}) + g|\psi(\mathbf{r}, t)|^2 + g_{\text{dd}} \int d\mathbf{r}' V_{\text{dd}}^{3D}(\mathbf{r} - \mathbf{r}') |\psi(\mathbf{r}', t)|^2 \right) \psi(\mathbf{r}, t). \quad (11)$$

The coefficient  $g_{\text{dd}}$  sets the strength of the dipolar interaction. Other generalizations of the NLS equation (1) will be discussed in sections 3 (two-component BEC) and 4 (BEC with the coupling of spin and angular momentum).

In order to further illustrate wide applicability of the NLS equation, here we present its form used in nonlinear fiber optics. Starting with Maxwell's equations and using several justified approximations, it can be shown that the

propagation of the pulse field envelope  $A(x, y, t, z)$  along the  $z$  direction is governed by the NLS equation of the form [20, 4]

$$i \frac{\partial A(x, y, t, z)}{\partial z} = \left[ -\frac{1}{2} \left( \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \beta \frac{\partial^2}{\partial t^2} \right) + \gamma |A(x, y, t, z)|^2 \right] A(x, y, t, z), \quad (12)$$

where the coefficient  $\beta$  describes dispersive effects and the coefficient  $\gamma$  describes nonlinear effects. In the following, we mainly consider equations (1) and (11), but we note that all the discussed algorithms and their numerical implementations can be used to simulate equation (12) as well.

## 2.2. Algorithm

As the NLS equation (1) has been widely used and solved numerically, various types of numerical methods have been applied to this purpose [5]. A recent review paper [21] provides extensive comparison of the properties of different numerical methods used to solve the NLS equation. Broadly speaking, numerical methods for solving equation (1) are either finite difference approximations, pseudo spectral methods or finite element methods. Depending on the applied time–discretization scheme, a finite difference method can be explicit, implicit, semi–implicit or split–step approximation. Here we follow derivations from reference [22, 23] and present a simple and efficient numerical method to integrate the NLS equation (1). From the numerical point of view, real–time propagation according to equation (1) and imaginary–time propagation (7) can be implemented in a very similar manner.

In order to integrate equations (1) and (7) in time, we perform time discretization  $t \equiv t_0 + n\varepsilon$ ,  $n = 0, 1, 2, \dots$ ,  $\psi(\mathbf{r}, t) \rightarrow \psi(\mathbf{r})_n$ . It turns out that optimal ways to deal with spatially local and non–local terms of these equations are different, hence we further split integration over time into two consecutive pieces. This approximation formally corresponds to a replacement of equation (1) by a set of two equations

$$i \frac{\partial \psi(\mathbf{r}, t)}{\partial t} = (V(\mathbf{r}) + g|\psi(\mathbf{r}, t)|^2) \psi(\mathbf{r}, t) \quad (13)$$

$$i \frac{\partial \psi(\mathbf{r}, t)}{\partial t} = -\frac{1}{2} \Delta \psi(\mathbf{r}, t), \quad (14)$$

assuming that the result of integration of equation (13) is the initial condition for solving equation (14). The split-step approximation given by (13) and (14) is valid only for short-enough time interval  $\varepsilon$  and introduces a numerical error of the order  $O(\varepsilon^2)$ .

For practical reasons in a general three-dimensional situation, we further split equation (14) into three parts treating spatial derivatives over the three spatial coordinates separately. Thus, in this way we obtain a set of four consecutive equations

$$i \frac{\partial \psi(\mathbf{r}, t)}{\partial t} = H_i \psi(\mathbf{r}, t) \quad (15)$$

where  $i \in \{1, 2, 3, 4\}$  and we introduced

$$H_1 = V(\mathbf{r}) + g|\psi(\mathbf{r}, t)|^2, \quad (16)$$

$$H_2 = -\frac{1}{2} \frac{\partial^2}{\partial x^2}, \quad H_3 = -\frac{1}{2} \frac{\partial^2}{\partial y^2}, \quad H_4 = -\frac{1}{2} \frac{\partial^2}{\partial z^2}. \quad (17)$$

Accordingly, between configurations  $n$  and  $n + 1$ , we introduce three intermediate states labeled by  $n + \frac{1}{4}$ ,  $n + \frac{2}{4}$ ,  $n + \frac{3}{4}$  such that

$$\psi(\mathbf{r})_{n+\frac{1}{4}} = \exp \left[ -i\varepsilon \left( V(\mathbf{r}) + g|\psi(\mathbf{r})_n|^2 \right) \right] \psi(\mathbf{r})_n, \quad (18)$$

and

$$\psi(\mathbf{r})_{n+\frac{2}{4}} - \psi(\mathbf{r})_{n+\frac{1}{4}} = \frac{i\varepsilon}{4} \frac{\partial^2}{\partial x^2} \left[ \psi(\mathbf{r})_{n+\frac{2}{4}} + \psi(\mathbf{r})_{n+\frac{1}{4}} \right], \quad (19)$$

$$\psi(\mathbf{r})_{n+\frac{3}{4}} - \psi(\mathbf{r})_{n+\frac{2}{4}} = \frac{i\varepsilon}{4} \frac{\partial^2}{\partial y^2} \left[ \psi(\mathbf{r})_{n+\frac{3}{4}} + \psi(\mathbf{r})_{n+\frac{2}{4}} \right], \quad (20)$$

$$\psi(\mathbf{r})_{n+1} - \psi(\mathbf{r})_{n+\frac{3}{4}} = \frac{i\varepsilon}{4} \frac{\partial^2}{\partial z^2} \left[ \psi(\mathbf{r})_{n+1} + \psi(\mathbf{r})_{n+\frac{3}{4}} \right]. \quad (21)$$

The result (18) gives an exact, analytical solution of equation (13). In this way the nonlinear term has been properly taken into account providing a stability of the employed numerical scheme for arbitrary strong nonlinearity coefficient  $g$ . In the equations (19, 20, 21), we have used the semi-implicit Crank–Nicolson method [24], which is second-order accurate in the time step  $\varepsilon$  and unconditionally stable method. The accuracy of the approximation is  $O(\varepsilon^2)$ .

To reduce the obtained differential equations to the algebraic form, we additionally perform space discretizations with the discretization steps  $h_x$ ,  $h_y$ , and  $h_z$ . To this end, we introduce integer indices  $i, j, k$ , which take values  $0 \leq i \leq N_x$ ,  $0 \leq j \leq N_y$ ,  $0 \leq k \leq N_z$ , such that  $\psi(x_i, y_j, z_k)_n \equiv \psi_{i,j,k,n}$ , where  $x_i \equiv \left(-\frac{N_x}{2} + i\right) h_x$ , and similarly for  $y$  and  $z$ . We also approximate second-order spatial derivatives in the standard way via the central difference



formula:

$$\frac{\partial^2 \psi(\mathbf{r})_n}{\partial x^2} \rightarrow \frac{1}{h_x^2} (\psi_{i+1,j,k,n} - 2\psi_{i,j,k,n} + \psi_{i-1,j,k,n}), \quad (22)$$

with the error of the order of  $h_x^2$ . We consider large enough grids such that we can use boundary conditions  $\psi(0, j, k)_n \approx 0$ ,  $\psi(N_x, j, k)_n \approx 0$  and similarly in the other two spatial directions. As a result of using approximation (22) in equation (19), we obtain a tridiagonal system of equations

$$-A\psi_{i+1,j,k,n+\frac{2}{4}} + B\psi_{i,j,k,n+\frac{2}{4}} - A\psi_{i-1,j,k,n+\frac{2}{4}} = \delta_i, \quad (23)$$

for each  $j$  and  $k$ , where  $A = \frac{i\varepsilon}{4h_x^2}$ ,  $B = 1 + \frac{i\varepsilon}{2h_x^2}$  and

$$\delta_i = A\psi_{i+1,j,k,n+\frac{1}{4}} + B^*\psi_{i,j,k,n+\frac{1}{4}} + A\psi_{i-1,j,k,n+\frac{1}{4}}.$$

A solution of this tridiagonal system of equations can be cast in the form

$$\psi_{i+1,j,k,n+\frac{2}{4}} = \alpha_i \psi_{i,j,k,n+\frac{2}{4}} + \beta_i, \quad (24)$$

and from this ansatz we find the recursive relations for the solution:

$$\alpha_{i-1} = \frac{A}{B - A\alpha_i}, \quad \beta_{i-1} = \frac{\delta_i + A\beta_i}{B - A\alpha_i}. \quad (25)$$

From the boundary condition  $\psi_{N_x,j,k,n+\frac{2}{4}} = \psi_{N_x,j,k,n+\frac{1}{4}} = 0$ , we derive initial values for  $\alpha$  and  $\beta$ ,  $\alpha_{N_x-1} = \beta_{N_x-1} = 0$ , which we use to solve the recursive equations (25). With another boundary condition  $\psi_{0,j,k,n+\frac{2}{4}} = 0$ , we finally solve equation (24). Similar procedure is afterwards followed in  $y$  and  $z$  direction. Due to the tridiagonal form of the above system of equations, the complexity of the overall algorithm is proportional to the number of discretization points, i.e.,  $O(N_x N_y N_z)$ .

Once the ground-state solution  $\psi_0(\mathbf{r})$  is calculated using the imaginary-time version of the above algorithm, the corresponding value of the ground-state energy  $E_0$  of the system can be calculated by numerically integrating expression (5). The described algorithm is implemented using the Fortran programming language [22] and the C programming language [23]. Further computational optimization is achieved by employing the state-of-the-art parallelization techniques for different computing platforms: OpenMP (for shared memory multiprocessing) [25] and hybrid OpenMP/MPI (for distributed memory systems, where each compute node uses OpenMP to maximize performance) [26].

In order to solve the dipolar NLS equation (11) numerically, we need to deal with the integral on the right-hand side of this equation efficiently. To this end, we exploit the fact that the Fourier transform of a convolution integral is given by a product of Fourier transforms of its components

$$I_{\text{dd}}(\mathbf{r}, t) = \int d\mathbf{r}' V_{\text{dd}}(\mathbf{r} - \mathbf{r}') n(\mathbf{r}', t) = \int \frac{d\mathbf{k}}{(2\pi)^3} e^{-i\mathbf{k}\mathbf{r}} \tilde{V}_{\text{dd}}(\mathbf{k}) \tilde{n}(\mathbf{k}, t), \quad (26)$$

where we define the direct and the inverse Fourier transformation as:

$$\tilde{A}(\mathbf{k}) = \int d\mathbf{r} A(\mathbf{r}) e^{i\mathbf{k}\mathbf{r}}, \quad A(\mathbf{r}) = \frac{1}{(2\pi)^3} \int d\mathbf{k} \tilde{A}(\mathbf{k}) e^{-i\mathbf{k}\mathbf{r}}. \quad (27)$$

The Fourier transform of the geometric factor (10) stemming from the dipolar interaction can be obtained analytically and is given by

$$\tilde{V}_{\text{dd}}(\mathbf{k}) = \frac{4\pi}{3} \left( \frac{3k_z^2}{\mathbf{k}^2} - 1 \right). \quad (28)$$

The Fourier transform of  $n(\mathbf{r}', t)$  and the inverse Fourier transform from equation (26) are calculated numerically using optimized libraries such as FFTW [27]. Finally, the obtained numerical value  $I_{\text{dd}}(\mathbf{r}, t)$  of the integral (26) is used in the first part of the split-step approach (18) as:

$$\psi(\mathbf{r})_{n+\frac{1}{4}} = \exp \left[ -i\varepsilon \left( V(\mathbf{r}) + g|\psi(\mathbf{r})_n|^2 + g_{\text{dd}} I_{\text{dd}}(\mathbf{r})_n \right) \right] \psi(\mathbf{r})_n. \quad (29)$$

Numerical codes implementing the described algorithm for solving the time-dependent dipolar NLS equation are written in the Fortran and in the C programming languages [28]. CUDA parallel version of the algorithm is developed for single GPGPUs [29], as well as a hybrid CUDA/MPI version for distributed memory systems, where each compute node has GPGPU installed [30].

### 2.3. Parallelization of the Algorithm

Let us note that the tridiagonal systems of equations (23) can be solved separately for different values of  $j$  and  $k$ . This feature is essential and can be directly exploited to parallelize numerical implementation of the above described algorithm. In fact, this has been done for practically all available modern computer architectures and platforms, and we briefly outline all implementations.

Initial algorithm for solving the NLS equation with the contact interaction term using Crank-Nicolson method was published in reference [22], which also




Figure 1. Example of 1D, 2D and 3D meshes with  $N_x$ ,  $N_y$  and  $N_z$  discretization points in  $x$ ,  $y$  and  $z$  direction, respectively, and the corresponding spacings  $dx$ ,  $dy$  and  $dz$ .

includes a serial implementation in Fortran. This was followed by the implementation written in C [23], which includes serial and parallel implementation for shared memory systems consisting of one or more multi-core CPUs using OpenMP approach. Fortran OpenMP implementation was also provided [25], as well as a full parallel MPI version for distributed memory systems, i.e., computer clusters with multi-core CPU nodes [26]. An extension of the algorithm to include the dipolar interaction term was published in reference [28], where the authors provided two serial implementations, in C and Fortran. This was followed by an implementation written in Nvidia CUDA for hardware accelerators in the form of GPUs [29], and later by a comprehensive suite of programs that includes OpenMP, MPI, CUDA, as well as their hybrid combinations [30].

To introduce the parallelization strategy, we start with the description of the discretization scheme used by the algorithm. Let  $N_x$ ,  $N_y$  and  $N_z$  be the number of points in each direction, corresponding to the  $x$ ,  $y$  and  $z$  directions in one, two or three dimensions (1D, 2D and 3D, respectively). In other words, the number of points in 1D is defined by  $N_x$ , in 2D by  $N_x$  and  $N_y$ , and in 3D by  $N_x$ ,  $N_y$  and  $N_z$ . Discretization points in each direction are equidistant, with their spacing defined in dimensionless units by variables  $dx$ ,  $dy$  and  $dz$ , illustrated in Figure 1.

The mesh determines how the space is discretized, but the actual values of




Figure 2. Main loop of the real-time propagation algorithm in 3D.

the wave function, which is also discretized by the mesh, need to be stored in a separate array. For convenience, we refer to this array as `psi` variable. This variable is defined as a vector of size  $N_x$  in 1D, as  $N_x \times N_y$  matrix in 2D and as  $N_x \times N_y \times N_z$  tensor in 3D. In real-time propagation, the wave function values are complex, while in imaginary-time propagation they are purely real, which determines the type of `psi` variable.

Time is discretized according to  $t_n = t_0 + n\varepsilon$ , with the time step  $\varepsilon$  stored in a variable `dt`. Taking into account the time step and the total amount of physical time we wish to simulate, we determine the number of iterations required. The main part of the algorithm is a loop where each iteration corresponds to one time step of the propagation. Inside the loop, we update the wave function by propagating it in time with respect to different parts of the Hamiltonian, according to the split-step scheme described in section 2.2 This is illustrated in Figure 2.

In each step of the main loop, wave function values are computed along the discretization mesh in several smaller steps. First such substep is the calculation of the dipolar interaction term (for programs where it is taken into account, otherwise it is omitted). We have two further substeps in 1D, pertaining to propagation w.r.t.  $H_1$  and  $H_2$ . In 2D and 3D we follow a similar procedure and add further substeps for  $H_3$  and  $H_4$ . Each substep produces intermediate values of the wave function that are used in the next substep. The order of the substeps dealing with parts of the Hamiltonian with spatial derivatives ( $H_2$ ,  $H_3$  and  $H_4$ ) can be arbitrary, a feature useful in hybrid and distributed memory algorithms. Propagation of the wave function w.r.t.  $H_1$  relies on the availability of the dipolar interaction term value, which is computed in the first substep, before we start with the update of the wave function.

When necessary, the dipolar interaction term is computed using the discrete Fourier transform (DFT). DFT of a sequence is commonly implemented by relying on the *Fast Fourier transform* (FFT) algorithm via an external library. From equation (26) it follows that we first need a DFT of a sequence consisting




Figure 3. Main loop of the imaginary-time propagation algorithm in 3D, which introduces an additional step of normalization.

of absolute squares of wave function values. The resulting complex-valued sequence is then multiplied by the Fourier transform of the dipolar potential. This transform is a known function (28), precomputed for a chosen spatial mesh and stored in an array. Next, by performing the inverse DFT we get the dipolar interaction term stored in the resulting array, which we then use to propagate the wave function w.r.t.  $H_1$ . In 2D and 3D, we need multidimensional DFT, which can be computed by the composition of a sequence of 1D DFTs along each direction.

The substep in which the wave function is propagated w.r.t. the Hamiltonian part without spatial derivatives ( $H_1$ ) proceeds by employing further nested loops over mesh points. Part of the  $H_1$  is the trap potential  $V(\mathbf{r})$ , which, like Fourier transform of the dipolar potential, can be computed only once during the initialization, and then reused. We assume that the trap potential is static, but in case of a time-varying trap potential, the algorithm can be modified to update the trap potential before this substep, i.e., at the beginning of each iteration. The remaining substep, or substeps in 2D and 3D, propagates the wave function by relying on the Crank-Nicolson scheme. We need to perform a backward sweep of the mesh in each direction to determine the corresponding Crank-Nicolson coefficients  $\alpha$ ,  $\beta$  and  $\delta$  for each direction, followed by a forward sweep to determine the solution for the entire space range. The algorithm for imaginary-time propagation is similar to the one discussed above, with the addition of the normalization step at the end of each iteration of the main loop, as illustrated in Figure 3. To normalize the wave function, we first need to compute its norm and then to divide the wave function values with that norm.

We see that the algorithm is computationally very demanding, as it requires multiple passes over the entire mesh in each iteration of the main loop. In 2D and 3D, the total number of discretized points sharply increases, making this problem even greater. As we have identified above, several variables can be moved out of the main loop and computed only once, however the remaining calculations are still very demanding, hence the need for parallelization.

The described substeps cannot be executed concurrently, as the intermediate values they produce are used in subsequent substeps. Instead, parallelization of the algorithm can only be achieved by focusing on the mesh loops inside substeps. This is an obvious choice for shared memory systems, as all data stored in the memory is local and accessible to all participating processes, removing the need for data transfers between processes. The loops themselves can easily be parallelized if they do not contain recursive relations. This is true for the substeps dealing with the computation of the dipolar term and propagation w.r.t.  $H_1$ . Substep involving  $H_2$  cannot be easily parallelized in 1D, as it has recursive relations in both backward and forward sweeps. In general, recursive relations can be parallelized using the *scan* algorithm [31] (also known as the generalization of *prefix sum* algorithm [32]), however the implementation complexity of such an algorithm made us discard this approach. In 2D and 3D, we can exploit the fact that the recursive relations appear only in the innermost loops of the substeps involving  $H_2$ ,  $H_3$  and  $H_4$ . Thus, we can achieve parallelization by dividing the work among the processes at the level of the outermost loop.

In addition to the above outlined parallelization that can be naturally implemented on a single computer using OpenMP approach, one can also utilize modern GPUs to parallelize the algorithm on a single node, and even combine CPU and GPU into a hybrid algorithm. We first discuss CUDA-based parallelization, introduced in reference [29]. The significant differences in architecture of the GPU and its accompanying execution model on one side, and CPU architecture with its conventional execution model on the other side, do not result in significant changes to the main algorithm for the case of Nvidia CUDA implementation. From the algorithm's point of view, both CPU and GPU are shared memory systems with multiple processing units, even though they are vastly different, and therefore can be used in a similar manner. Main loop of the algorithm remains unchanged, as well as each substep. However, one has to replace the FFTW library used by C/OpenMP programs with the cuFFT library, which is available in CUDA. cuFFT provides two interfaces: the native one and the FFTW one [33]. FFTW interface is intended to be used as a drop-in replacement for FFTW, allowing programs written primarily with FFTW in mind to use CUDA GPUs with minimal modifications to the source code. While this interface could be used as a temporary solution during the development of CUDA programs, the goal of using CUDA for all computation relating to the propagation of the wave function could only be reached with the native cuFFT

interface. This interface is modeled after FFTW but differs from it in the way plans are created and executed. In cuFFT, precision and data type of the transform is determined by the functions initiating execution of plans, which results in a slightly different sequence of function calls and in improved performance.

In order to use CPU and GPU at the same time for hybrid algorithms, the work has to be divided efficiently between them. If we put aside the obvious hardware architecture difference between CPU and GPU, we can consider the computer with GPU as being similar to the distributed memory system consisting of two computing nodes with different characteristics. Both CPU and GPU have their own memory, and are connected through a fast interconnect, in this case the PCI-Express. This means that we need an approach similar to the one distributed memory algorithms use to achieve the desired parallelization. The most important difference here is that all data is available in main (CPU) memory, and only portions of it need to be transferred to GPU memory, as opposed to the true distributed memory systems, where no single part of the system has all the data. Therefore, the general approach to perform this type of hybrid computation is to:

1. designate a portion of data to be processed by GPU,
2. copy that data to GPU while simultaneously using CPU for computation over the remaining data,
3. transfer the data back from GPU, overwriting old data, and
4. synchronize CPU and GPU.

This flow is illustrated in Figure 4.

Note that data can be offloaded from CPU to GPU in various ways [34, 35] and, in general, any data distribution scheme which targets distributed memory systems may be used. However, we settled on a simple approach using 1D decomposition, also known as *slab decomposition*. When using slab decomposition, data is distributed along one dimension, usually the slowest-changing one, and the remaining dimension (in 2D programs) or dimensions (in 3D programs) of data remain local. This means that we can perform computation on the local data efficiently, and no data exchanges are necessary. Depending on a data access pattern, in our programs we rely on decompositions either along the  $x$  direction (discretized with the  $N_x$  spatial points), or along the  $y$  direction (discretized with the  $N_y$  spatial points). In 3D programs, there is no need

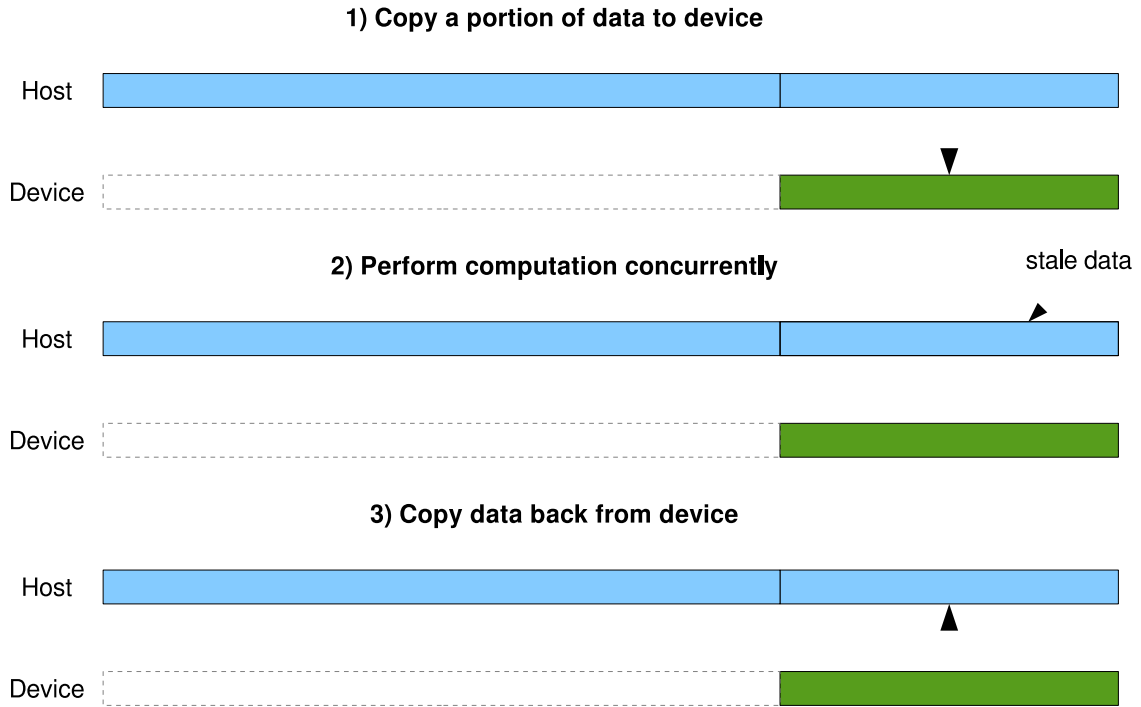


Figure 4. Flow of data between host and device.

to decompose the data along the  $z$  direction. In case we need data from the distributed dimension to become local in order to perform some computation (e.g., to update the wave function values when data are decomposed along the  $x$  direction), we have to reassemble data in host memory and decompose again along the appropriate dimension. Figure 5 illustrates this concept.

With the decomposition scheme in place, we can consider how to divide the computation of the wave function propagation in time between CPU and GPU. Once the data have been distributed between CPU and GPU, on the CPU side only the exit condition of the outer loop needs to be adjusted, so that the CPU processes a smaller number of elements (e.g., by replacing  $N_x$  with  $\text{cpu}N_x$ ). Similar changes are required on the GPU side as well, however we also need to ensure that GPU has valid data to work with, and that data are reassembled in host memory after the computation on GPU is done.

Therefore, in the hybrid implementation we initially distribute the data along the slowest-changing dimension. This allows CPU and GPU to simultaneously perform computation of wave function propagation w.r.t.  $H_1$  and  $H_3$  parts of the Hamiltonian (and  $H_4$  when working in 3D). The propagation of the wave function w.r.t.  $H_2$ , corresponding to the  $x$  direction, can be done locally



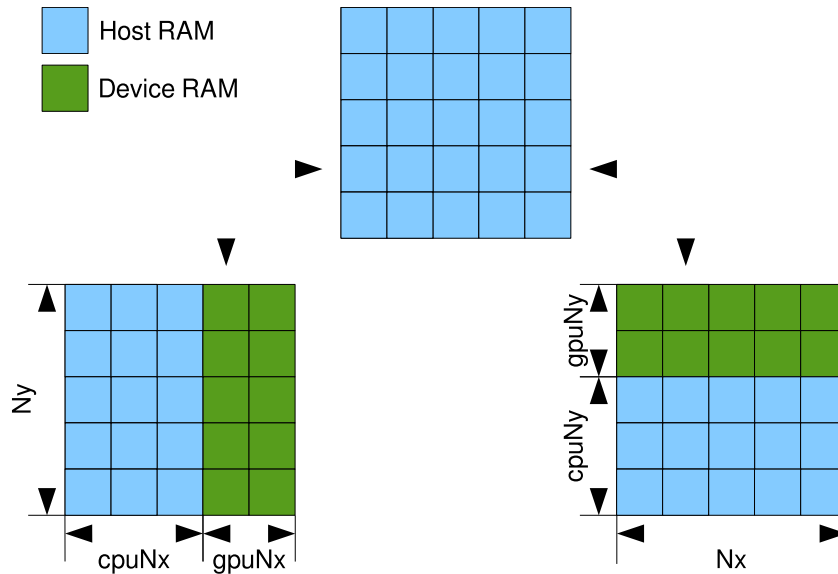


Figure 5. Two offload patterns used. If offloaded along  $x$  direction, the  $y$  direction remains local. If we need to access whole  $x$  direction data, we need to decompose along  $y$  direction.

only if we decompose the data along the  $y$  direction beforehand. Note that we do not have to transfer data from GPU back to host memory at the end of each function, due to the fact that time propagation w.r.t.  $H_2$ ,  $H_3$  and  $H_4$  parts of the Hamiltonian can be done in arbitrary order in each step. Without this, the time propagation workflow of 3D programs would have to involve the following steps, according to Figure 2:

1. calculate the dipolar term,
2. transfer the data decomposed along  $x$  direction to GPU,
3. propagate the wave function w.r.t.  $H_1$
4. reassemble the data in host memory,
5. transfer the data decomposed along  $y$  direction to GPU,
6. propagate the wave function w.r.t.  $H_2$ ,
7. reassemble the data in host memory,
8. transfer the data decomposed along  $x$  direction to GPU,

9. propagate the wave function w.r.t.  $H_3$ ,
10. propagate the wave function w.r.t.  $H_4$ ,
11. reassemble the data in host memory.

Rearranging the order of time-propagation substeps allows us to remove one transfer of data to GPU and its subsequent reassembly in host memory:

1. calculate the dipolar term,
2. transfer the data decomposed along  $x$  direction to GPU,
3. propagate the wave function w.r.t.  $H_1$ ,
4. propagate the wave function w.r.t.  $H_3$ ,
5. propagate the wave function w.r.t.  $H_4$ ,
6. reassemble the data in host memory,
7. transfer the data decomposed along  $y$  direction to GPU,
8. propagate the wave function w.r.t.  $H_2$ ,
9. reassemble the data in host memory.

Therefore, we use the above, optimized sequence in all 3D programs, and similarly in 2D.

To complete the hybrid algorithm, we also need to distribute computation of the dipolar term between CPU and GPU, where complexity arises in performing DFT on distributed data. Currently available FFT libraries target either CPU or GPU for their computation, but unfortunately not both at the same time. There are numerous attempts to develop specialized FFT libraries which would enable this [36, 37, 38], however a full-featured library with support for R2C transforms with advanced data layout is still not available. The approach we used is to rely on existing libraries for actual transforms, but perform data distributions manually. The libraries we use, FFTW on CPU and cuFFT on GPU, support advanced data layouts as well as working on a subset of the whole data, allowing for an efficient implementation.

To perform the Fourier transform simultaneously on CPU host and GPU device, we have to split the single multidimensional transform into a series of

1D transforms along each dimension of the input data, which can be computed on CPU and GPU independently. This approach is known as the *row-column* algorithm [39], and is often used in FFT libraries. The essence of this algorithm can best be summarized in an example of a 2D FFT. Given a matrix  $N_x \times N_y$ , we compute the DFT in the following way:

1. Transfer portion of the input array, decomposed along the  $x$  direction, to device memory. We transfer last  $g_{\text{pu}}N_x \times N_y$  consecutive array elements to GPU. This can be done as a single memory copy operation due to the flat allocation that was used. The choice whether to offload data to GPU memory from the beginning or the end of the input array is arbitrary.
2. Perform DFT along the  $y$  direction on both CPU and GPU concurrently. CPU will perform  $c_{\text{pu}}N_x$  such transformations, while the GPU will perform  $g_{\text{pu}}N_x$  1D DFTs. Each of these transforms will take a subset of the input array, which are  $N_y/2 + 1$  elements apart. Note that CPU will not do anything with the last  $g_{\text{pu}}N_x \times N_y$  elements.
3. Copy the array with the transform back from GPU to CPU, writing over the stale data with the relevant portion transformed on GPU. After this step, we have the complete array transformed along the  $y$  direction residing in host memory.
4. Transfer portion of the input array, decomposed along the  $y$  direction, to GPU memory. Similarly to step 1, we transfer last  $N_x \times g_{\text{pu}}N_y$  elements to GPU.
5. Perform DFT along the  $x$  direction on both CPU and GPU. This time, CPU will perform  $c_{\text{pu}}N_y/2 + 1$  such transformations, while GPU will perform remaining  $g_{\text{pu}}N_y/2$  transformations. The halving of the number of transformations is due to use of R2C transformations. In each transform, elements are  $N_y/2 + 1$  places apart, while the first element of each transform is adjacent to the previous one.
6. Copy the array with the transform back from GPU to CPU. With this step completed, we have the full FFT of the input array residing in host memory. This step may be omitted if the computation that follows does not require the FFT of input data to be fully assembled in host memory. This is the case in our algorithm, as we use the resulting transformed array

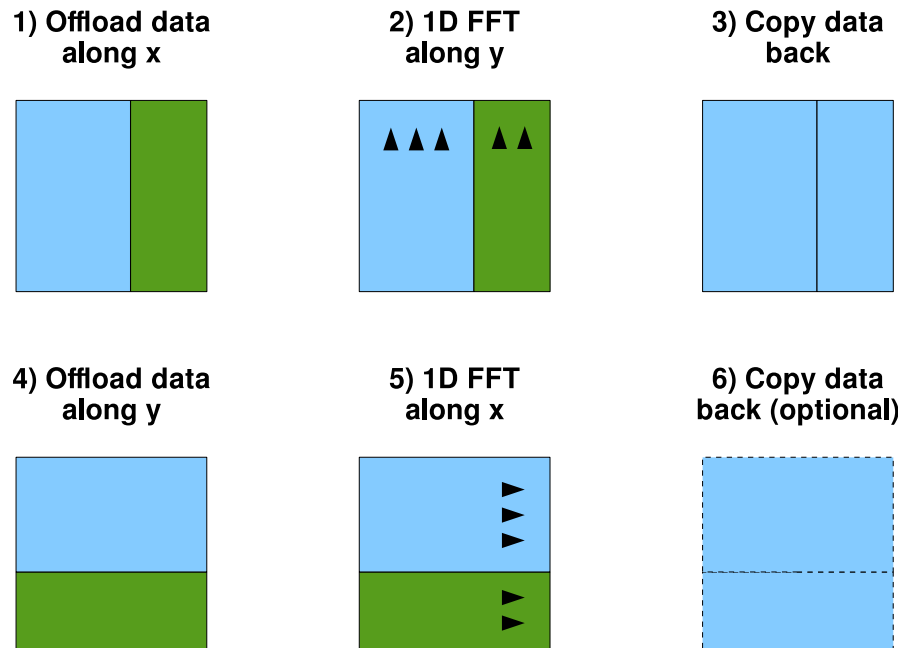


Figure 6. Hybrid algorithm for concurrent FFT on host and device.

only for subsequent computation on GPU. Since the memory transfer is expensive, omitting this step leads to significant performance improvement.

The inverse Fourier transform can be done in an analogous way. In 3D, the general principle remains the same, however, we do not separate the DFT into three 1D transforms, but to one 2D and one 1D transform. We do this due to better performance of 2D transforms, which in both libraries used is found to be better than two 1D transforms. Figure 6 illustrates the DFT algorithm described above.

And finally let us mention that several distributed memory implementations, targeting computer clusters, were developed as well:

1. pure CPU version, which is built on top of the shared memory algorithm (OpenMP/MPI),
2. pure GPU version, which is built on top of the CUDA implementation of the shared memory algorithm (CUDA/MPI),
3. hybrid version, which is based on the algorithm described in Chapter (Hybrid/MPI).

Two of the three implementations, the OpenMP/MPI and CUDA/MPI implementations, have been published in reference [30] and are available for download [40, 41]. The Hybrid/MPI implementation is also publicly available [41]. The main challenges in the development of distributed memory and hybrid algorithms were devising an efficient data distribution scheme that facilitates fast data shuffling between computing resources. In the MPI implementations this was done via a transpose operation realized in two ways, one relying on FFTW and the other relying on MPI vector types and collective communications. Further details can be found in reference [42].

## **2.4. Performance of Parallel Algorithms: Measurement Results and Modeling**

In previous section we presented various parallel implementations of programs solving the dipolar and contact interaction GP equation without providing details of their measured performance. This section is dedicated to filling that gap. We tested various scaling scenarios, strong and weak, from a single computing node to the cluster. All testing was done at the PARADOX supercomputing facility located at the Scientific Computing Laboratory, Center for the Study of Complex Systems of the Institute of Physics Belgrade.

Performance of hybrid implementations is highly dependent on the amount of work offloaded to GPU(s), so the key to maximizing performance of these versions is proper selection of parameters which control this feature. We used evolutionary computation techniques to find the optimal solutions, as described here. In the remaining subsections we give an overview of the tests and the methodology they relied on, present the performance results and their modeling, and finally discuss how to select the optimal algorithm for a given hardware platform.

### **2.4.1. Optimization of Input Parameters for Hybrid Implementations**

Hybrid implementations, running on a single computer or a cluster with multiple computing nodes, can potentially offer the best performance of all algorithms presented in this thesis, by utilizing all allocated computing resources. For this to happen, work must be divided between CPU and GPU in such a way as to maximize their throughput and minimize their idle time. A question that naturally arises is how to achieve this for any possible mesh size for a given CPU/GPU combination. Unfortunately, there is no single best way to divide the

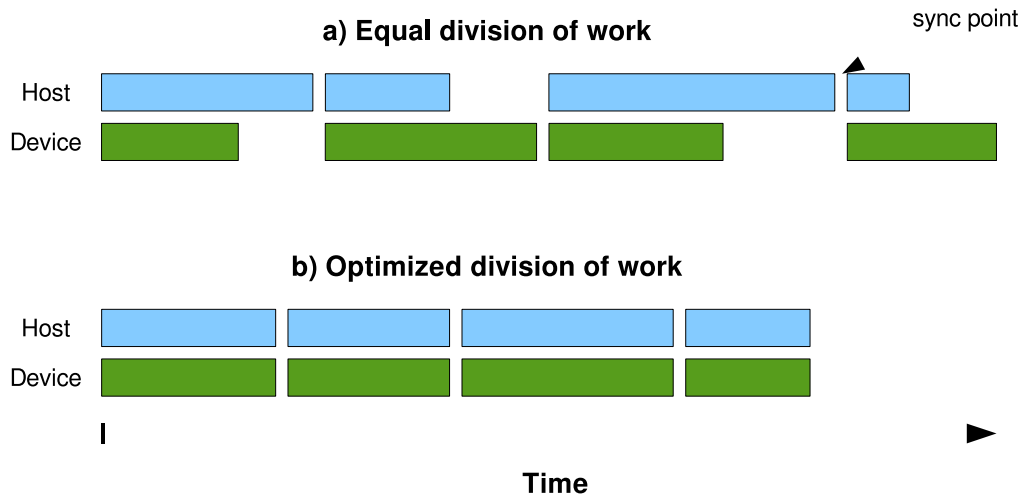


Figure 7. Execution timeline in a hybrid system: (a) with equal data distribution, leading to unbalanced computation time; (b) with optimized data distribution and ideal computation load.

work, since processing powers of CPUs and GPUs vary significantly. For example, in a computer where a powerful new GPU is paired with an older type of CPU, more work should be offloaded to the GPU, and vice versa for computers with a powerful CPU and low-performing GPU. Even if the installed CPU and GPU offer similar performance in terms of floating point operations per second (as is the case for the PARADOX cluster), the differing architectures mean that certain portions of the algorithm are better suited to CPU or GPU. If we were to divide the work naively, e.g., equally, this would lead to unbalanced computation, as illustrated in Figure 7(a), which in turn increases the computation time.

To get over this problem, the described hybrid implementations have an option to manually specify each parameter that controls the work being done on a GPU. Ideally, this flexibility would allow us to overlap computation on CPU and GPU as much as possible, thus minimizing the execution and idle time of each resource. The execution timeline for the ideal case looks like the illustration in Figure 7(b).

The aforementioned parameters include the total amount of data transferred to GPU, the number of chunks, kernel grid size parameters and a special parameter controlling whether some functions will even be offloaded to GPU or not, which is useful in situations where the CPU is much more powerful than the

GPU. In 3D programs, this amounts to 33 integer-valued parameters for which we would like to find the optimal values for the desired mesh size on a given computing system (with given hardware characteristics). Furthermore, there are constraints that the parameters must satisfy. For example, the total amount of data transferred to GPU must be divisible by the number of chunks, and the size of a block that kernel launches must fit within the limits imposed by the underlying GPU.

To evaluate any combination of parameters, the programs have to be executed, and their execution time measured. The search space of 33 parameters is clearly too large to be exhaustively traversed in a reasonable amount of time. Fortunately, we can reduce the number of parameters we have to search through at any single time by grouping them based on the synchronization point. After the CPU and GPU synchronize, data have to be divided again, so the only relevant parameters in the region between two synchronization points are the ones controlling the division of work and the ones controlling the kernels which are executed in that region. During a single iteration, CPU and GPU have to synchronize four times:

1. after the FFT on local dimensions,
2. after the inverse FFT on the non-local dimension,
3. before the function performing propagation w.r.t.  $H_2$ , and
4. after that function (which represents the end of a single iteration).

This gives us four distinct parameter sets to optimize:  $PS_1$  with 6 parameters,  $PS_2$  also with 6 parameters,  $PS_3$  with 15 parameters and  $PS_4$  with 5 parameters. We can consider each parameter set independently.

We investigated several approaches to optimization of parameter sets: naive brute-force search, iterative optimization via gradient descent, and metaheuristic via genetic algorithm. Our goal was not to find the best possible optimization algorithm for our problem, but rather to implement a reasonably good one. By this we mean the algorithm which is not difficult to implement, which will give us a set of parameters that are close to the optimal ones, and will not take a long time to finish. Our focus was on 3D variants of both single node and MPI versions of hybrid programs. The concepts presented here apply equally to 2D variants and there is no important difference between the single node and

the MPI version, so we will not make this distinction in the remainder of this section. We will now describe how the three approaches were implemented.

A brute-force search (BFS), or exhaustive search, considers every possible combination of parameter values in order to find the optimal one. Even with parameters divided into smaller sets, the exhaustive search takes too long, and is only feasible if we narrow down the ranges of all parameters involved. Unfortunately, this is only possible when we have a deep understanding of the performance of the CPU and GPU models used, when we can provide a reasonable guess for the optimal solution. An alternative is to consider only certain values along the range of a single parameter, rather than every value in the range. This approach is feasible and we were able to find reasonably good parameter values using it. The downside is that this way we may miss the optimal values of parameters, as they often lie between the selected test points. For instance, on one PARADOX computing node and a mesh size  $256 \times 256 \times 256$ , with the work divided between CPU and GPU along the outermost,  $x$  dimension, if we test only for the values of offloaded data to GPU that are multiples of 32 (i.e.,  $32 \times 256 \times 256$ ,  $64 \times 256 \times 256$ , etc.), we will miss the optimal value which lies around  $140 \times 256 \times 256$  (assuming all other parameters are fixed). Even though BFS is often not feasible, it is useful as the baseline for the evaluation of other methods we implemented.

With the BFS implemented as a baseline, we switched our focus to the implementation of an iterative optimization algorithm. A well-known and widely used algorithm is *gradient descent* (GD) [15]. It is an iterative procedure in which every iteration aims to get closer to the minimum of the given function  $F(\mathbf{x})$ , where  $\mathbf{x}$  represents a vector of function parameters. If the function  $F(\mathbf{x})$  is defined and differentiable around some point  $\mathbf{a}$ , then  $F(\mathbf{x})$  decreases fastest in the direction of the negative gradient of  $F$  at  $\mathbf{a}$ . GD exploits this fact and defines  $\mathbf{b}$  as

$$\mathbf{b} = \mathbf{a} - \gamma \nabla F(\mathbf{a}), \quad (30)$$

so that  $F(\mathbf{b}) \leq F(\mathbf{a})$ , for  $\gamma$  small enough. This observation is used as a building block for the algorithm. We start from an initial guess  $\mathbf{x}_0$  for a location of the minimum of  $F$  and construct a sequence  $\mathbf{x}_n (n \geq 1)$  such that:

$$\mathbf{x}_n = \mathbf{x}_{n-1} - \gamma \nabla F(\mathbf{x}_{n-1}). \quad (31)$$

This sequence will eventually converge to the local minimum, if one exists. GD is most applicable to functions for which the gradient can be computed analytically, however it can also be used when the derivative can only be obtained



numerically. To apply it here, the execution time of our programs will be the function  $F$  that we minimize, and its arguments are the  $n = 33$  parameters that control the GPU. The execution time is measured as the average time of a single iteration of the main time-propagation loop, sampled over a given number of iterations. As explained earlier, we can divide the problem into four separate minimization problems, with four sets of parameters of dimensionality  $m_i$  ( $i = 1, 2, 3, 4$ ), in the same way as in BFS. To numerically compute the  $k$ -th partial derivative of one of the minimization functions at a set of candidate parameters  $a_1, a_2, \dots, a_m$ , we use first order approximation:

$$\frac{\partial F}{\partial x_k}(a_1, \dots, a_m) \approx \frac{F(a_1, \dots, a_k + h_k, \dots, a_m) - F(a_1, \dots, a_k, \dots, a_m)}{h_k}, \quad (32)$$

where  $h_k$  is the increment of a given parameter. Since the parameters are integer-valued, we have to carefully choose the value of  $h_k$ , to be as small as possible, while still satisfying all the constraints that the corresponding parameter may have w.r.t. other parameters. We note that in order to evaluate the full gradient of  $F$  at  $(a_1, \dots, a_m)$  we need to execute our programs  $m + 1$  times.

We stress that execution times of our programs are not always the same due to the hardware, software and OS scheduling issues, making the minimization functions noisy. This noise affects the calculation of the gradient and could point the algorithm in the wrong direction. We can detect this by checking if the value of minimization function has increased in subsequent GD iteration, and discard this move if necessary. While useful, resorting to this tactic means that in case the GD gets stuck in a local minimum, it will not be able to get out and thus will never reach the global minimum. A naive way to check if this is the case would be to start over from a different initial point, and see if the GD algorithm converges to the same minimum. There are more sophisticated approaches to addressing this issue, which we discuss later.

In general, it is not possible to completely eliminate noise, due to inherent problems of accurately measuring time on a computer. However, we can try to minimize it by increasing the precision of execution time measurement of our programs. This can be achieved by averaging execution times over larger number of iterations of the main time-propagation loop. Unfortunately, this also means that the total execution time of the GD algorithm will significantly increase.

Another issue in our implementation of GD arises when we have to select the next values of our parameters based on the output values of previous GD

iteration. These output values are real-valued, meaning that we need to convert them to integers somehow. Simply rounding them up or down to the nearest integer will not be enough, especially for small value of  $\gamma$ . For example, if we set the initial value of some parameter to 10, and after one GD iteration the proposed value of that parameter is 10.2, rounding it down will reset the value back to 10, effectively discarding the whole GD iteration. To avoid getting stuck in this way, our implementation selects the next possible value in the direction of change. In the example above it would mean selecting 11 as the new value. This change makes it impossible for the algorithm to converge, but it will usually stay close the minimum where all proposed solutions are of similar quality.

With the GD algorithm implemented as described above, we are able to get a set of optimized parameters from a random set of initial values. However, such optimized parameters are often suboptimal, and do not have the best performance. This is due to the fact that in most cases the GD will converge to the nearest local minimum, and there is no guarantee that that local minimum is also the global one. Using our previous example from BFS, a mesh of  $256 \times 256 \times 256$ , with GD we obtain that local minima exist for all power of two values, i.e., parameters suggest offloading  $2^n \times 256 \times 256$  to the GPU. If we randomly select a small initial value, GD will get stuck in the nearest local minimum, which in this example will be far from the global one, thus producing a very bad solution. We can attempt to avoid getting stuck in a local minimum by adapting the parameter  $\gamma$ , also known as *learning rate* in machine learning literature. Keeping the learning rate high for a first few iterations would allow the algorithm to find the general location of the minimum, after which we could gradually lower the learning rate until convergence is achieved, in a process called *annealing* [43]. In practice this did not completely solve the problem, as we found that the optimal learning rate to start with varies with the initial set of parameters and the mesh size, and thus has to be manually selected. This complicates attempts to automate the process of finding the optimal parameters for a range of mesh sizes, which we needed as part of the tests of both hybrid implementations. A method to remove the manual tuning of the learning rate exists [44], but we find it to be too complex to implement for our programs, because it would require significant changes to the way the parameters are selected.

From the discussion above, we conclude that the GD is not the best-suited method for the optimization in our case, mostly because our minimization function is noisy and the numerical computation of the gradient is costly. Derivative-free optimization methods would be better suited to the problem, e.g., stochastic

approximation algorithms like simultaneous perturbation (SPSA) [45, 46], or metaheuristics like genetic algorithms. We decided to use a genetic algorithm approach, as it is simple to understand and implement, as well as easy to adjust to get the desired behavior.

Genetic algorithm (GA) is an optimization method based on natural selection that mimics the process found in biological evolution [47]. GA works by creating a population of individual solutions, which it then evaluates and modifies, creating a new population, and iterating this process. Unlike the classical algorithms like GD, which iterate a single candidate solution towards the optimal one, the GA iterates a population of solutions in which the best individuals approach the optimal solution. The initial population is usually created at random, giving the GA different points in the search space to start from. The individual solutions are evaluated using a user-supplied fitness function, giving each individual a score based on how well they perform the given task. Individuals with the highest score are then selected to “reproduce” and create new offspring, after which they may be mutated randomly. The offspring form a new population, and the process can be repeated again. The GA continues until a suitable solution is found, or after a certain number of generations has passed.

To implement a GA, we first have to decide how to represent an individual. Individuals are created based on their blueprint, called *chromosome* in GA terminology. The most often used representation is bit-string [48], where all properties of an individual (its *genes*) are serialized to an array of bits, and then concatenated. More advanced representations exist, e.g., for encoding real values, permutations and general data structures [49, 50], however since our individuals are sets of integer parameters, we did not develop any special representation and instead we used ordinary arrays of integers as a chromosome. Therefore, each gene in a chromosome is a single GPU parameter of our programs. Individuals need to be evaluated using a fitness function. In our case, this means executing the programs with the parameters extracted from the individual’s chromosome, and reporting the execution time as fitness score, with lower execution time being better.

Next step is the implementation of the three GA operators: selection, reproduction (*crossover* in GA terminology) and mutation. Each operator can be implemented in different ways. Most common type of the selection operator is *roulette wheel* selection. This is a type of fitness-proportionate selection, with the idea to give each individual a slice of the circular roulette wheel based on their fitness. The wheel is then spun, and when the roulette ball stops, the in-

dividual in whose region the ball stopped is selected. In this way, the fittest individuals have the greatest chance of being selected for reproduction, while the ones with very low score quickly die out. We also tested an alternative type of selection, the *tournament* selection. In this type of selection, we randomly select several individuals, and host a tournament for them. The individual with the best fitness score wins, and is selected for reproduction. By changing the size of the tournament we can control the selection pressure, e.g., by using small tournament size in the first few iterations we can prevent premature convergence and increase it in later generations when we have explored the search space enough. In our tests, the tournament selection gave slightly better results for smaller populations, by keeping the population diverse in early generations. For larger population size, both selection rules performed equally. Alongside the main selection algorithm, we also used *elitist* selection, where first few individuals are copied to the next generation without changing their chromosome. This prevented the loss of the best individuals in the next generation, but has to be used carefully as it may lead to premature convergence to a suboptimal solution.

To produce the next generation, selected individuals should combine their chromosomes and produce offspring. This is the task of the crossover operator. Crossover exchanges parts of the chromosomes, mimicking the biological recombination from nature. In its most basic form, crossover works by randomly selecting a point and exchanging segments before and after the point to create two new offspring from two parents. This type of crossover is called *single-point crossover*, and is illustrated in Figure 8(a). Other popular crossover techniques are *two-point crossover* and *uniform crossover*, illustrated in Figures 8(b) and 8(c). Two-point crossover is similar to the single-point crossover, just with two points instead of one. On the other hand, in the uniform crossover each gene of the offspring is selected randomly, either from the first parent or from the second one, with some fixed probability, typically 0.5. Using the uniform crossover leads to a wider exploration of the search space [51], but this may not always result in better performance of the operator [52]. We tested single- and two-point crossover and found that there is little difference in terms of performance between them. Since some of our parameters have constraints, it is important to select only the crossover points which lead to an allowed recombination of genes. Unfortunately, this made the implementation of uniform crossover impossible, and we did not pursue it further.

The final step in producing the next generation is to apply the mutation operator on the new population. The mutation operator randomly changes genes




Figure 8. Three different crossover techniques in GA: a) single-point crossover, b) two-point crossover and c) uniform crossover.

to new values, which is equivalent of performing a random walk through the search space. Changing every gene would not be desired, so mutation operates with a very small probability, meaning most of the genes will be left unchanged. When chromosomes are implemented as bits, the mutation would be equivalent to flipping a random bit. However, we could not implement mutation this way, as random changes to the chromosome would often result in non-functioning individuals, e.g., parameter constraints would not be satisfied. Special care must be taken to ensure that the mutation produces a healthy individual, similarly to the crossover operator. The importance of mutation operator and its relation to the crossover is often debated [53, 54], with its role being defined as “to maintain diversity within the population and inhibit premature convergence” [55], as the crossover operator does not introduce new information to the population. In our tests, the mutation plays a crucial role if the population size is small, e.g., less than 20 individuals, in accordance with references [56, 57]. If the population size is large, e.g., over 200 individuals, the positive effects of mutation on the population fitness are not that evident.

GA is typically iterated for a fixed number of generations, as is the case in our implementation. Alternatively, we could have implemented some exit clause in the main loop of the GA which would stop the evolution after the best individuals have not been improved for some number of generations. After each run of the GA, there is usually several highly fit individuals in the population. Since randomness is an integral part of any GA, different runs of the algorithm produce slightly different results. Therefore, our GA does not converge to a single solution, but produces candidate solutions that have very similar parameter values and overall fitness. Among the candidate solutions, there may be

small differences in the amount of data offloaded to GPU or the kernel parameters, with the execution times negligibly different. The number of generations required to produce a good solution also varies due to the randomness in the initial population and the population size. In our tests, it has a value between 10 and 50.

The implementation described above provides only the basis for a successful application of GA to the problem at hand. As can be seen, there are several important parameters of GA to tune in order to get optimal results. These include the population size, the number of generations, mutation rate, and the selection parameters (e.g., elite selection rate, the tournament size). We did not perform thorough testing of the performance of the GA that would allow us to obtain the best values for these parameters, as we were more focused on the quality of the candidate solutions GA creates. However, we observed that our GA finds good solutions faster if the initial population size is between 100 and 200, with the number of generations between 10 and 20, depending on the set of parameters we wish to optimize.

To get a perspective of how the three optimization methods perform, we tested them by comparing their final solutions, as well as by recording the number of program executions needed to get to the optimal set of parameters. For this test, we used 3D real-time propagation hybrid program on a single computing node. All three methods were allowed to execute the program up to 1000 times. The mesh sizes used range from  $80 \times 80 \times 80$  to  $600 \times 600 \times 600$ , all of which could potentially be offloaded to GPU. We tested parameter set  $PS_3$  (with 15 parameters), which controls the execution of FFT and the kernels in the subsequent functions that perform propagation w.r.t.  $H_1$ ,  $H_3$ , and  $H_4$ . The BFS algorithm was used as the baseline. Note that the range of parameters and a small number of allowed executions implies that the BFS algorithm takes the values of parameters with large stride, potentially missing the optimal solution. GD method is used as described, with the learning rate  $\gamma$  initially set to a higher value, which was gradually decreased. We performed the GD for 60 iterations, which amounts to 960 program executions. The GA was run on a population of 100 individuals, for 10 generations, amounting to the same the number of program executions as the BFS. We have used mutation rate of 5%, and the tournament selection, with the tournament size equal to 10% of the population size. Elitism was also included, with the top 2% of the population copied over. All algorithms were tested five times, to minimize the effects of random initialization of both GD and GA methods. The results are shown in Table 1.

**Table 1. A comparison of three optimization methods. Reported times are given in milliseconds, for a single iteration of the main time-propagation loop, averaged over 50 iterations. The last column contains the minimal execution times obtained by manual tuning. The reported time for each algorithm represents the minimal achieved value in five test runs**

Mesh size	BFS	GD	GA	Best
$80 \times 80 \times 80$	7	6	6	6
$128 \times 128 \times 128$	18	27	16	16
$240 \times 240 \times 240$	154	272	135	126
$256 \times 256 \times 256$	181	199	169	156
$360 \times 360 \times 360$	343	387	312	298
$480 \times 480 \times 480$	981	1049	868	829
$512 \times 512 \times 512$	1452	1628	1312	1242
$600 \times 600 \times 600$	2591	2984	2227	2159

From the results we conclude that the GA was the most effective optimization method of the three approaches. As expected, GA found better solutions than BFS due to large strides BFS had to use. The randomness in the initial population has a big effect on the convergence of the GA method, sometimes enabling it to find the optimal solution after just three generations. Even if the fitness of the initial population is very bad, the GA still converges to very good solutions after 10 generations. On the other hand, GD performed very poorly, mainly because the noise in program execution times has often thrown it in the wrong direction. Also, GD would get stuck in the nearest local minimum, which often was not the global one. When the initial position of GD is near the global minimum it converged to toward the optimal solution, which was rarely the case. For mesh sizes larger than the tested ones (relevant for MPI-based implementations), corresponding to a larger range of parameters, GD would be even less effective. Both BFS and GD can be made more usable if the range of parameters can be narrowed, i.e., if we know the relative performance of the GPU in comparison to the CPU. However, the GD would still be somewhat inefficient, due to its higher susceptibility to noise.

Since the GA method is shown to be superior to the other two approaches, we use GA as our optimization method of choice in the next two sections.

### 2.4.2. Testing Methodology

All programs were tested on the PARADOX-IV cluster, which is a part of the PARADOX supercomputing facility. This cluster is comprised of computing nodes with two Intel Xeon E5-2670 Sandy Bridge CPUs (with a total of  $2 \times 8 = 16$  cores), with 32 GB of RAM and one Nvidia Tesla M2090 GPU with 6 GB of GPU RAM, each connected by InfiniBand QDR interconnect. We used Intel's compiler (version 2016) to compile the serial and OpenMP programs, and CUDA 7.5 for the GPU portions of the CUDA and hybrid programs. MPI-based implementations were compiled with Open MPI (version 1.10), which itself relied on underlying Intel and CUDA compilers. In the case of Hybrid/MPI programs, we performed tests using both FFTW and our own transpose routines. We found the minimal execution times to be about the same for both approaches, but the FFTW transpose would sometimes exhibit very bad performance due to the creation of suboptimal communication plan. For this reason, the execution times reported for Hybrid/MPI implementation are obtained using only our own transpose routines.

The base of all performance evaluations was the measured execution time of critical regions of the programs, i.e., the portions performing wave function propagation in imaginary or in real time. This measurement excluded the time spent in other parts of the programs, e.g., initialization of OpenMP/CUDA/MPI environment, memory allocation and deallocation, creation and destruction of FFTW plans, initialization of variables and I/O operations. Measuring average execution time of a single iteration of the main time-propagation loop allows us to predict the performance and total execution time of a given simulation, as the number of iterations is specific to the problem at hand and may vary significantly between different simulations. All measurements were collected using high precision timers based on `clock_gettime` POSIX function on the CPU side, and CUDA event API on the GPU side. The execution time of a single iteration of serial programs depends on the mesh size, controlled by variables  $N_x$ ,  $N_y$  and  $N_z$ . For parallel programs, we can measure speedup and scaling efficiency as a function of the varying number of processing elements (OpenMP threads or MPI processes).

We tested performance of 1D, 2D and 3D programs on a range of mesh sizes, for a varying number of OpenMP threads and MPI processes, as shown in Table 2. Mesh sizes were chosen from the corresponding range, and we did not focus solely on mesh sizes which maximize performance of the programs, to



obtain a more realistic assessment. We varied the number of OpenMP threads from 1 to 16, in increments of one thread. Similarly, we tested the MPI-based implementations by varying number of MPI processes, each bound to a different cluster node. In this way we tested MPI-based implementations on 2, 4, 8, 12, 16, 20, 24, 28 and 32 computing nodes. Note that varying the number of processing elements is not applicable to the single-GPU implementations of the programs, as they always use all available processing resources on a GPU. Only one of the two 1D and 2D programs were tested (corresponding to  $x$  direction in 1D and  $x$ - $y$  plane in 2D), because, performance-wise, there is no difference between them (e.g., `imag1dX-th` vs. `imag1dZ-th`).

The main performance indicator is the execution time, the wall-clock time of one iteration of the main loop, averaged over 5 executions of 1000 iterations, reported in milliseconds. Using the results obtained, we calculated the speedup of all programs compared to the published serial C implementation. We were also interested in examining the scaling efficiency, or *scalability*, of the OpenMP and MPI programs. We tested both *strong* scaling, when the mesh size stays the same but the number of processing elements varies, and *weak* scaling, when the amount of work each processing element performs stays the same while the number of processing elements increases. More formally, given the execution time of a single iteration of serial programs ( $T(1)$ ), and the corresponding execution time for parallel programs performed with  $N$  processing elements ( $T(N)$ ), we calculated speedup as  $S(N) = T(1)/T(N)$  and strong scaling efficiency as  $E(N) = S(N)/N$ . Weak scaling is computed as  $E_W(N) = T_W(1)/T_W(N)$ , where  $T_W(1)$  is the execution time of a program using single processing element performing the work assigned to it, while  $T_W(N)$  is the execution time of a program using  $N$  processing elements performing  $N$  times more work. We achieve this by increasing the mesh size.

### 2.4.3. Performance Test Results and Modeling of Single Node Programs

In this section we present the results obtained for single computing node OpenMP, CUDA and hybrid programs, and compare them to the previously published [23] serial implementation.

**Table 2. Performance testing matrix, showing the mesh sizes and numbers of processing elements (threads or processes) used to test the programs, as well as the baseline program used for comparison**

Program	Mesh size		Processing elements		Baseline
	Min	Max	Min	Max	
OpenMP programs					
imag1dX-th	1000	1000000	1	16	imag1d
real1dX-th	1000	1000000			real1d
imag2dXY-th	1000 × 1000	15000 × 15000	1	16	imag2dXY
real2dXY-th	1000 × 1000	13000 × 13000			real2dXY
imag3d-th	50 × 50 × 50	800 × 800 × 800	1	16	imag3d
real3d-th	50 × 50 × 50	800 × 800 × 800			real3d
CUDA programs					
imag2dXY-cuda	1000 × 1000	15000 × 15000	1	1	imag2dXY
real2dXY-cuda	1000 × 1000	13000 × 13000			real2dXY
imag3d-cuda	50 × 50 × 50	600 × 600 × 600	1	1	imag3d
real3d-cuda	50 × 50 × 50	540 × 540 × 540			real3d
Hybrid programs					
imag2dXY-hetero	1000 × 1000	15000 × 15000	16+1	16+1	imag2dXY-th
real2dXY-hetero	1000 × 1000	13000 × 13000			real2dXY-th
imag3d-hetero	50 × 50 × 50	600 × 600 × 600	16+1	16+1	imag3d-th
real3d-hetero	50 × 50 × 50	600 × 600 × 600			real3d-th
OpenMP/MPI programs					
imag3d-mpi	480 × 480 × 250	1920 × 1920 × 960	1 × 16	32 × 16	imag3d-th
real3d-mpi	480 × 480 × 250	1920 × 1920 × 960			real3d-th
CUDA/MPI programs					
imag3d-mpicuda	480 × 480 × 250	1920 × 1920 × 960	1	32	imag3d-cuda
real3d-mpicuda	480 × 480 × 250	1920 × 1920 × 960			real3d-cuda
Hybrid/MPI programs					
imag3d-mpihetero	480 × 480 × 250	1920 × 1920 × 960	1 × (16+1)	32 × (16+1)	imag3d-hetero
real3d-mpihetero	480 × 480 × 250	1920 × 1920 × 960			real3d-hetero

**Table 3. Wall-clock execution times of a single iteration of the main time-propagation loop of single-node OpenMP programs (in milliseconds) for different number of OpenMP threads  $N_{\text{th}}$  and speedup  $S(16)$  in strong scaling tests. The speedup is calculated w.r.t. the execution times of previously published serial versions of programs [23], given in the second column**

Program	Serial	$N_{\text{th}} = 1$	$N_{\text{th}} = 2$	$N_{\text{th}} = 4$	$N_{\text{th}} = 8$	$N_{\text{th}} = 16$	$S(16)$
imag1dX-th	9.1	7.1	4.7	3.4	2.9	2.8	2.5
real1dX-th	15.2	14.2	10.5	8.2	7.3	7.2	2.0
imag2dXY-th	13657	7314	4215	2159	1193	798	9.2
real2dXY-th	17281	11700	6417	3271	1730	1052	11.1
imag3d-th	16064	9353	5201	2734	1473	888	10.5
real3d-th	22611	17496	9434	4935	2602	1466	11.9

Strong scaling performance test results for the OpenMP-based implementation using methodology described in the previous section are given in Table 3 and Figure 9. They show the obtained execution times, speedups and strong scaling efficiencies for different number of OpenMP threads. Columns  $N_{\text{th}} = 1$ ,  $N_{\text{th}} = 4$ ,  $N_{\text{th}} = 8$  and  $N_{\text{th}} = 16$  in Table 3 correspond to the number of threads used, while the last column shows the obtained speedup  $S(16)$  with 16 OpenMP threads compared to one OpenMP thread. Strong scaling efficiency  $E(N_{\text{th}}) = S(N_{\text{th}})/N_{\text{th}}$  in Figure 9 is calculated as a fraction of the obtained speedup compared to a theoretical maximum. The mesh size used in 1D is  $10^5$ , in 2D  $10^4 \times 10^4$ , while in 3D the mesh size is  $480 \times 480 \times 480$ . Execution times and speedups of imag1dZ-th, real1dZ-th, imag2dXZ-th, and real2dXZ-th (not reported here) are similar to those of imag1dX-th, real1dX-th, imag2dXY-th, and real2dXY-th, respectively.

The change from C2C to R2C FFT routine has a big impact on the execution time of single-threaded ( $N_{\text{th}} = 1$ ) programs compared to the previous serial programs. As we can see from the table, these improvements alone yield a speedup of 1.3 to 1.9 in 2D and 3D programs, and somewhat smaller speedup for 1D programs, 1.1 to 1.3. The use of additional threads brings about further speedup (reported in the last column) of 2 to 2.5 for 1D programs, and 9 to 12 for 2D and 3D programs. In Figure 9 we see that the efficiency rapidly decreases for 1D programs, even though speedup increases with the number of threads used.

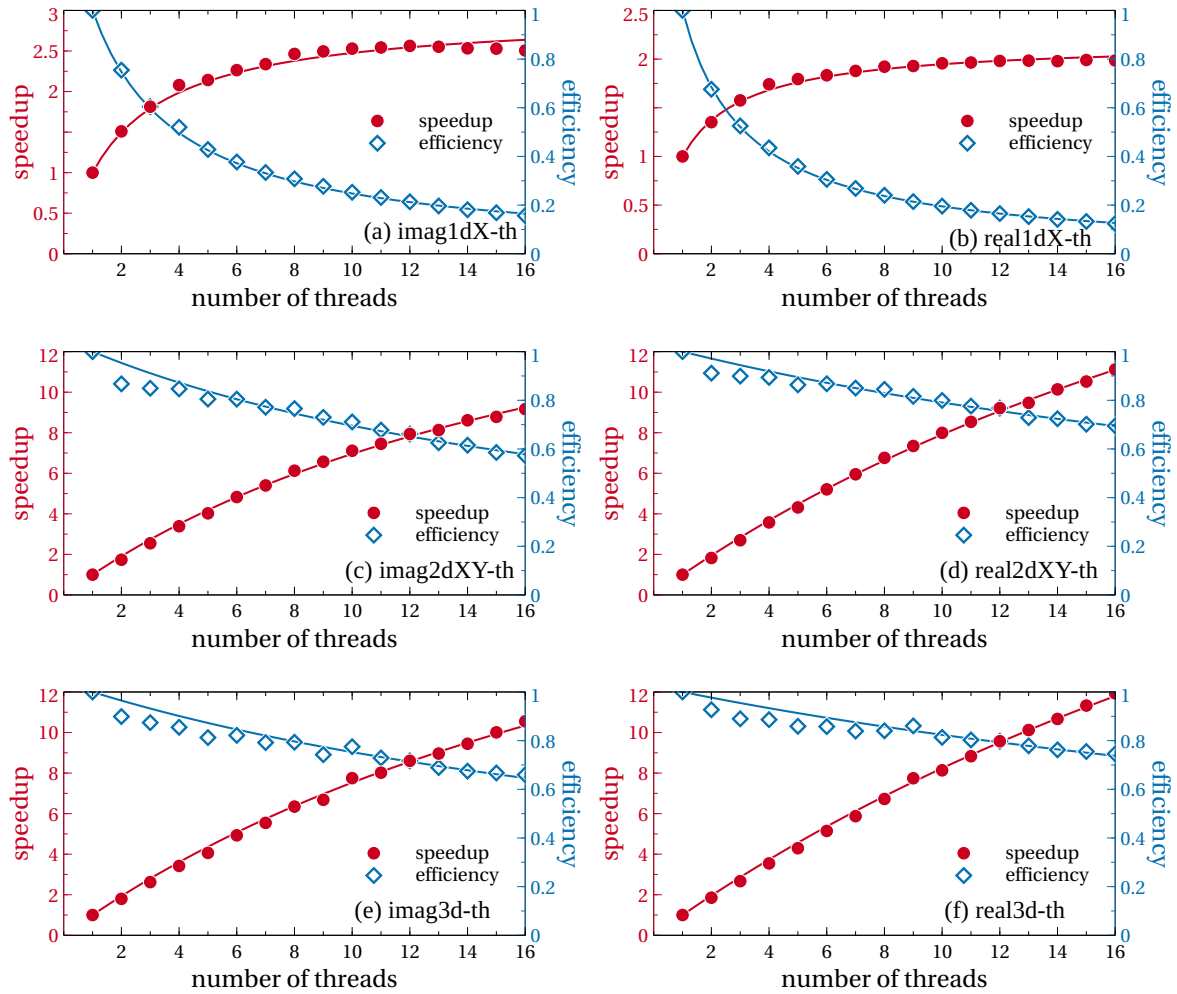


Figure 9. Speedup in the execution time and strong scaling efficiency of OpenMP programs compared to single-threaded runs for: (a) imag1dX-th, (b) real1dX-th, (c) imag2dXY-th, (d) real2dXY-th, (e) imag3d-th, (f) real3d-th. Solid lines represent fits to measured data, where fit model functions are given in the text and obtained fit parameters are listed in Table 4. Note that the model parameter  $p$  is fitted only to the speedup data, and then used to plot the efficiency model curve.

This is expected, as parts of the algorithm dealing with the recursive relations for calculation of the CN coefficients are inherently serial. In 1D, already with  $N_{\text{th}} = 4$  threads we almost achieve the maximal speedup, while still keeping the efficiency around 50%. We also see that, as expected, speedup and efficiency of multidimensional programs behave quite well as we increase the numbers of threads. In particular, we note that the efficiency always remains above 60%,

making the use of all available CPU cores worthwhile.

From Figure 9 we observe that the speedup of 1D programs saturate quickly due to inherent serial nature of the portion of the algorithm, while in 2D and 3D the speedup behaves almost linearly. Despite their obvious differences, all curves in Figure 9 can be successfully modeled based on Amdahl's law [58]. Namely, the measured execution time  $T(N_{\text{th}})$  of one iteration of the main loop can be expressed as

$$T(N_{\text{th}}) = T(1) \left( s + \frac{p}{N_{\text{th}}} \right), \quad (33)$$

where  $N_{\text{th}}$  is the number of threads used,  $T(1)$  is the execution time of a single-threaded run,  $s$  is the serial fraction of the loop code, and  $p$  is the corresponding parallel fraction. By definition,  $s + p = 1$ , and therefore the speedup can be modeled by

$$S(N_{\text{th}}) = \frac{T(1)}{T(N_{\text{th}})} = \frac{1}{1 - p + p/N_{\text{th}}}, \quad (34)$$

where the parallel fraction of the main loop code  $p$  is the only fit parameter. Table 4 gives the obtained model fit parameter values for the data from Figure 9. As we can see, the fits match the obtained measurement data very well. Note that the efficiency can be expressed in terms of the speedup model

$$E(N_{\text{th}}) = \frac{S(N_{\text{th}})}{N_{\text{th}}} = \frac{1}{N_{\text{th}}} \frac{1}{1 - p + p/N_{\text{th}}}, \quad (35)$$

and is not fitted independently. Figure 9 shows that this model with the parameter  $p$  fitted on the speedup data matches very well with the efficiency data points.

From the obtained values of the parallel fraction of the algorithm, we see that 2D and 3D programs are almost ideally parallelizable, with  $p$  over 95%. In 1D case, however, the parallel fraction is around 66% for imaginary-time propagation and around 54% for real-time propagation, due to the fact that calculation of CN coefficients in the function `calclux` cannot be parallelized. Note that the parallel fraction of imaginary-time 1D program is higher than that of the real-time 1D program due to the larger amount of arithmetic operations required to process complex-valued data in the real-time version of the serial function `calclux`.

OpenMP programs were also tested for weak scalability. We were mostly interested in 3D variants of the programs, which we tested by fixing the amount

**Table 4. Values of obtained strong scaling model fit parameter  $p$  and the estimated fit errors for OpenMP-parallelized programs**

Program	$p$	$\Delta p$
imag1dX-th	0.662	0.004
real1dX-th	0.541	0.003
imag2dXY-th	0.9514	0.0007
real2dXY-th	0.9706	0.0004
imag3d-th	0.9635	0.0008
real3d-th	0.9762	0.0006

of work to 6,912,000 spatial points of the mesh, which corresponds to a mesh size of  $240 \times 240 \times 120$ . By increasing the number of OpenMP threads, we also increase the mesh size to be the multiple of 6,912,000. This means that the mesh has to be increased in such a way that, when divided among the threads, each thread gets 6,912,000 spatial points to process. The mesh sizes that satisfy this requirement cannot always be obtained by multiplying all three array dimensions with the same number, so we have to work with mesh sizes in which the size of each dimension may be different. For such mesh sizes, we tested all possible combinations, however no significant difference has been observed. Table 5 contains mesh sizes we used for testing for  $N_{\text{th}} = 1, 2, 4, 8, 12, 16$  threads, along with the execution times (in milliseconds) and the weak scaling efficiency. The testing was also done for  $N_{\text{th}} = 6$  and 10 threads, but the corresponding mesh sizes are omitted from the table for brevity. However, Figure 10 presents complete data collected in this test, for all  $N_{\text{th}}$  values, averaged over mesh sizes used. The figure shows that the real-time 3D programs have better weak scaling efficiency, which is about 75% at 16 threads, while imaginary-time programs demonstrate smaller, but still significant efficiency of about 60%.

To model weak scaling efficiency, we compare execution times of a single iteration of the main loop  $T_W(N_{\text{th}})$ , performed with  $N_{\text{th}}$  threads, where for each value of  $N_{\text{th}}$  the total amount of work is  $N_{\text{th}}$  times the work being performed in a single-threaded run, i.e., the amount of work per thread is constant. The weak scaling efficiency is defined as

$$E_W(N_{\text{th}}) = \frac{T_W(1)}{T_W(N_{\text{th}})}, \quad (36)$$

where  $T_W(1) = T(1)$  is the execution time of a single-threaded run for a given

**Table 5. Weak scaling efficiency of OpenMP-parallelized 3D programs. Wall-clock execution times  $T_W$  are given in milliseconds, efficiencies  $E_W$  in percents**

Mesh size	imag3d-th		real3d-th	
	$T_W$	$E_W$	$T_W$	$E_W$
$N_{\text{th}} = 1$				
$120 \times 240 \times 240$	699	98.2	1259	99.5
$240 \times 120 \times 240$	711	96.5	1271	100
$240 \times 240 \times 120$	686	100	1275	98.2
$N_{\text{th}} = 2$				
$240 \times 240 \times 240$	718	95.5	1254	99.9
$N_{\text{th}} = 4$				
$240 \times 240 \times 480$	755	90.9	1312	95.5
$240 \times 480 \times 240$	778	88.2	1301	96.3
$480 \times 240 \times 240$	754	91	1310	95.6
$N_{\text{th}} = 8$				
$240 \times 480 \times 480$	875	78.4	1405	89.9
$480 \times 240 \times 480$	847	81	1393	89.1
$480 \times 480 \times 240$	886	77.5	1431	87.7
$N_{\text{th}} = 12$				
$360 \times 480 \times 480$	1010	67.9	1532	81.7
$480 \times 360 \times 480$	981	69.9	1514	82.7
$480 \times 480 \times 360$	1004	68.4	1581	79.2
$N_{\text{th}} = 16$				
$480 \times 480 \times 480$	1155	59.4	1680	74.5

workload (in our case, 6,912,000 spatial points). The expected execution time of a workload assigned to a simulation with  $N_{\text{th}}$  threads, executed in a single-threaded run, is  $N_{\text{th}}T(1)$ . In weak scaling tests, this workload is executed with  $N_{\text{th}}$  threads and therefore the expected execution time can be modeled by

$$T_W(N_{\text{th}}) = N_{\text{th}}T(1) \left( 1 - p + \frac{p}{N_{\text{th}}} \right). \quad (37)$$

According to this argument, we model the weak scaling efficiency by a single-parameter function

$$E_W(N_{\text{th}}) = \frac{1}{p + (1 - p)N_{\text{th}}}. \quad (38)$$




Figure 10. Weak scaling efficiency of OpenMP-parallelized 3D programs, averaged over all mesh sizes tested for each value of  $N_{th}$ . Solid lines represent fits to measured data, where fit model functions are given in the text.

We fitted this model to data presented in Figure 10 and the obtained fit parameters are given in Table 6. As we can see from the figure, the above model is an excellent fit to experimental data, and both 3D OpenMP-parallelized programs have high (above 95%) parallel fraction of the code, in agreement with the results of strong scaling tests (Table 4).

**Table 6. Values of obtained weak scaling model fit parameter  $p$  and the estimated fit errors for OpenMP-parallelized programs**

Program	$p$	$\Delta p$
imag3d-th	0.958	0.002
real3d-th	0.9792	0.0009

Overall, we conclude that imaginary-time OpenMP-parallelized programs show smaller speedup and efficiency across all tests, except in 1D. This can be attributed to the additional step of wave function normalization in each iteration of the imaginary-time propagation, as well as the fact that real-time programs work with complex-valued data that require more arithmetic operations for the same mesh size. Since much of the computation inside loops requires simple arithmetic, the throughput of the CPU is often not fully exploited in imaginary-time programs, thus the pressure of memory bandwidth makes these programs less efficient. The real-time programs are also affected by this, however to a somewhat lesser extent.



Next, we consider the CUDA implementation of the shared memory algorithm. GPU functions as a single processing element, therefore we cannot test CUDA implementation by varying both the mesh size and the number of processing elements. However, just varying the mesh size gives us valuable insight into the behavior of this implementation, due to the difference in programming models and the libraries used. Tables 7 and 8 show the execution times (in milliseconds) for a number of mesh sizes tested, as well as the average speedup  $\langle S \rangle$  compared to the serial programs [23]. Figure 11 shows the speedup obtained for all mesh sizes tested, and horizontal lines represent average speedups obtained for each program. Note that the dispersion of data is due to the use of `FFTW_ESTIMATE` flag in library calls to `FFTW` in the serial programs. Use of this flag results in a choice of suboptimal FFT algorithm for some mesh sizes. The vertical lines in Figure 11 denote the change in `POTMEM` parameter. The speedups left of the first vertical line are obtained with `POTMEM=2` and thus demonstrate the best speedup. Second group of results, between the two vertical lines, is obtained with `POTMEM=1`, and we note that the speedup decreases slightly, while the results right of the second vertical line are obtained with `POTMEM=2`, and show the smallest speedup due to the use of mapped memory. The 2D programs in  $x$ - $z$  plane exhibited very similar performance to those of  $x$ - $y$  plane, and therefore we did not include them in the figures.

**Table 7. Wall-clock execution times of a single iteration of the main time-propagation loop of single-node 2D CUDA programs (in milliseconds) for different mesh sizes, and average speedup w.r.t. to the execution times of serial 2D programs [23]**

Program	2000 <sup>2</sup>	4000 <sup>2</sup>	6000 <sup>2</sup>	8000 <sup>2</sup>	10000 <sup>2</sup>	12000 <sup>2</sup>	Avg. speedup
imag2dXY-cuda	24.1	104.3	235.2	386.1	657.1	1150.4	10
real2dXY-cuda	29.9	112.4	266.4	444.0	749.0	1528.3	14

For small mesh sizes, the GPU remains underutilized, resulting in a smaller speedup. For large mesh sizes, where the GPU memory usage approaches the limit, we also see declining speedup. This is due to the inevitable use of `POTMEM` parameter, which keeps some arrays in the host memory when they cannot fit in the memory of the GPU. Overall, the CUDA implementation shows execution times similar to the OpenMP implementation for imaginary-time propagation, and slightly lower for real-time propagation. We stress that these results strongly depend on the type of GPU used and that the obtained speedups

**Table 8. Wall-clock execution times of a single iteration of the main time-propagation loop of single-node 3D CUDA programs (in milliseconds) for different mesh sizes, and average speedup w.r.t. to the execution times of serial 3D programs [23]**

Program	$100^3$	$200^3$	$300^3$	$400^3$	$500^3$	Avg. speedup
imag3d-cuda	10.6	79.3	298.8	674.5	1260.2	7.1
real3d-cuda	10.9	84.1	302.5	682.4	1467.4	13.5

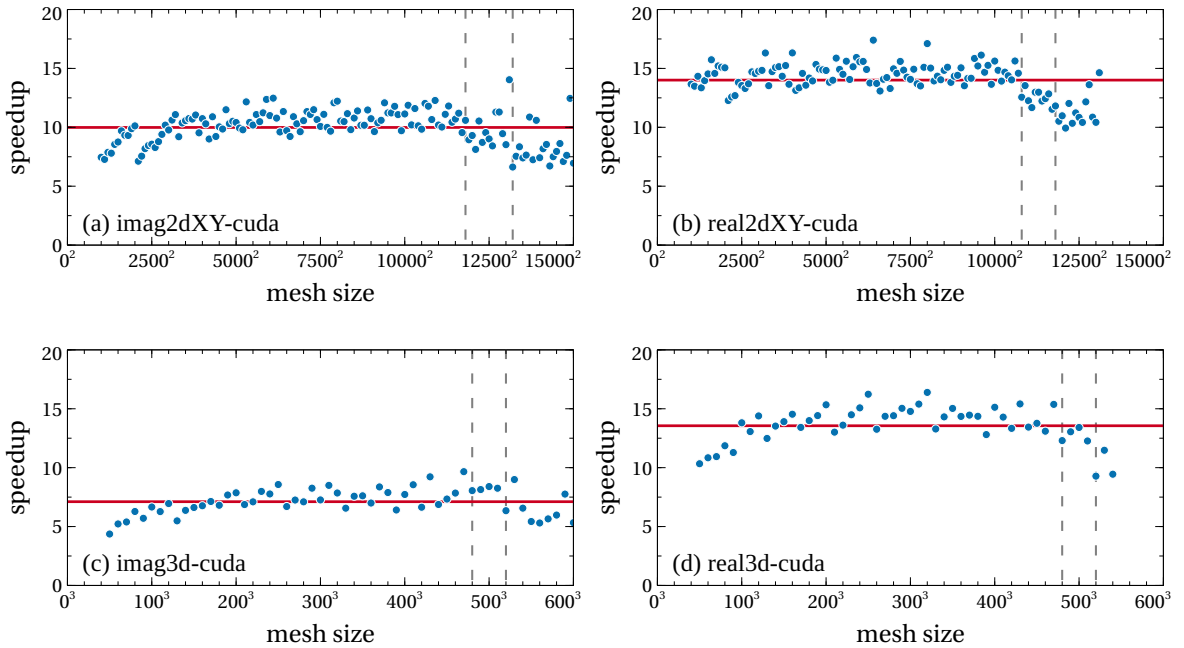


Figure 11. Speedup in the execution time of CUDA programs compared to the serial programs for all tested mesh sizes: (a) *imag2dXY-cuda*, (b) *real2dXY-cuda*, (c) *imag3d-cuda*, (d) *real3d-cuda*. Horizontal lines represent  $\langle S \rangle$ , while dashed vertical lines represent different values of POTMEM parameter (see text).

may be even better for newer-generation GPUs (e.g., *Kepler*-, *Maxwell*-, *Pascal*- and *Volta*-based GPUs).

Hybrid OpenMP/CUDA implementation was tested on a range of mesh sizes, similarly to the OpenMP- and CUDA-based implementations. Since OpenMP implementation showed that all cores should be used for 2D and 3D programs, we kept the number of OpenMP threads fixed at 16 in our hybrid algorithm. The parameters governing the amount of data offloaded to GPU were optimized using our GA method from section 2.4.1 An illustrative subset of the

results for 3D programs is shown in Table 9, which gives execution times for a single iteration of the main loop, for different mesh sizes. A comparison of execution times of hybrid and pure OpenMP programs with  $N_{\text{th}} = 16$  threads is given in Figure 12(a), while Figure 12(b) compares the performance of hybrid and pure CUDA programs, for all tested mesh sizes.

**Table 9. Wall-clock execution times of a single iteration of the main time-propagation loop of single-node 3D hybrid programs (in milliseconds) for different mesh sizes. Offload parameters are optimized using the GA optimization method**

	$80^3$	$200^3$	$320^3$	$440^3$	$560^3$	$680^3$	$800^3$
imag3d-hetero	6.3	67.1	271.4	678.9	1493.3	2750.9	4461.6
real3d-hetero	8.0	94.7	374.1	922.8	1985.2	4020.2	6676.6

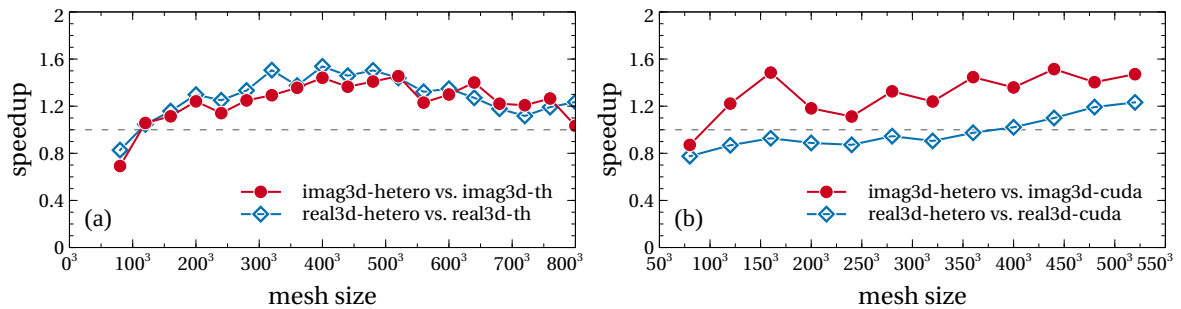


Figure 12. Speedup in the execution time of 3D hybrid programs compared to: (a) OpenMP programs with  $N_{\text{th}} = 16$  threads, (b) CUDA programs. Dashed horizontal lines correspond to a speedup value  $S = 1$ .

From Figure 12 we see that the hybrid implementation outperforms the OpenMP one for all mesh sizes except for the smallest one. The same applies to the comparison of imaginary-time hybrid and CUDA programs, while real-time hybrid program outperforms the corresponding CUDA program only for mesh sizes larger than  $400 \times 400 \times 400$ . Although one would expect that the speedup of the optimized hybrid algorithm is always equal to or larger than one, we see that this is not the case for all mesh sizes in Figure 12. This is due to the fact that the hybrid FFT algorithm is always employed in hybrid programs, meaning

that even if no or all data are offloaded to GPU, the splitting of FFT will still take place. Therefore, in these limiting cases we do not obtain pure OpenMP or CUDA algorithms that would yield maximal performance. Splitting the computation of FFT along one direction in `calcpsidd2` function disables some of the potential optimizations that libraries like FFTW and cuFFT exploit, and introduces data copies between host and device which cannot be fully offset using CUDA streams. The hybrid algorithm can compensate for this if the amount of offloaded data to GPU is sufficiently large, which is not the case for the smallest mesh size.

**Table 10. Optimal fraction of total data offloaded to GPU for different mesh sizes in hybrid 3D programs, obtained using the GA optimization method**

Mesh size	GPU portion of data [%]							
	imag3d-hetero				real3d-hetero			
	PS <sub>1</sub>	PS <sub>2</sub>	PS <sub>3</sub>	PS <sub>4</sub>	PS <sub>1</sub>	PS <sub>2</sub>	PS <sub>3</sub>	PS <sub>4</sub>
80 × 80 × 80	20	3	60	33	23	3	40	10
120 × 120 × 120	12	10	25	28	23	10	25	27
160 × 160 × 160	15	7	31	24	44	19	34	23
200 × 200 × 200	38	27	37	28	44	26	46	26
240 × 240 × 240	38	28	40	30	38	28	47	25
280 × 280 × 280	40	31	30	33	46	22	37	25
320 × 320 × 320	33	38	29	36	45	28	69	44
360 × 360 × 360	41	33	27	36	49	33	43	31
400 × 400 × 400	38	25	30	38	46	24	47	21
440 × 440 × 440	31	35	29	39	45	35	43	32
480 × 480 × 480	16	19	22	41	45	33	62	35
520 × 520 × 520	40	31	28	46	45	25	40	37
560 × 560 × 560	38	41	30	43	36	24	35	32
600 × 600 × 600	37	27	30	43	29	29	29	28
640 × 640 × 640	36	40	30	43	24	22	24	23
680 × 680 × 680	26	38	28	38	20	16	20	20
720 × 720 × 720	23	26	25	32	17	13	17	17
760 × 760 × 760	21	22	25	27	14	14	14	13
800 × 800 × 800	19	23	23	21	11	8	12	12

Given that pure CUDA programs performing imaginary-time propagation demonstrated smaller speedup than the corresponding OpenMP programs (Figure 11 vs. Figure 9), it is expected that CUDA portion of the hybrid imaginary-time propagation programs would yield smaller improvement, which is evident when we compare hybrid implementation with the OpenMP one in Figure 12(a), except for the largest mesh sizes. We can reach the same conclusion if we consider that, in real-time propagation, CUDA implementation shows the highest speedup of all single-node programs, and therefore the corresponding hybrid programs show better speedup than imaginary-time propagation programs parallelized with OpenMP. On the other hand, when compared with the CUDA implementation in Figure 12(b), the situation is reversed, and we see much better performance of the hybrid imaginary-time propagation programs, while the hybrid real-time propagation programs achieve speedup larger than one only for mesh sizes larger than  $400 \times 400 \times 400$ . This result can be attributed to the fact that in real-time propagation the amount of data copied between host and device is larger, given that some of the copied arrays are complex-valued.

Also, we can observe that the speedup in Figure 12(a) declines with mesh size. This is due to the memory saturation of the GPU device, as only a small portion of the data can be offloaded when the mesh is large. The amount of data offloaded to GPU for all tested mesh sizes is given in Table 10, where we can see that the total amount of data processed by the GPU declines with increasing mesh size. As a result, the computation is unbalanced, given that the host has to work with a much larger portion of the data. This is reflected in the optimization step, where our GA quickly converges to the highest possible mesh size that can be offloaded to the device, implying that a greater amount of device memory would be required to get a balanced computation.

To make an optimal choice of programs to be used on a single node for a particular hardware platform, one has to perform detailed tests using the methodology presented in section 2.4.2. The results presented in this section apply to hardware comparable to the one available at PARADOX supercomputing facility, but one can expect similar behavior for all modern types of CPU and GPU.

#### **2.4.4. Performance Test Results and Modeling of MPI Programs**

Following on from the previous section, here we present the results obtained for MPI-parallelized programs, executed on a computer cluster consisting of varying number of computing nodes, and compare them with the corresponding

single-node programs.

MPI programs are highly dependent on the configuration of the cluster, mainly on the speed of interconnect, but also on the distribution of processes and threads, NUMA configuration, MPI configuration, etc. Getting the best performance out of the programs requires some experimentation with several different configurations. The results presented are obtained without extensive tuning of the cluster or MPI runtime, with the aim to show the base performance.

Strong scaling of MPI-based implementations is tested in a similar way to the single-node testing, by varying number of cluster nodes  $N_p$  from 2 to 32. On each node, a single MPI process is launched, which then uses the node's resources, either by further spawning 16 OpenMP threads in case of OpenMP/MPI and Hybrid/MPI implementations, or by invoking the CUDA kernels utilizing the node's GPU in case of CUDA/MPI implementation. In this test, the baseline used for comparison was the equivalent single-node implementation, i.e., OpenMP/MPI programs are compared with OpenMP programs executing with  $N_{th} = 16$  threads, CUDA/MPI programs are compared to the single-node pure CUDA programs, and Hybrid/MPI programs is compared to the single-node hybrid programs.

Mesh size we use in this test is  $480 \times 480 \times 250$ . Since the values of  $N_x$  and  $N_y$  parameters must be divisible by the number of processes  $N_p$ , this mesh cannot be distributed over 28 processes. In this case we use slightly modified mesh, taking into account the multiple of 28 that is closest to 480, yielding a mesh size of  $476 \times 476 \times 250$ . Data is never distributed along  $z$  direction, so no such requirement exists for  $N_z$ . The resulting mesh has slightly fewer spatial points, and thus less work per process, potentially allowing for a better performance, but in our tests the smaller mesh did not have a significant impact on the measured execution time.

Note that strong scaling tests, where the mesh size is fixed, inevitably lead to saturation and decrease in measured speedup and efficiency values. Whatever the chosen mesh size, increase in the number of cluster nodes used will eventually yield insufficient amount of work per MPI process, such that communication will start to dominate over computation. Therefore, in our strong scaling tests we can expect that execution times initially decrease with increasing values of  $N_p$ , but eventually performance of MPI-based programs will decline. This issue is addressed by considering weak scaling tests.

**Table 11. Wall-clock execution times of a single iteration of the main time-propagation loop of OpenMP/MPI, CUDA/MPI and Hybrid/MPI programs (in milliseconds) for different number of MPI processes  $N_p$  and speedup  $S(32)$  in strong scaling tests. The speedup is calculated w.r.t. the baseline execution times of the corresponding single-node programs (OpenMP, CUDA and hybrid, respectively), given in the second column.**

Program	Baseline	$N_p = 4$	$N_p = 8$	$N_p = 16$	$N_p = 24$	$N_p = 32$	$S(32)$
imag3d-mpi	1124	541	262	134	89	64	17.5
real3d-mpi	2140	700	358	207	155	98	21.8
imag3d-mpicuda	579	438	210	103	71	59	9.8
real3d-mpicuda	800	609	291	142	95	79	10.1
imag3d-mpihetero	489	299	162	99	84	81	6.0
real3d-mpihetero	613	407	255	154	135	101	6.0

The illustrative portion of obtained execution times of a single iteration of the main loop for the three MPI-parallelized implementations are shown in Table 11, together with the execution times of the corresponding baseline single-node program. Columns  $N_p = 4$ ,  $N_p = 8$ ,  $N_p = 16$  and  $N_p = 32$  correspond to the number of cluster nodes used, while the last column shows the obtained speedup  $S(32)$  with  $N_p = 32$  nodes compared to baseline single-node programs.

The maximal speedup of OpenMP/MPI implementation ranges from 17 for imaginary-time propagation to 22 for real-time propagation programs for  $N_p = 32$ . The complete measurement results for speedup and efficiency are depicted in Figure 13(a) and 13(b), where we see that the speedup grows linearly with the number of nodes used, while the efficiency remains mostly constant in the range between 40% and 60%, thus making the use of OpenMP/MPI programs highly advantageous for simulations with large mesh size. In general, we can expect even better efficiency for larger mesh sizes.

Similar behavior is observed for CUDA/MPI implementation. The obtained speedup with  $N_p = 32$  nodes ranges from 9 to 10, with the slightly lower efficiency, between 30% and 40%, as shown in Figure 13(c) and 13(d). Even though the efficiency is lower for this implementation, the speedup still grows linearly and the execution times are lower than for the OpenMP/MPI implementation. This makes CUDA/MPI programs ideal choice for use on GPU-enabled computer clusters. Additional benefit of using CUDA/MPI programs is their low CPU usage (using only one CPU core per cluster node), allowing for the possibility that the same cluster nodes are used for other CPU-intensive simulations in a time-sharing fashion.

As we see in Figure 13(e) and 13(f), the linear growth of speedup is also present for the Hybrid/MPI implementation, however, as the amount of work per-process shrinks, the efficiency drastically drops down to 20% with  $N_p = 32$  nodes. For the mesh size used in this test, we observe that the Hybrid/MPI implementation performs very well on  $N_p < 16$  processes, providing the lowest execution times, but with  $N_p \geq 16$  processes its execution times become larger than for the CUDA/MPI implementation (especially for real-time propagation programs), and eventually even for the OpenMP/MPI implementation, as illustrated in Figure 14. This is again due to the insufficient amount of work each computing node performs after data are distributed among the MPI processes, which can be seen from the amount of data offloaded to GPU for all tested values of  $N_p$  in Table 12. Similar saturation in performance will eventually happen



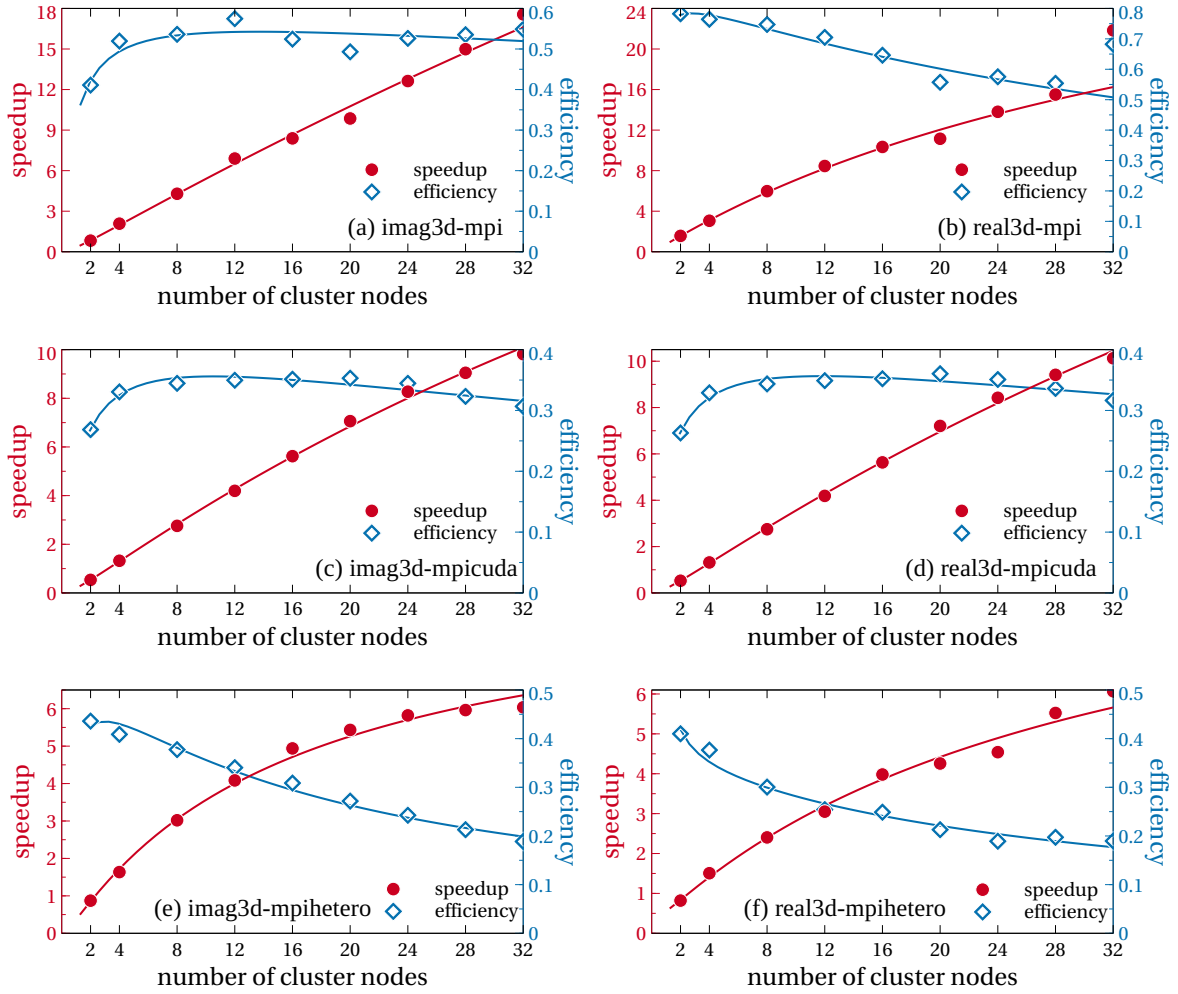


Figure 13. Speedup in the execution time and strong scaling efficiency of MPI-based programs compared to single-node runs: (a) `imag3d-mpi`, (b) `real3d-mpi`, (c) `imag3d-mpicuda`, (d) `real3d-mpicuda`, (e) `imag3d-mpihetero`, (f) `real3d-mpihetero`. Solid lines represent fits to measured data, where fit model functions are given in the text and obtained fit parameters are listed in Table 13. Note that the model parameters are fitted only to the speedup data, and then used to plot the efficiency model curve.

for the OpenMP/MPI and CUDA/MPI implementations for larger values of  $N_p$ .

We thus conclude that Hybrid/MPI implementation has the best performance of the three MPI-based implementations if the amount of work per process remains high enough to justify the use of hybrid algorithm. Otherwise, either of the other two MPI-based implementations should be considered first. The energy efficiency of this and other MPI-based implementations was not




Figure 14. Speedup in the execution time of Hybrid/MPI programs in strong scaling tests compared to the other two MPI-based programs: (a) OpenMP/MPI, (b) CUDA/MPI.

explored due to the difficulty in making precise measurements. If the energy consumption is not an issue, Hybrid/MPI implementation will yield the best performance, providing the cluster has powerful GPUs installed and the mesh size used is large enough.

**Table 12. Optimal fraction of total data offloaded to GPU per process for different values of  $N_p$  in Hybrid/MPI programs, obtained using the GA optimization method in strong scaling tests**

$N_p$	GPU portion of data [%]							
	imag3d-mpihetero				real3d-mpihetero			
	PS <sub>1</sub>	PS <sub>2</sub>	PS <sub>3</sub>	PS <sub>4</sub>	PS <sub>1</sub>	PS <sub>2</sub>	PS <sub>3</sub>	PS <sub>4</sub>
1	30	17.5	34.4	36.7	45.2	32.5	55	35
2	25	10	17.5	16.7	20.8	11.5	23.8	16
4	10	5.2	9.2	9.4	10	5.2	15	6.7
8	5	2.9	5.8	6.7	6.3	2.1	6.7	5
12	3.3	1.7	4.2	4.2	4	1.3	5	3.3
16	2.9	1.3	2.9	2.1	3.1	0.4	2.9	2.1
20	1.7	0.6	2.5	1	2.5	0.4	2.5	1.7
24	1.7	1.3	2.1	1.7	2.1	1.3	2.5	0.8
28	1.7	0.4	2.1	1.3	1.7	0.4	1.7	1.3
32	0.8	0.4	1.7	0.8	1.5	0.4	2.1	0.8

We now model the speedup and strong scaling efficiency of MPI-based programs. In general, the execution time of a program can be expressed as

$$T(N_p) = \alpha + \beta L + \gamma \frac{M}{N_p} + \beta V \frac{M}{N_p^2}, \quad (39)$$

where  $\alpha$  represents the average time to perform serial portion of the code,  $L$  is the communication latency associated with one MPI message,  $\beta$  is the frequency of MPI messaging,  $\gamma$  is the average time to process one spatial point,  $V$  is related to data transfer speed (throughput), and  $M$  is the mesh size ( $N_x \times N_y \times N_z$ ). The communication overhead of one all-to-all message passing instance is equal to  $L + VM/N_p^2$ , where each of  $N_p$  processes communicates its  $M/N_p$  part of the mesh evenly to all other processes, leading to a message size of  $M/N_p^2$ . Taking into account that in the strong scaling tests the mesh size is fixed, the above model can be simplified to

$$T(N_p) = T(1) \left( a + \frac{b}{N_p} + \frac{c}{N_p^2} \right), \quad (40)$$

while the speedup can be modeled by

$$S(N_p) = \frac{T(1)}{T(N_p)} = \frac{1}{a + b/N_p + c/N_p^2}. \quad (41)$$

**Table 13. Values of obtained strong scaling model fit parameters  $a$ ,  $b$  and  $c$ , as well as the estimated fit errors for MPI-parallelized programs**

Program	$a$	$\Delta a$	$b$	$\Delta b$	$c$	$\Delta c$
imag3d-mpi	0.008	0.006	1.6	0.2	1.4	0.4
real3d-mpi	0.026	0.006	1.1	0.1	0.2	0.3
imag3d-mpicuda	0.025	0.008	2.3	0.2	2.8	0.6
real3d-mpicuda	0.020	0.008	2.3	0.2	2.8	0.6
imag3d-mpihetero	0.104	0.009	1.7	0.2	0.9	0.5
real3d-mpihetero	0.09	0.02	2.7	0.3	-1.0	0.6

This model is fitted to the obtained strong scaling measurement data and the results are presented in Table 13, while the corresponding model curves are

shown as solid lines in Figure 13. As we see, the proposed model agrees very well with the experimental data. The only exception is the speedup and efficiency of real-time OpenMP/MPI program (real3d-mpi) on  $N_p = 32$  processes, where FFTW library creates an optimal transform plan that works very well with the given mesh size. However, changing the mesh size even slightly gives the performance comparable to the model prediction. We note that the value of the model parameter  $c$  is negative for real-time Hybrid/MPI program (real3d-mpihetero), which should not be the case since this parameter is related to data transfer speed. However, due to uncertainty of the optimization choice by the GA method and the fact that the fraction of total data offloaded to GPU gradually decreases (Table 12), leading from real hybrid algorithm to almost pure OpenMP/MPI one with hybrid FFT, it is not surprising that the obtained value differs from expected. Taking into account relatively large fit error  $\Delta c/c = 60\%$  for this parameter, we can still use expression (41) to model performance of Hybrid/MPI programs.

Next, we test weak scaling of MPI-based implementations. The same number of cluster nodes (and thus MPI processes) was used as in previous tests, while the starting mesh, which corresponds to a unit workload, had a size of  $480 \times 480 \times 480$ , amounting to 110,595,000 spatial points. This number of spatial points was kept constant per process, similarly to the weak test of OpenMP programs. There is an exception to this scheme, the case when programs are executed on  $N_p = 28$  nodes, as the scaled number of spatial points cannot be evenly distributed among 28 processes. In this case we compare the weak scalability by scaling up a starting mesh of  $476 \times 476 \times 476$ , which has 107,850,176 spatial points, just slightly less than for other values of  $N_p$ . The base of comparison was the execution time of a single iteration of the main loop with the corresponding MPI-based program running as a single process ( $N_p = 1$ ). While we used this configuration as the baseline, we do not recommend running MPI-based programs with  $N_p = 1$  processes, because no special handling of such case has been implemented. This means that MPI-parallelized programs are always transposing the data, which is unnecessary with  $N_p = 1$  as all data is local to the single process. Instead, we recommend using single-node variants of the programs outside cluster environment.

**Table 14. Weak scaling of MPI-based programs. Wall-clock execution times  $T_W$  are given in milliseconds, efficiencies  $E_W$  in percents**

Mesh size	imag3d-mpi		real3d-mpi		imag3d-mpicuda		real3d-mpicuda		imag3d-mpihetero		real3d-mpihetero	
	$T_W$	$E_W$	$T_W$	$E_W$	$T_W$	$E_W$	$T_W$	$E_W$	$T_W$	$E_W$	$T_W$	$E_W$
$N_p = 1$												
480 × 480 × 480	1758	100	2956	100	4338	100	6346	100	1971	100	2877	100
$N_p = 2$												
480 × 480 × 960	3858	45.6	4910	60.2	3785	114.6	5430	116.9	2044	96.4	2878	100
480 × 960 × 480	3704	47.5	5035	58.7	3825	113.4	5573	113.9	1997	98.7	2928	98.3
960 × 480 × 480	3758	46.8	5506	53.7	3871	112.1	5609	113.1	2019	97.6	2980	96.5
$N_p = 4$												
480 × 960 × 960	4267	41.2	6033	49.2	2995	144.8	4278	148.3	2149	91.7	3009	95.6
960 × 480 × 960	3877	45.4	5223	56.6	3042	142.6	4318	146.9	2218	88.9	3274	87.9
960 × 960 × 480	4000	44	5063	58.4	3078	140.9	4424	143.4	2148	91.8	3106	92.6
$N_p = 8$												
960 × 960 × 960	4327	40.6	6285	47	2952	146.9	4167	152.3	2369	83.2	3338	86.2
$N_p = 12$												
1440 × 960 × 960	4400	40	6729	43.9	2902	149.5	4254	149.2	2335	84.4	3504	82.1
960 × 1440 × 960	4268	41.2	6263	47.2	2953	146.9	4167	152.3	2412	81.7	3440	83.6
960 × 960 × 1440	4234	40.2	6299	46.9	2865	151.4	4153	152.8	2413	81.7	3351	85.9
$N_p = 16$												
1920 × 960 × 960	4356	40.4	6419	46	3001	144.5	4278	148.3	2392	82.4	3434	83.8
960 × 1920 × 960	4266	41.2	6376	46.4	2970	146.1	4170	152.2	2552	77.2	3525	81.6
960 × 960 × 1920	4234	40.2	6508	45.4	2888	150.2	4148	153	2639	74.7	3606	79.8
$N_p = 24$												
1440 × 1920 × 960	4677	40	6895	42.9	2924	148.4	4266	148.8	2552	77.3	3770	76.3
1440 × 960 × 1920	4796	38.3	7076	41.8	2984	145.4	4235	149.8	2611	75.5	3699	77.8
1920 × 1440 × 960	4470	40.6	6439	45.9	3007	144.2	4286	148.1	2546	77.4	3504	82.1
1920 × 960 × 1440	4450	40.8	6425	46	2984	145.3	4251	149.3	2551	77.3	3550	81.1
960 × 1440 × 1920	4364	41.1	6846	43.2	2879	150.7	4105	154.6	2682	73.5	3725	77.2
960 × 1920 × 1440	4346	41	6806	43.4	2878	150.7	4133	153.5	2501	78.8	3562	80.8
$N_p = 32$												
1920 × 1920 × 960	4452	39.5	6704	44.1	3183	136.3	4518	140.5	2585	76.3	3715	77.5
1920 × 960 × 1920	4520	38.9	6802	43.4	3163	137.2	4480	141.7	2632	74.9	3769	76.3
960 × 1920 × 1920	4301	40.9	6663	44.4	3014	143.9	4351	145.9	2776	71	3933	73.0

Table 14 lists mesh sizes used and execution times obtained for  $N_p = 1, 2, 4, 8, 12, 16$  and 32 cluster nodes. We also tested weak scaling on  $N_p = 20, 24$  and 28, but have excluded them from the table for brevity. A complete comparison is shown in Figure 15. The ordering of mesh dimensions had a slightly greater impact than with single-node OpenMP programs, varying up to 10% for the OpenMP/MPI programs, and less for the other two MPI-based implementations. This means that for meshes of the same size, thus implying equal work, execution times were mostly lower for those with larger values of  $N_x$  (e.g., the execution time for a  $960 \times 480 \times 960$  mesh is lower than for a  $480 \times 960 \times 960$  one). However, no distinct pattern emerges that would give us a clue as to which order is the most favorable. Upon inspection, we find that the difference in execution times is due to the different FFT plans employed by the FFTW and cuFFT libraries. The communication time remains mostly the same, as is expected since the amount of data exchanged is the same.

The results of weak scaling tests warrant further discussion as the execution times and plotted efficiencies appear misleading. On one hand, OpenMP/MPI programs show poor scaling results, achieving only 40% efficiency, while on the other hand the CUDA/MPI show efficiency well above 100%, a seemingly impossible result. The lower efficiency of the OpenMP/MPI programs is due to the very good performance of the baseline run. Here, the FFTW creates a very good FFT and transpose plans, which exploit the fact that all data are local, and thus have very good execution time. With  $N_p = 2$  nodes the execution time increases significantly, and efficiency drops to about 45-55%. However, adding more work and nodes has much smaller impact on efficiency, which only drops to 40-45% with  $N_p = 32$  nodes. If we use  $N_p = 2$  as the baseline, we see that the efficiency remains above 80%, a very good result.

In contrast to the OpenMP/MPI version, CUDA/MPI version performs very badly when executed on one node. Due to message sizes, MPI implementation used (Open MPI) relies on asynchronous copies through host memory [59] to perform the required communication. This results in bad performance due to a lack of overlap when copying data, as there is only one process involved. With  $N_p = 2$  processes the situation improves, as multiple streams are used, much like how we employ streams to overlap computation and data transfers in implementation of the hybrid algorithm. Between  $N_p = 4$  and 16 processes we get the best performance, and further increasing the number of processes only slightly reduces efficiency due to the size and amount of messages passed. If we compare the weak scaling efficiency with the best result as the baseline

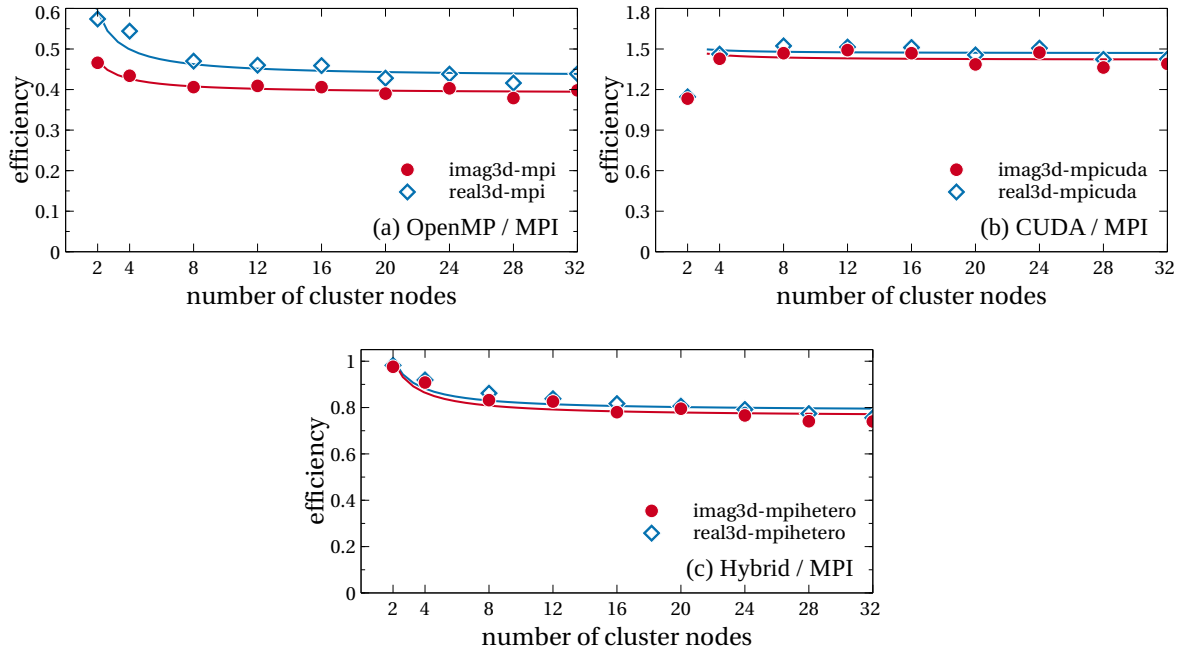


Figure 15. Weak scaling efficiency of three MPI-based implementations, averaged over all mesh sizes tested for each value of  $N_p$ . Solid lines represent fits to measured data, where fit model functions are given in the text.

(obtained with  $N_p = 16$  processes), we see that the efficiency is lower than for the OpenMP/MPI version, due to the different transpose routine used and divided multidimensional FFT that has to be employed. However, we stress that, in terms of absolute execution times, CUDA/MPI is faster than OpenMP/MPI version for  $N_p \geq 4$ .

The Hybrid/MPI version behaves as expected and demonstrates weak scaling efficiency of about 70-75% with  $N_p = 32$  nodes. Since data are transposed only in host memory, this version does not suffer the penalty of memory copies like the CUDA/MPI version, and therefore achieves much better efficiency with small number of cluster nodes. In terms of absolute execution times, Hybrid/MPI version is the fastest of all three implementations for  $N_p \geq 2$  and is therefore the algorithm of choice for distributed memory systems. Of course, it assumes previous optimization of offloading parameters using the GA method, which itself is time-consuming and has to be taken into consideration when making the choice. Table 15 provides data on the optimal fraction of total data offloaded to GPU in our weak scaling tests. As the number of processes increases, one expects that the amount of data offloaded to GPU remains constant, which is the case, as we see from the table. The only exceptions are when the

**Table 15. Optimal fraction of total data offloaded to GPU in weak scaling tests for different mesh sizes in Hybrid/MPI programs, obtained using the GA optimization method**

Mesh size	imag3d-mpihetero				real3d-mpihetero			
	PS <sub>1</sub>	PS <sub>2</sub>	PS <sub>3</sub>	PS <sub>4</sub>	PS <sub>1</sub>	PS <sub>2</sub>	PS <sub>3</sub>	PS <sub>4</sub>
$N_p = 1$								
480 × 480 × 480	35	20	31	44	46	30	40	35
$N_p = 2$								
480 × 480 × 960	46	33	33	47	38	29	37	38
480 × 960 × 480	50	67	17	47	45	60	19	40
960 × 480 × 480	50	15	70	34	43	15	67	35
$N_p = 4$								
480 × 960 × 960	40	65	18	47	40	63	17	37
960 × 480 × 960	50	10	67	42	40	12	75	33
960 × 960 × 480	33	29	33	40	45	22	33	40
$N_p = 8$								
960 × 960 × 960	53	25	33	47	40	27	40	28
$N_p = 12$								
1440 × 960 × 960	37	20	53	45	42	18	60	63
960 × 1440 × 960	40	38	27	47	45	35	27	35
960 × 960 × 1440	50	25	38	45	45	25	40	40
$N_p = 16$								
1920 × 960 × 960	47	13	67	27	42	13	80	30
960 × 1920 × 960	47	47	35	37	47	53	20	25
960 × 960 × 1920	50	20	33	33	47	20	40	27
$N_p = 24$								
1440 × 1920 × 960	47	40	30	53	47	33	25	63
1440 × 960 × 1920	50	17	60	45	47	13	60	70
1920 × 1440 × 960	40	20	50	40	45	20	53	27
1920 × 960 × 1440	53	13	60	45	45	10	80	25
960 × 1440 × 1920	50	30	27	47	45	30	27	27
960 × 1920 × 1440	40	50	18	45	45	50	20	30
$N_p = 32$								
1920 × 1920 × 960	47	27	33	30	47	17	40	27
1920 × 960 × 1920	50	13	60	40	47	10	80	40
960 × 1920 × 1920	47	47	40	33	47	40	23	27



value of  $N_Y$  is greater than  $N_X$  where we see a larger fraction of data offloaded to GPU in  $PS_2$  and  $PS_4$ , and when the value of  $N_Y$  is smaller than  $N_X$  where we see a smaller fraction of data offloaded to GPU in  $PS_2$  and  $PS_4$ .

To model the obtained weak scaling results, we start from equation (39) for the execution time of a single iteration of the main loop. In weak scaling tests, the mesh size is increased proportionally to the number of MPI processes  $N_p$ , and therefore the execution time is given by

$$T_W(N_p) = \alpha + \beta L + \gamma \frac{MN_p}{N_p} + \beta V \frac{MN_p}{N_p^2} = \alpha + \beta L + \gamma M + \beta V \frac{M}{N_p}. \quad (42)$$

The weak scaling efficiency

$$E_W(N_p) = \frac{T_W(1)}{T_W(N_p)}, \quad (43)$$

can thus be modeled by

$$E_W(N_p) = \frac{1}{a + b/N_p}. \quad (44)$$

We fitted this model to data presented in Figure 15 and the obtained fit parameters are given in Table 16. Note that the model is fitted to CUDA/MPI scaling data only for  $N_p \geq 4$ . The agreement between the model and measurement data is very good, although we note that the parameter  $b$  has negative values, contrary to its expected relation to the communication cost. Therefore, we conclude that the above expression can be successfully used for modeling of the performance of all three MPI-based implementations.

To summarize, in all test results of MPI-based programs, we note an expected declining of the scaling efficiency. This is due to the introduction of distributed transposes of data, creating overhead that negatively impacts scaling efficiency. It is most evident in both strong and weak scaling tests of the CUDA/MPI version, as the transpose algorithm is inferior to the one provided by FFTW, used in OpenMP/MPI implementation. In our tests, all three MPI versions of programs failed to achieve actual speedup ( $S(N_p) > 1$ ) on less than  $N_p = 4$  nodes, due to the introduction of these transpose routines. We therefore recommend using MPI versions only on  $N_p = 4$  or more cluster nodes.

**Table 16. Values of obtained weak scaling model fit parameters  $a$  and  $b$ , as well as the estimated fit errors for MPI-parallelized programs**

Program	$a$	$\Delta a$	$b$	$\Delta b$
imag3d-mpi	2.56	0.06	-0.9	0.3
real3d-mpi	2.32	0.05	-1.3	0.2
imag3d-mpicuda	0.70	0.02	-0.1	0.2
real3d-mpicuda	0.68	0.02	-0.1	0.2
imag3d-mpihetero	1.32	0.03	-0.6	0.2
real3d-mpihetero	1.27	0.03	-0.6	0.2

### 2.4.5. Selecting the Optimal Algorithm

In previous sections we presented results of detailed performance tests and associated models for all developed programs. Here we provide general guidelines for obtaining the best performance from each implementation.

We note that the extensive testing performed shows that the best performance can be achieved by evenly distributing the workload among the MPI processes and OpenMP threads, and by using mesh sizes which are optimal for FFT. In particular, the single-node OpenMP programs have the best performance if the number of spatial points in all directions, controlled by parameters  $N_x$ ,  $N_y$  and  $N_z$ , is divisible by the number of OpenMP threads used. Similarly, the OpenMP/MPI implementation achieves the best performance if  $N_x$  and  $N_y$  are divisible by a product of the number of MPI processes and the number of OpenMP threads used.

On the other hand, CUDA implementation works best if all the data for a given mesh size can fit into the GPU memory. In the case of a small mesh size, CUDA programs may not be able to saturate all Streaming Multiprocessors (SM) of a given GPU, and in this case the OpenMP programs may be considered first. For CUDA/MPI programs, the best performance is achieved if  $N_x$  and  $N_y$  are divisible by a product of the number of MPI processes and the number of SMs in the GPU used. Note that CUDA-based implementations write output files (e.g., density profiles) by transferring data from GPU memory to host memory, where a single thread writes to a file. Therefore, these implementations may not be suitable for simulations that require the output to be written frequently, after a small number of time propagation steps.

Best performance of hybrid implementations can be obtained by following

the same guidelines. Furthermore, we recommend using these programs with the optimization methods described in section 2.4.1 to optimally divide the work between host and device. If manual tuning of offload parameters is required, we recommend that the guidelines above are followed for both CPU and GPU portions of the mesh. In the case of large disparity in the performance of host or device, hybrid versions will not provide the lowest execution times, and in such cases the pure CPU or GPU implementations with or without MPI could be better suited.

In addition to the guidelines for mesh sizes presented above, all programs benefit from the mesh size which is also optimal for FFT. The best FFT performance on CPU with FFTW library is obtained if  $N_x$ ,  $N_y$  and  $N_z$  can be expressed as  $2^a 3^b 5^c 7^d 11^e 13^f$ , where  $e$  and  $f$  are either 0 or 1, and the other exponents are non-negative integer numbers [60]. Similarly, for cuFFT the best performance is achieved for transform sizes of the form  $2^a 3^b 5^c 7^d$ . In hybrid implementations, the same applies to the host and device portions of  $N_x$  and  $N_y$ , i.e.,  $\text{cpu}N_x$ ,  $\text{cpu}N_y$ ,  $\text{gpu}N_x$  and  $\text{gpu}N_y$ .

## 2.5. Applications

The applicability and accuracy of the NLS equation for the description of BEC properties has been established in early studies addressing time-of-flight expansion and collective modes of the condensate [61, 1]. More recently, numerical algorithms [22, 23, 25, 26] for solving the NLS equation (1), have been used to investigate: nonlinear frequency shifts of the collective mode induced by a harmonic modulation of interaction [62]; onset and features of Faraday waves in BEC [63, 64, 65]; phase-separation dynamics of two-component BECs [66]; geometric resonances in a BEC with two- and three-body interactions [67]; non-equilibrium dynamics of a stirred BEC [68, 69, 70]; localization of a two-component BEC in a random potential [71]; vortex lattice melting in the presence of impurities [72]; classification of fractional vortices in a spinor BEC [73]; BEC dynamics in a dimple trap [74, 75, 76]; dynamics of localized impurity in BEC [77]; analysis of localized states in a multi-component BEC [78]; BEC dynamics induced by a sudden change of an optical lattice geometry [79]; BEC dynamics in a gravito-optical surface trap [80]; emergence of a Bose glass phase in a BEC in the presence of disorder [81]; vortex dynamics [82]; vortex generation in a stirred BEC [83]; dynamics of a self-bound droplet in the presence of an attractive two-body and a very small repulsive three-body interaction

[84, 85]. The dynamics of light bullets in a cubic–quintic nonlinear medium based on an extension of equation (12) has been published in reference [86]. As an illustration, in section 3 we present results of the study addressing the phase–separation dynamics of two–component BECs [66].

Furthermore, the codes [22, 23] have been extended to incorporate effects of the spin–orbit coupling which are currently intensely explored in the systems of cold atoms. In particular, the following subjects have been investigated: phase separation in a spin–orbit coupled BEC [87]; soliton dynamics in spin–orbit coupled BEC [88, 89, 90, 91]; dynamics of a BEC with spin–angular momentum coupling [92]. In section 4 we present recently obtained results about dipole mode and breathing mode excitations of a BEC in the presence of the coupling of pseudo spin and angular momentum [66].

The algorithms [28, 29, 30] for solving the dipolar NLS equation (11) have been applied to the study of: vortex dynamics in a dipolar BEC [93, 94, 95, 96]; phase separation in a binary dipolar BEC [97]; dynamics of a dipolar BEC droplet [98]; stability and dynamics of a dipolar Bose–Fermi mixture [99, 100]; solitons in a dipolar BEC [101, 102, 103, 104, 105, 106, 107]; self–trapping of a dipolar BEC in a double–well [108]. As a final illustration, in section 5 we present a study of vortex formation in a dipolar BEC.

### 3. Phase Separation Dynamics of a Two–Component BEC

Systems of multi–component BECs involve an additional degree of freedom and enrich the physics that can be explored within cold atom setups. These systems are either realized as mixtures of different bosonic atoms or different hyperfine states of the same atomic species. In this section we extend the previously described numerical algorithm and apply it for a study of a two–component system, described by two condensate wave functions  $\psi_1(x, t)$  and  $\psi_2(x, t)$  in a quasi one–dimensional geometry. The ground state configuration can be obtained as a minimum of the following energy functional

$$E_0 = \int dx \left[ (\psi_1^*(x) \ \psi_2^*(x)) H_0 \begin{pmatrix} \psi_1(x) \\ \psi_2(x) \end{pmatrix} + \frac{g_{11}}{2} |\psi_1(x)|^4 + \frac{g_{22}}{2} |\psi_2(x)|^4 + g_{12} |\psi_1(x)|^2 |\psi_2(x)|^2 \right], \quad (45)$$

where  $H_0$  corresponds to the non-interacting Hamiltonian

$$H_0 = \left( -\frac{1}{2} \frac{\partial^2}{\partial x^2} + V(x) \right) \mathcal{I}_2, \quad (46)$$

with  $\mathcal{I}_2$  being a  $2 \times 2$  identity matrix. As the total number of atoms of each species is conserved, the normalization condition is implemented as  $\int dx |\psi_1(x)|^2 = \int dx |\psi_2(x)|^2 = 1$ . The expression (45) encompasses three interaction constants:  $g_{11}$  for atoms of the first species,  $g_{22}$  for atoms of the second species and  $g_{12}$  interaction constant between the atoms of the two species. We assume that all these interactions are repulsive, i.e.,  $g_{11} \geq 0$ ,  $g_{22} \geq 0$ ,  $g_{12} \geq 0$ . For  $g_{12} = 0$ , the two species do not "feel" each other and energy defined in (45) consists of two copies of single-component energy (5). In this case density distributions of the two species  $|\psi_1(x)|^2$  and  $|\psi_2(x)|^2$  may fully overlap in space. As intra-component repulsion  $g_{12}$  gets stronger, two-component mixtures may exhibit phase separation [13]. Simply, as it is energetically costly for the two species to overlap, they spatially separate instead. In the homogeneous case  $V(x) = 0$ , the necessary condition for the onset of phase separation is given by [109]

$$g_{11}g_{22} < g_{12}^2. \quad (47)$$

Often the two-component condensate is not prepared as the ground state of (45), but using a dynamical protocol [110, 111, 112] that we describe in the following. The dynamics of a two-component condensate in the mean-field regime is governed by two coupled Gross-Pitaevskii equations [113]:

$$i \frac{\partial \psi_1(x, t)}{\partial t} = \left[ -\frac{1}{2} \frac{\partial^2}{\partial x^2} + \frac{1}{2} \lambda^2 x^2 + g_{11} |\psi_1(x, t)|^2 + g_{12} |\psi_2(x, t)|^2 \right] \psi_1(x, t), \quad (48)$$

$$i \frac{\partial \psi_2(x, t)}{\partial t} = \left[ -\frac{1}{2} \frac{\partial^2}{\partial x^2} + \frac{1}{2} \lambda^2 x^2 + g_{12} |\psi_1(x, t)|^2 + g_{22} |\psi_2(x, t)|^2 \right] \psi_2(x, t). \quad (49)$$

Wave functions of the condensates  $\psi_1(x, t)$  and  $\psi_2(x, t)$  are normalized to unity and  $\lambda$  is the strength of the harmonic trap  $V(x) = \lambda^2 x^2 / 2$ . The non-equilibrium dynamics of a two-component BEC captured by the coupled NLS equations (48) and (49) has attracted lot of theoretical interest [114, 115, 116, 117, 118].

Initially, a single-component condensate is produced. We model this situation as a two-component BEC with  $g_{11} = g_{22} = g_{12}$  and prepare initial state as the ground state of (45) for  $g_{11} = g_{22} = g_{12}$ . At  $t = 0$  a pulse is applied and half

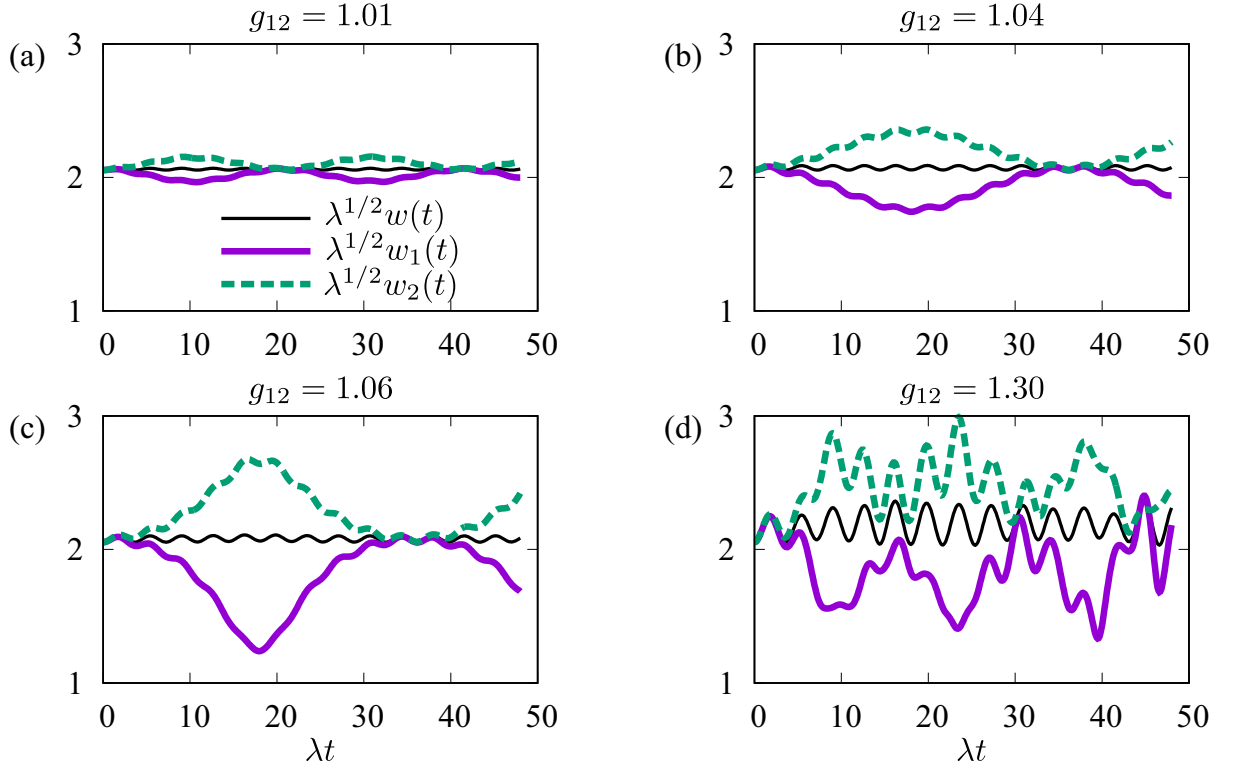


Figure 16. Time evolution of condensate widths for several values of  $g_{12}$ : (a)  $g_{12} = 1.01$ , (b)  $g_{12} = 1.04$ , (c)  $g_{12} = 1.06$ , (d)  $g_{12} = 1.3$ . Other parameters:  $\lambda = 10^{-3}$ ,  $g_{11} = 1$ ,  $g_{22} = 1.01$ .

the atoms are transferred into another hyperfine state. An important assumption of our analysis is that the process is completely spatially symmetric, i.e. the two BECs overlap perfectly at the initial moment. Due to the transfer, values of  $g_{22}$  and  $g_{12}$  are suddenly changed and dynamics of the two components is started. We monitor the time dependence of the widths of the two components,

$$w_{1,2}^2(t) = \int_{-\infty}^{\infty} dx x^2 |\psi_{1,2}(x, t)|^2, \quad (50)$$

as well as the total width  $w(t) = \sqrt{(w_1(t)^2 + w_2(t)^2)/2}$ . Typical results obtained for different final values of  $g_{12}$  and  $\lambda = 10^{-3}$ ,  $g_{11} = 1$ ,  $g_{22} = 1.01$  are presented in Figure 16. We observe fast oscillations in the total width with a frequency unaffected by a change in  $g_{12}$ . These oscillations correspond to the so-called breathing mode. In a quasi one-dimensional geometry it turns out that their period is  $2\pi/(\sqrt{3}\lambda)$ . On top of this, we notice slow oscillations in the width of each of the two component for  $g_{12} = 1.01$  as shown in Figure 16(a).

These out-of-phase oscillations correspond to the "spin" mode, and their frequency decreases with increase in  $g_{12}$ . In particular, the oscillation period of the spin mode is about 20 for  $g_{12} = 1.01$  (see Figure 16(a)) and about 30 for  $g_{12} = 1.04$  (see Figure 16(b)). Above a certain threshold value of  $g_{12}$ , oscillation amplitudes sharply increase and the width dynamics becomes aperiodic and irregular (see Figure 16(c) for  $g_{12} = 1.06$ ). Irregularity of the width dynamics is even more apparent for  $g_{12} = 1.3$  as shown in Figure 16(d).

In Figure 17 we show the dynamics of the "spin density"  $|\psi_1(x, t)|^2 - |\psi_2(x, t)|^2$  for the same parameters as in Figure 16. The plots Figure 17(a) and (b) show regular oscillations. By comparing plots Figure 17(b) and (c) we note a sharp change in the color scales that matches the sharp increase in the oscillation amplitude observed earlier. Finally, in Figure 17(d) we observe a clear onset of pattern dynamics. By applying a suitable linear stability analysis, it was shown that the onset of aperiodic motion is marked by emergence of an unstable mode around a referent stationary state [66]. This behavior is known as modulation instability and by investigating a range of values  $\lambda \in (10^{-5}, 10^{-1})$ , we found that in a realistic experimental situation, stronger trapping potential strength  $\lambda$  shifts the instability condition (47) toward higher values of  $g_{12}$ .

## 4. Excitations of a BEC with Spin-Angular Momentum Coupling

Experimental realization of an effective spin-orbit coupling in ultracold atom systems [119, 120, 121, 122, 123, 124, 125] has allowed for new quantum phases to be explored. Bosonic systems with spin-orbit coupling are interesting as they have no direct analogues in condensed matter systems and provide a new research playground [126, 127]. In recent references [128, 129, 130, 131] the following Hamiltonian for a two-component bosonic system has been introduced:

$$H_{\text{SOC}} = \left( -\frac{1}{2}\Delta + \frac{r^2}{2} \right) \mathcal{I}_2 + \frac{\Omega^2 r^2}{2} \begin{pmatrix} 1 & e^{-2i\phi} \\ e^{2i\phi} & 1 \end{pmatrix}, \quad (51)$$

where  $\mathcal{I}_2$  is a  $2 \times 2$  identity matrix and the effective spin 1/2 comes from the two bosonic components involved. The system is assumed to be effectively two-dimensional (tightly trapped in the longitudinal direction). The value of  $\Omega$  is set by the intensity of two co-propagating Laguerre-Gauss laser beams




Figure 17. Spin dynamics illustrated via the density difference given by  $\lambda^{-1/2} (|\psi_1(x, t)|^2 - |\psi_2(x, t)|^2)$ , induced by a quench of  $g_{12}$  and  $g_{22}$ . Parameters: (a)  $g_{12} = 1.01$ , (b)  $g_{12} = 1.04$ , (c)  $g_{12} = 1.06$ , (d)  $g_{12} = 1.3$  and  $\lambda = 10^{-3}$ ,  $g_{11} = 1$ ,  $g_{22} = 1.01$ .

that carry a unit of angular momentum in the opposite rotational directions. The last,  $\phi$ -dependent term, where  $\phi$  is the polar angle, provides the coupling between the spin and angular momentum, as can be explicated by using a proper unitary transformation [129]. In this section we extend the algorithms presented in subsection 2.2 and use them to investigate excitations of a BEC with spin-angular momentum coupling as defined in equation (51).

The total energy per particle of the condensed state with the order parameter  $(\psi_1(\mathbf{r}) \ \psi_2(\mathbf{r}))^T$  is given by

$$E_0 = \int d\mathbf{r} \left[ (\psi_1^* \ \psi_2^*) H_{\text{SOC}} \begin{pmatrix} \psi_1 \\ \psi_2 \end{pmatrix} + \frac{g}{2} |\psi_1|^4 + \frac{g}{2} |\psi_2|^4 + g |\psi_1|^2 |\psi_2|^2 \right], \quad (52)$$

where we assume spin-symmetric interactions  $g_{11} = g_{22} = g_{12} = g$ . Due to off-diagonal terms in (51), which allow for the conversion of one atomic state into the other, only the total number of atoms is conserved and the normalization condition is given by  $\int d\mathbf{r} (|\psi_1(\mathbf{r})|^2 + |\psi_2(\mathbf{r})|^2) = 1$ .



To implement the spin–angular momentum coupling in the propagation routine, we split the step of imaginary–time propagation into four pieces. As usual, in the first step we take into account harmonic potential and nonlinear terms, according to equation (13). The second step comprises the spin–angular momentum coupling, and in the last two steps we perform integrations of the type (14) with respect to the two spatial variables  $x$  and  $y$ . Explicitly, the off–diagonal terms from (51) are treated separately starting from

$$-\frac{\partial\psi_1(\mathbf{r}, \tau)}{\partial\tau} = e^{-2i\phi} c(r)\psi_2(\mathbf{r}, \tau), \quad (53)$$

$$-\frac{\partial\psi_2(\mathbf{r}, \tau)}{\partial\tau} = e^{2i\phi} c(r)\psi_1(\mathbf{r}, \tau), \quad (54)$$

where  $c(r) = \frac{1}{2}\Omega^2 r^2$ . By differentiating the two last equations we obtain two independent equations

$$\frac{\partial^2\psi_1(\mathbf{r}, \tau)}{\partial\tau^2} = c(r)^2\psi_1(\mathbf{r}, \tau), \quad \frac{\partial^2\psi_2(\mathbf{r}, \tau)}{\partial\tau^2} = c(r)^2\psi_2(\mathbf{r}, \tau) \quad (55)$$

that are solved by

$$\begin{aligned} \psi_1(\mathbf{r}, \tau) &= A_1(\mathbf{r}) e^{-c(r)\tau} + B_1(\mathbf{r}) e^{c(r)\tau}, \\ \psi_2(\mathbf{r}, \tau) &= A_2(\mathbf{r}) e^{-c(r)\tau} + B_2(\mathbf{r}) e^{c(r)\tau}. \end{aligned} \quad (56)$$

From the initial conditions for the second part of the  $n$ –th integration step at  $\tau' = (n + \frac{1}{4}) \varepsilon$  we obtain a set of algebraic equations for the coefficients

$$\begin{aligned} A_1(\mathbf{r}) + B_1(\mathbf{r}) &= \psi_1(\mathbf{r}, \tau'), \\ -c(r)A_1(\mathbf{r}) + c(r)B_1(\mathbf{r}) &= -e^{-2i\phi} c(r) \psi_2(\mathbf{r}, \tau'), \\ A_2(\mathbf{r}) + B_2(\mathbf{r}) &= \psi_2(\mathbf{r}, \tau'), \\ -c(r)A_2(\mathbf{r}) + c(r)B_2(\mathbf{r}) &= -e^{2i\phi} c(r) \psi_1(\mathbf{r}, \tau'), \end{aligned}$$

that finally leads to

$$\begin{aligned} A_1(\mathbf{r}) &= \frac{1}{2} \left[ \psi_1(\mathbf{r}, \tau') + e^{-2i\phi} \psi_2(\mathbf{r}, \tau') \right], \\ B_1(\mathbf{r}) &= \frac{1}{2} \left[ \psi_1(\mathbf{r}, \tau') - e^{-2i\phi} \psi_2(\mathbf{r}, \tau') \right], \\ A_2(\mathbf{r}) &= \frac{1}{2} \left[ \psi_2(\mathbf{r}, \tau') + e^{2i\phi} \psi_1(\mathbf{r}, \tau') \right], \\ B_2(\mathbf{r}) &= \frac{1}{2} \left[ \psi_2(\mathbf{r}, \tau') - e^{2i\phi} \psi_1(\mathbf{r}, \tau') \right]. \end{aligned} \quad (57)$$

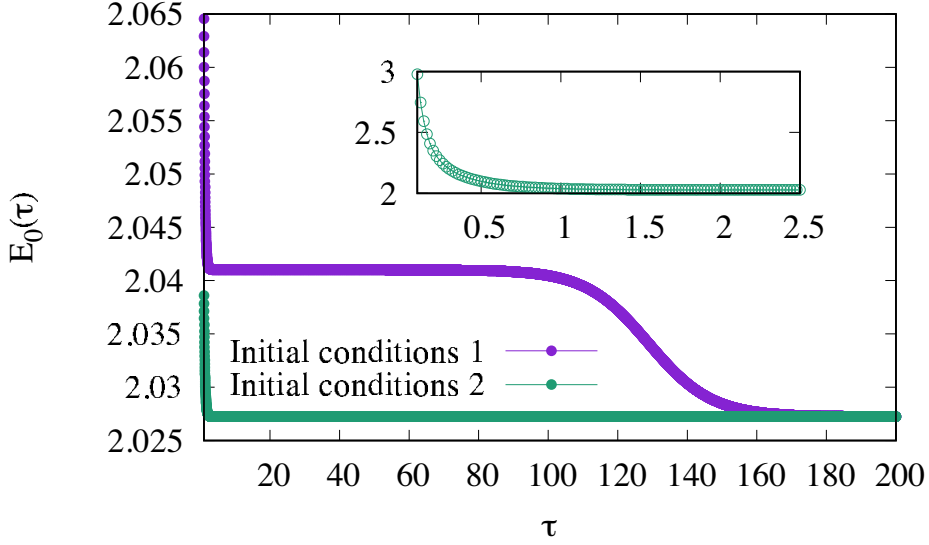


Figure 18. The energy of the ground state  $E_0(\tau)$  defined in equation (52) vs. imaginary time  $\tau$  obtained for different initial conditions. In the first case, we start from a random initial state and observe slow convergence toward the global minimum. In the second case, initial conditions have proper dependence on the polar angle  $\phi$ , as given in equation (58) and the convergence toward the ground state is exponentially fast as shown in the inset. Parameters:  $\Omega = 3.2$  and  $g = 1$ .

In our numerical implementation we make direct use of equations (56) and (57).

Before turning to numerical simulations, we first briefly present an important analytical result available for  $g = 0$ . General symmetry considerations imply that the ground state configurations of (52) for  $g = 0$  can be expressed in a form

$$\phi_m(r, \phi) = \frac{e^{im\phi}}{\sqrt{2\pi}} \begin{pmatrix} f_m(r)e^{-i\phi} \\ g_m(r)e^{i\phi} \end{pmatrix}, \quad (58)$$

where  $m$  is an integer. In particular for  $\Omega < \Omega_c$ , with  $\Omega_c = 3.35$  there are two degenerate ground states  $m = \pm 1$  and for  $\Omega > \Omega_c$  the ground state configuration corresponds to  $m = 0$ . The  $m = \pm 1$  states are denoted as half-vortices, and  $m = 0$  state is named vortex-antivortex pair. In order to explore ground state configurations for a finite value of interaction constant  $g$ , we perform minimization of energy functional (52) using imaginary time propagation. How quickly we reach the ground state strongly depends on the choice of initial configuration as shown in Figure 18. Starting with random initial values, we find

that propagation routine may first reach a metastable configuration (local minimum of  $E_0$ ) and eventually, only after very long propagation time, it selects the right minimum. With a better choice of initial configuration, motivated by exact results (58), the right minimum is found with substantially less iterations. Considerations of this kind are relevant whenever there are several configurations that are nearly degenerate – i.e., have close values of the corresponding energies (52) and represent local minima of the functional (52).




Figure 19. The  $m = 1$  ground state obtained for  $\Omega = 3.2$ ,  $g = 1$ . We plot: (a) the density distribution of the first component  $|\psi_1(\mathbf{r})|^2$ , (b) the density distribution of the second component  $|\psi_2(\mathbf{r})|^2$ , (c) the argument of  $\psi_1(\mathbf{r})$  and (d) the argument of  $\psi_2(\mathbf{r})$ .

Two examples of the ground state configurations are given in Figs. 19 and 20. We find that at finite values of  $g$ , ground states can also be distinguished according to their dependence on polar angle  $\phi$ , as given in equation (58). In Figure 19 we present the ground state solution for  $\Omega = 3.2$  and  $g = 1$ . This is a half-skyrmion configuration characterized by  $m = 1$ . We note that the

first component is confined around the trap center and its phase dependence is trivial. In contrast, the second bosonic species exhibits double phase winding, see Figure 19(d). In Figure 20 we present the ground state solution for  $\Omega = 3.5$  and  $g = 1$ . The densities of the two components overlap in space, Figure 20(a) and (b), and exhibit phase winding in opposite directions, Figure 20 (c) and (d). This is an example of  $m = 0$  state called a vortex–antivortex pair.

In general, quantum phases can be distinguished according to their low-lying spectra. In cold atom experiments routinely accessible with great precision are breathing mode and dipole mode excitations. To probe these modes, initially the system is prepared in its ground state and then a sudden quench of the potential  $V(\mathbf{r})$  is performed. We now theoretically address these excitations for the half-skyrmion and vortex–antivortex state.




Figure 20. The  $m = 0$  ground state obtained for  $\Omega = 3.5$ ,  $g = 1$ . We plot: (a) the density distribution of the first component  $|\psi_1(\mathbf{r})|^2$ , (b) the density distribution of the second component  $|\psi_2(\mathbf{r})|^2$ , (c) the argument of  $\psi_1(\mathbf{r})$  and (d) the argument of  $\psi_2(\mathbf{r})$ .

The dynamics of a BEC with spin–angular momentum coupling is governed by two coupled time–dependent Gross–Pitaevskii equations for  $\psi_1(\mathbf{r}, t)$  and  $\psi_2(\mathbf{r}, t)$ :

$$i\frac{\partial\psi_1}{\partial t} = \left[ -\frac{1}{2} \left( \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \right) + V(r) + \frac{r^2\Omega^2}{2} + g(|\psi_1|^2 + |\psi_2|^2) \right] \psi_1 + \frac{r^2\Omega^2}{2} e^{-2i\phi} \psi_2, \quad (59)$$

$$i\frac{\partial\psi_2}{\partial t} = \left[ -\frac{1}{2} \left( \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \right) + V(r) + \frac{r^2\Omega^2}{2} + g(|\psi_2|^2 + |\psi_1|^2) \right] \psi_2 + \frac{r^2\Omega^2}{2} e^{2i\phi} \psi_1. \quad (60)$$

The trapping potential is assumed to be harmonic, initially set to  $V(r) = r^2/2$ . To induce a breathing mode, we perturb the trap strength

$$V_{\text{pert}} = (1 + \eta) \frac{r^2}{2}, \quad (61)$$

and calculate the time evolution of the width of the probability distribution,

$$\langle r^2(t) \rangle = \int_0^{2\pi} d\phi \int_0^\infty dr r^3 [|\psi_1(\mathbf{r}, t)|^2 + |\psi_2(\mathbf{r}, t)|^2]. \quad (62)$$

By inspecting the corresponding Fourier transform, we identify the lowest excited frequency as the breathing mode  $\omega_B$ .

In a typical harmonically trapped two–dimensional system, without the spin–orbit coupling  $\Omega = 0$ , the Gross–Pitaevskii description predicts that the breathing mode frequency is  $\omega_B = 2$  (in the units of the harmonic trap frequency), independent of interaction constant  $g$  [132]. Analytical arguments and numerical results (not presented here) show that this general conclusion holds true for the vortex–antivortex state. On the other hand, the breathing mode frequency of a half–skyrmion state depends on both the coupling  $\Omega$  and on the interaction strength  $g$ . Illustrative results for  $\langle r^2(t) \rangle$  are presented in Figure 21(a) for  $\Omega = 3.2$ ,  $\eta = 0.1$  and several values of  $g$ . From the related Fourier transforms shown in Figure 21(b), we find that the breathing mode frequency increases for several percent with  $g$ .

To excite a dipole mode, we consider a shift of the trap bottom in  $x$  direction,

$$V_{\text{pert}} = \frac{r^2}{2} - \delta x \frac{r}{2} \left( e^{i\phi} + e^{-i\phi} \right), \quad (63)$$

and monitor the motion of the center of mass of the system in that direction,

$$\langle x(t) \rangle = \int_0^{2\pi} d\phi \frac{e^{i\phi} + e^{-i\phi}}{2} \int_0^\infty dr r^2 [|\psi_1(\mathbf{r}, t)|^2 + |\psi_2(\mathbf{r}, t)|^2], \quad (64)$$

as well as  $\langle y(t) \rangle$ .

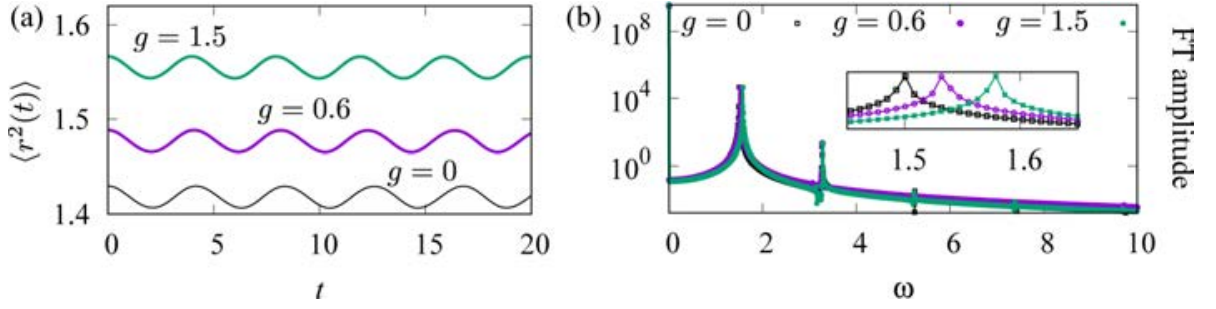


Figure 21. Breathing mode oscillations in half-skyrmion phase: (a)  $\langle r^2(t) \rangle$  versus  $t$  and (b) corresponding Fourier transform. From the inset we observe increase of the breathing mode frequency with  $g$ . Motion is induced by changing harmonic trap potential as  $\frac{r^2}{2} \rightarrow 1.01\frac{r^2}{2}$ ,  $\Omega = 3.2$ .

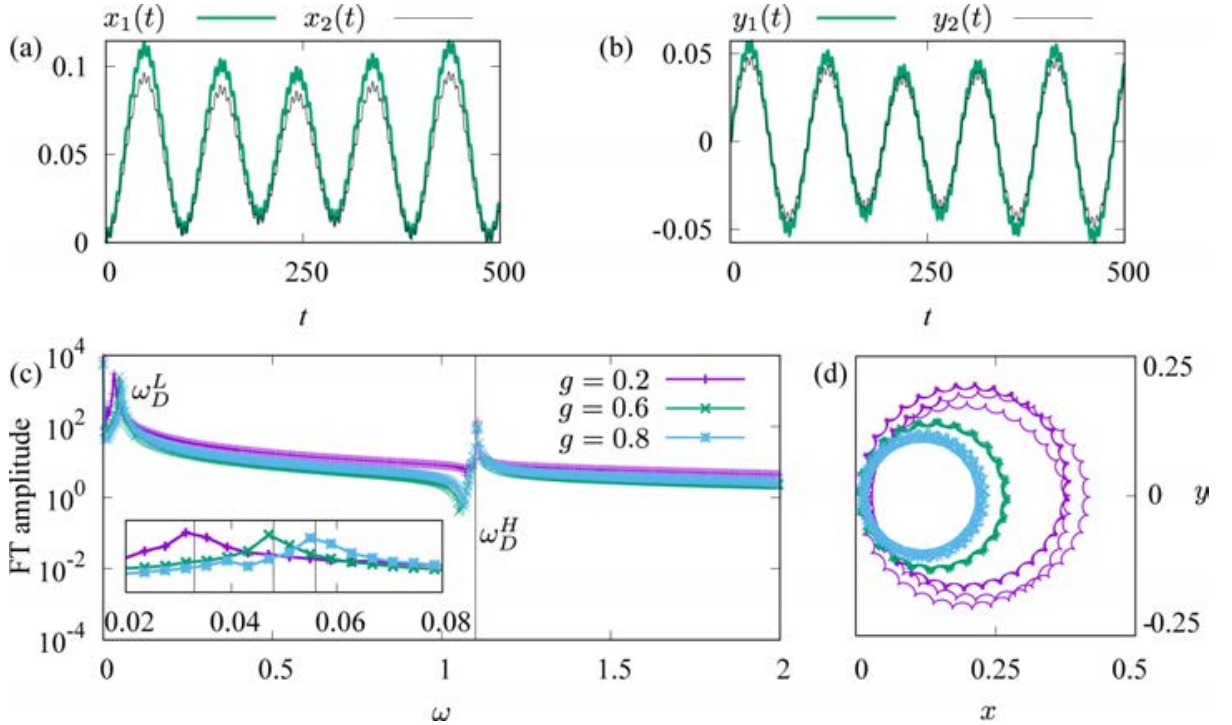


Figure 22. Dipole mode oscillations of half-skyrmion state in interacting case for  $\Omega = 3.2$ . Motion is induced by shifting the harmonic trap bottom for  $\delta x = 0.01$ . In (a) and (b)  $g = 1$ . In (c) motion of the center of mass,  $y(t)$  versus  $x(t)$ , is plotted. The trajectory radius gets smaller with increasing  $g$ . In (d) vertical lines give results for  $\omega_D^L$  (in the inset) and  $\omega_D^H$  (in the main panel) obtained using the Bogoliubov method, and dots represent Fourier transform of  $x(t)$ .

In the conventional harmonically trapped system ( $\Omega = 0$ ), this perturbation excites the Kohn mode – an oscillation with the trap frequency along  $x$

axis unaffected by interactions. Illustrative results for a half-skyrmion state are presented in Figure 22 for  $\Omega = 3.2$  and they obviously exhibit more complex behavior in comparison to the results for  $\Omega = 0$ . We observe that oscillatory motion in  $x$  direction is coupled to the motion along  $y$  direction. Moreover, we find that there are two frequencies involved. We label them as  $\omega_D^L$  (low) and  $\omega_D^H$  (high). In Figure 22(a) and 22(b) we see that two bosonic components oscillate in-phase in both directions. The frequency  $\omega_D^L$  exhibits a strong increase with  $g$ , as is depicted in Figure 22(c). As the frequency  $\omega_D^L$  gets larger, the induced oscillation amplitude gets weaker, Figure 22(d). For the considered case, the frequency  $\omega_D^H$  is found to be almost independent of  $g$ , see Figure 22(c). Results based on the complementary Bogoliubov approach, match quite well to the numerical data obtained from direct numerical simulations of equations (59)–(60), see Figure 22(c).

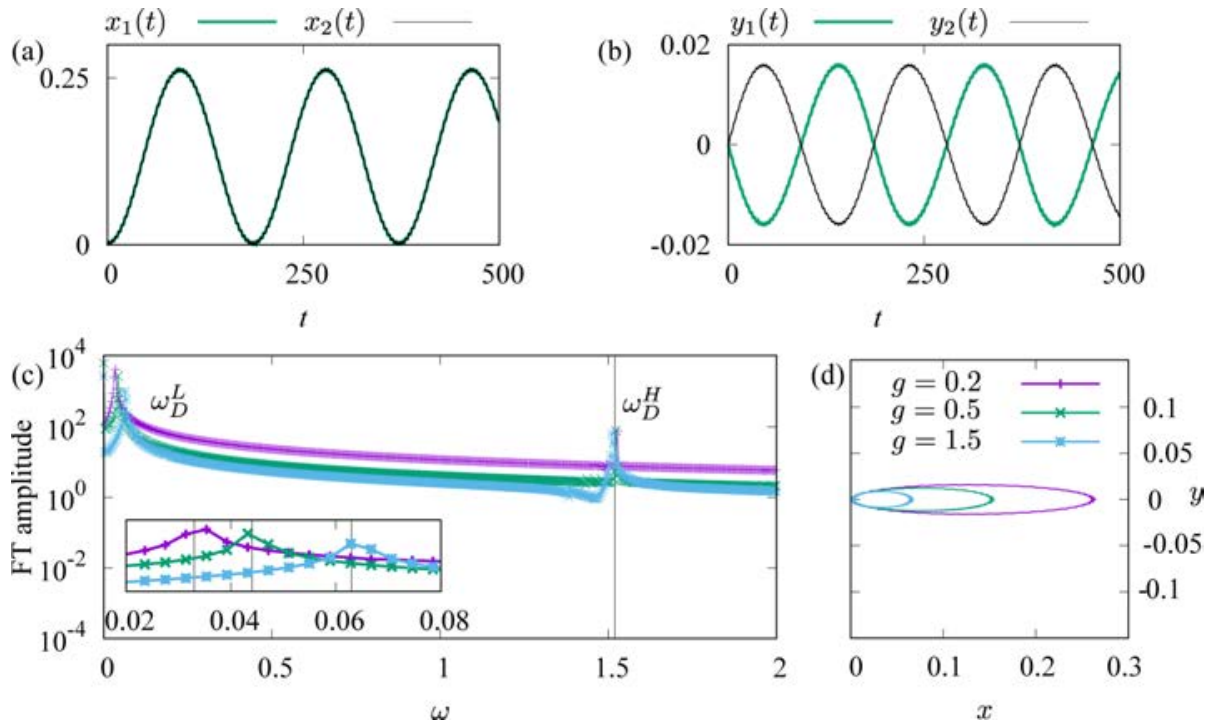


Figure 23. Dipole mode oscillations of the  $m = 0$  solution in interacting case for  $\Omega = 3.5$ . Motion is induced by shifting harmonic trap bottom for  $\delta x = 0.01$ . In (a) and (b)  $g = 0.2$ . In (c) trajectory of the center of mass of a single bosonic component,  $y_1(t)$  versus  $x_1(t)$ , is plotted. In (d) vertical lines give results for  $\omega_D^L$  (in the inset) and  $\omega_D^H$  (in the main panel) obtained using Bogoliubov method and dots represent Fourier transform of  $x(t)$ .



In the vortex–antivortex state, the two bosonic components exhibit an out-of-phase oscillation in  $y$  direction, see Figure 23(b), and consequently the center of mass only oscillates in  $x$  direction with the frequency  $\omega_D^L$  that exhibits an increase with  $g$ , see Figure 23(c). The trajectory of the center of mass of each of the components is given by an ellipse, which is strongly elongated in  $x$  direction, see Figure 23(d).

## 5. Study of Critical Velocity for Vortex Formation

As we have already seen, the programs presented in this chapter can be used to model and study a variety of systems. In this section we demonstrate the versatility of developed programs on a simulation of vortex formation in a BEC. We use the MPI-based programs to simulate the effects of a moving obstacle in an oblate atomic BEC. The obstacle, a repulsive Gaussian laser beam, moves through the condensate and sheds quantum vortices, elementary excitation of a superfluid [14, 13]. The vortices appear only when the obstacle is moving above some critical speed, consistent with the Landau’s criterion of superfluidity. At low obstacle velocities above a critical value, vortex dipoles emerge, and as the speed is increased further, individual vortices and rotating vortex pairs are also formed.

In the first subsection we document how our programs can be used to simulate the experiment reported in reference [133], and compare experimental and numerical results. The experiment measured critical velocity for the emergence of vortex dipoles and rotating vortex pairs in a BEC of sodium atoms ( $^{23}\text{Na}$ ), which do not exhibit the dipolar interaction. Therefore, in this section we also demonstrate how our programs can be modified to switch off the dipolar interaction term when solving NLS equation. The agreement of the results obtained numerically and experimental observations provides an external check of the correctness of our algorithms and their implementations.

In the second subsection we investigate the formation of vortices in a dipolar BEC of dysprosium atoms ( $^{164}\text{Dy}$ ). Our simulations follow the same methodology as the experiment with sodium atoms, only with atomic species exhibiting strong dipolar interaction. We study effects of the dipolar interaction on the critical velocity for the emergence of vortices, as well as the interplay between contact and dipolar interaction. The visualization extensions presented in reference [42] have proven to be indispensable during these simulations, as they allow much easier study of the results and control of the simulation.



### 5.1. Formation of Vortices in a BEC

BEC is a superfluid quantum liquid and one of its hallmarks are quantized vortices, which appear as elementary excitations of the system. Quantization of vortices is connected to excitation spectrum and, according to Landau's criterion, leads to the existence of a minimal velocity an obstacle moving through the superfluid has to have in order to generate such elementary excitations. This critical velocity  $v_c$  can be experimentally measured and, in principle, depends on the experimental protocol used.




Figure 24. Illustration of the experimental setup used in reference [133] to study vortex formation in a BEC of sodium atoms. A repulsive Gaussian laser beam is initially at the center of the condensate and moves along  $y$  direction with a constant velocity  $v$  until it reaches its final position. The beam is then switched off during a period of 0.5 s. For sufficiently large velocity vortices and vortex pairs are generated.

Reference [133] reports the study of vortex formation in an oblate BEC of  $N_{\text{at}} = 3.2 \times 10^6$  sodium atoms. The trapping potential frequencies used are  $(\omega_x, \omega_y, \omega_z) = 2\pi \times (9, 9, 400)$  Hz, and the  $s$ -wave scattering length determining the strength of contact interaction was  $a_s = 51.9 a_0$ , where  $a_0 = 0.0529$  nm is the Bohr radius. The moving obstacle is realized by a repulsive Gaussian laser

beam corresponding to an additional potential of the form

$$V_B(\mathbf{r}; t) = V_0 \exp \left\{ -2 \frac{[y - y_0(t)]^2 + z^2}{\sigma^2} \right\}, \quad (65)$$

where  $V_0$  represents the strength of the beam,  $\sigma$  is the  $1/e^2$  beam waist, while  $y_0(t) = vt$  determines the center of the beam position  $(0, y_0(t), 0)$ . In our simulations, as in the experiment, the initial position of the beam is at the condensate center and then it moves with the velocity  $v$  along  $y$  direction for  $24 \mu\text{m}$ . Afterwards, the beam is switched off linearly during  $0.5 \text{ s}$ . Illustration of the experimental setup is given in Figure 24. Using imaginary-time propagation, we calculate the ground state of the system with the trapping potential  $V(\mathbf{r}) + V_B(\mathbf{r}; t)$ , where  $V(\mathbf{r})$  corresponds to the harmonic part, with the frequencies given above. The experiment measured critical velocity for the emergence of vortex dipoles, pairs of vortices of the opposite sign. We numerically addressed this setup and calculated this critical velocity for the values of the parameters  $\sigma = 1.3 l$  and  $V_0 = 250 \hbar \tilde{\omega}$ , where  $l = 6.988 \mu\text{m}$  is the harmonic oscillator length for the referent frequency  $\tilde{\omega} = \omega_x$ .

Before running simulations with the parameters described above, we had to implement the time-dependent potential  $V(\mathbf{r}) + V_B(\mathbf{r}; t)$ . Initialization of  $V(\mathbf{r})$  remains the same, i.e., it is initialized before the main loop and stored in an array, while calculation of  $V_B(\mathbf{r}; t)$  is implemented in a separate function that is called in each iteration within the main loop. Two distinct phases of  $V_B(\mathbf{r}; t)$  exist: one relating to the movement of the beam, and the other relating to the beam shutdown during  $0.5 \text{ s}$ . Each phase is implemented as a separate function, and the active phase is determined by the current time, i.e., iteration number.

Figure 25 shows several snapshots of a typical dynamical evolution of the system. Imaginary-time propagation yields ground state of the system, which is shown in Figure 25(a), where we plot integrated 2D density profile in  $y$ - $z$  plane. To model local inhomogeneities always present in the experiment, we add uniformly distributed random noise to the wave function of the order of 10%. Such modified ground state represents initial state of the system, and the beam starts to move along  $y$  direction at time  $t = 0$  with the speed  $v = 1.26 \text{ mm/s}$ . Figures 25(b) and 25(c) show 2D density profile of the system at times  $t = 52.17 \text{ ms}$  and  $t = 358.1 \text{ ms}$ , respectively. In Figure 25(b) the beam already reached its final position and its switch-off started. We can observe that several vortices (vortex dipoles) are generated and that the used speed exceeds the critical velocity. Generated vortices are stable and can be seen after extended




Figure 25. Time evolution of 2D density profile of BEC of sodium atoms with an obstacle moving at speed  $v = 1.26 \text{ mm/s} > v_c$ , when several vortices are generated. Each panel shows integrated 2D density profile in  $y$ - $z$  plane at different time  $t$ . All lengths are expressed in units of  $\mu\text{m}$ , and particle density is given in units  $N_{\text{at}}/l^2$ .

period of time, Figure 25(c). To confirm that we indeed have vortices and not just localized density minima, we plot the  $x = 0$  slice of the phase of the wave function in Figure 26. We observe characteristic braiding and jumps of the phase in the vicinity of density minima, which is a well-known hallmark of a vortex.

When the obstacle is moving with an under-critical velocity, no vortices are generated, as in Figure 27. The velocity  $v = 0.87 \text{ mm/s}$  is just slightly lower than the the critical one, and we can see precursors of vortices at the beam edge in Figures 27(a) and 27(b). At higher velocities, vortices emerge, as we have

seen in Figure 25. At even higher velocities, Figure 28, rotating vortex pairs are generated, in addition to individual vortices.

The obtained results are in good agreement with experimental findings of reference [133] and show that the programs developed within this thesis can be successfully used to model BEC systems with contact interaction, even for the most complex setup, when vortices are generated due to a moving obstacle.

## 5.2. Effects of Dipolar Interactions on Vortex Formation in a BEC

To further test our programs, we model vortex formation in a BEC of dysprosium atoms ( $^{164}\text{Dy}$ ), which exhibit strong magnetic dipole moment. In par-




Figure 26. Time evolution of the wave function phase of BEC of sodium atoms with an obstacle moving at speed  $v = 1.26 \text{ mm/s} > v_c$ , when several vortices are generated. Each panel shows  $x = 0$  slice of the phase of the wave function in  $y$ - $z$  plane at different time  $t$ . All lengths are expressed in units of  $\mu\text{m}$ .

particular, we study effects of the dipolar interaction strength for varying  $s$ -wave scattering lengths on the critical velocity for the emergence of vortices, which was not investigated experimentally. Dysprosium atoms have the largest magnetic dipole moment ( $m = 10\mu_B$ ) available in ultracold atom experiments, corresponding the characteristic dipole-dipole interaction length  $a_{\text{dd}} = 132 a_0$ . External magnetic field can be used to align all atomic dipoles in the same direction, as well as to tune their strength up to a maximal value given above. The same applies to the  $s$ -wave scattering length, which determines the contact




Figure 27. Time evolution of 2D density profile of BEC of sodium atoms with an obstacle moving at speed  $v = 0.87 \text{ mm/s} \lesssim v_c$ , when no vortices are generated. Each panel shows integrated 2D density profile in  $y$ - $z$  plane at different time  $t$ . All lengths are expressed in units of  $\mu\text{m}$ , and particle density is given in units  $N_{\text{at}}/l^2$ .

interaction strength and which can also be tuned using the Feshbach resonance technique. Following reference [134], we take the same range of possible values for  $a_s$  as for  $a_{\text{dd}}$ , i.e., up to a maximal value of  $132 a_0$ . With this we show how the developed programs can be used to model and theoretically address new physical phenomena, before they are studied experimentally. Such approach is essential for the design of many upcoming experiments, since otherwise it would be extremely difficult to predict what would be the relevant range of physical quantities to be measured, thus making it very challenging to perform




Figure 28. Time evolution of 2D density profile of BEC of sodium atoms with an obstacle moving at speed  $v = 1.4$  mm/s  $\gg v_c$ , when rotating vortex pairs are also generated. Each panel shows integrated 2D density profile in  $y$ - $z$  plane at different time  $t$ . A rotating vortex pair is highlighted in panel (c). All lengths are expressed in units of  $\mu\text{m}$ , and particle density is given in units  $N_{\text{at}}/l^2$ .



the experiments. Having the results of detailed numerical simulations for the particular system enables experiments to target appropriate range of all relevant quantities, and to focus on discovering new phenomena.

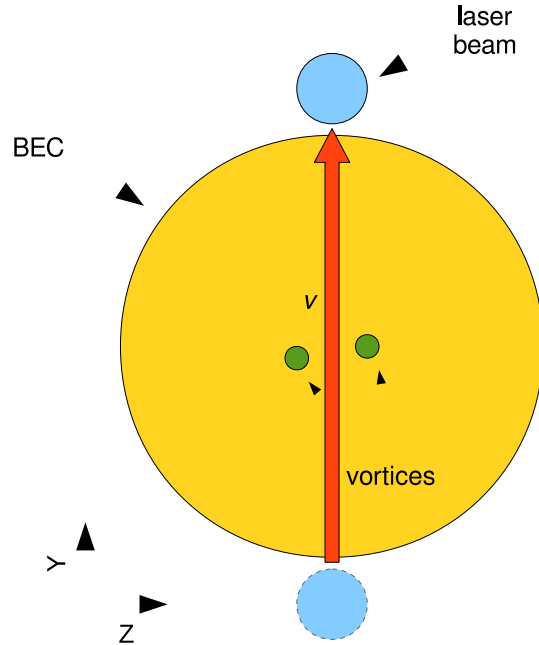


Figure 29. Illustration of the setup we use to study vortex formation in a dipolar BEC of dysprosium atoms. A repulsive Gaussian laser beam is initially outside of the condensate and moves along  $y$  direction with a constant velocity  $v$  until it reaches its final position on the other side. For sufficiently large velocity vortices are generated.

The experiment follows the same methodology as in previous section and uses the same parameters where applicable. Namely, the trap frequencies remain the same, i.e.,  $(\omega_x, \omega_y, \omega_z) = 2\pi \times (9, 9, 400)$  Hz, as well as strength of the beam  $V_0 = 250 \hbar \tilde{\omega}$ . Due to the strong dipolar interactions affecting the stability of the BEC, we had to use a much smaller number of atoms,  $N_{\text{at}} = 8 \times 10^4$ . This resulted in a much smaller BEC, so we had to reduce the  $1/e^2$  beam waist to  $\sigma = l$ , where  $l = 2.617 \mu\text{m}$  is harmonic oscillator length corresponding to dysprosium atoms. Also, we positioned the beam outside of the condensate, and move it all the way to the other side, as illustrated in Figure 29. The potential  $V_B(\mathbf{r}; t)$  has the same form (65) as in previous section, just the center of the beam is now given by  $y_0(t) = y_{00} + vt$ , where  $y_{00} = -15$  in units of  $l$ .

Modifications to the programs presented in the previous section can be

reused for this numerical experiment and we only need to re-enable previously disabled calculation of the dipolar term in the main time propagation loop. As before, the ground state of the system is obtained using imaginary-time program, which then serves as the initial state of the real-time propagation program, with uniformly distributed random noise of 10% added. To calculate critical velocity  $v_c$  for the emergence of vortices, we search for the minimal speed of the laser beam for which two vortices appear, as illustrated in Figure 29.

Figure 30 shows several 2D density profiles of a typical dynamical evolu-




Figure 30. Time evolution of 2D density profile of BEC of dysprosium atoms for  $a_s = 66a_0$  and  $a_{dd} = 44a_0$ , with an obstacle moving at speed  $v = v_c = 0.16$  mm/s, when two individual vortices are generated. Each panel shows integrated 2D density profile in  $y$ - $z$  plane at different time  $t$ . All lengths are expressed in units of  $\mu\text{m}$ , and particle density is given in units  $N_{\text{at}}/l^2$ .



tion of the system for velocity  $v = v_c = 0.16$  mm/s. The ground state, obtained through imaginary-time propagation is shown in Figure 30(a). The beam, initially outside of the condensate, moves along the  $y$  direction, as seen in Figures 30(b) through 30(f). In Figure 30(c) the precursors of the vortices form on the edges of the beam, which then separate from the beam if  $v \geq v_c$ , a situation seen in Figure 30(d). In Figures 30(e) and 30(f) we see that the generated vortices are stable, and survive for long propagation times.

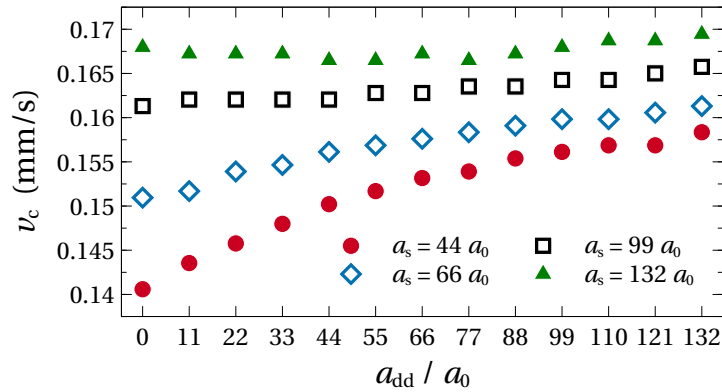


Figure 31. Critical velocity  $v_c$  for the emergence of vortices generated by a moving obstacle as a function of the characteristic dipolar interaction length  $a_{dd}$  (in units of  $a_0$ ) for varying values of the  $s$ -wave scattering length  $a_s$ .

Figure 31 shows results of our numerical study of the dipolar interaction effects on the critical velocity. We see that for large values of the  $s$ -wave scattering length, i.e., for the contact interaction comparable or larger than the dipole interaction, effects of decreasing  $a_{dd}$  are very small and probably could not be experimentally measured. On the other hand, when contact interaction is tuned down so that the dipole interaction starts to dominate the behavior of the system, critical velocity depends much stronger on  $a_{dd}$  and could be easily measured in future experiments.

The above study represents an example on how our programs can be used to verify and compare results of current experiments, as well as to theoretically investigate new phenomena and plan future experiments.

## Conclusion

In this chapter we have reviewed an efficient numerical algorithm for solving the nonlinear Schrödinger equation. The stability and efficiency of the used finite

difference approach is ensured through properly chosen discretization schemes: the split-step Crank-Nicolson method for the time derivatives and central difference formula for the spatial derivatives. In a computational implementation, the presented numerical algorithm is well suited for various parallelization techniques. We have exploited several parallelization approaches available on present-day computer architectures: OpenMP (for shared memory multiprocessing), OpenMP/MPI (for distributed memory systems, where each compute node uses OpenMP to maximize performance), CUDA (for single GPGPUs), CUDA/MPI (for distributed memory systems, where each compute node has GPGPU installed), as well as their hybrid combinations. For efficient hybrid implementations it is essential to divide the work in a way that is suited to the performances of the available computer architecture. To this purpose additional input parameters are introduced and it is shown that their optimal values can be set using a variant of a genetic algorithm. Both strong and weak scalability of the algorithm are explicitly demonstrated by detailed measurements and subsequent analysis. The scalability is proved to be of high importance as parallel versions of the algorithm allow for efficient studies of realistically large system sizes, which are needed in order to make a direct connection to actual physical experiments. As illustrative examples, three applications of the presented algorithms to the topics of current interest in the field of cold atomic BECs were presented. In section 3 a study of phase separation dynamics of a two-component BEC is performed. In section 4 we addressed excitations of a BEC with the coupling of spin and angular momentum and finally in section 5 we investigated vortex formation dynamics in a stirred BEC in the presence of dipolar interactions.

## Acknowledgments

The authors acknowledge valuable discussions with Sadhan K. Adhikari. This work was supported by the Ministry of Education, Science, and Technological Development of the Republic of Serbia under project ON171017. Numerical simulations were performed on the PARADOX supercomputing facility at the Scientific Computing Laboratory of the Institute of Physics Belgrade.

## References

- [1] Dalfovo, F., Giorgini, S., Pitaevskii, L. P., and Stringari, S. Theory of

- Bose-Einstein condensation in trapped gases. *Rev. Mod. Phys.* 71 (1999), 463–512.
- [2] Gross, E. P. Structure of a quantized vortex in boson systems. *Il Nuovo Cimento* 20 (1961), 454–477.
- [3] Pitaevskii, L. P. Vortex Lines in an Imperfect Bose Gas. *Sov. Phys. JETP* 13 (1961), 451.
- [4] Agrawal, G. P. Nonlinear fiber optics: its history and recent progress. *J. Opt. Soc. Am. B* 28 (2011), A1.
- [5] Taha, T. R., and Ablowitz, M. I. Analytical and Numerical Aspects of Certain Nonlinear Evolution Equations. II. Numerical, Nonlinear Schrödinger Equation. *J. Comput. Phys.* 55 (1984), 203–230.
- [6] Anderson, M. H., Ensher, J. R., Matthews, M. R., Wieman, C. E., and Cornell, E. A. Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor. *Science* 269 (1995), 198–201.
- [7] Davis, K. B., Mewes, M.-O., Andrews, M. R., van Druten, N. J., Durfee, D. S., Kurn, D. M., and Ketterle, W. Bose-Einstein condensation in a gas of sodium atoms. *Phys. Rev. Lett.* 75 (1995), 3969–3973.
- [8] Bradley, C. C., Sackett, C. A., Tollett, J. J., and Hulet, R. G. Evidence of Bose-Einstein Condensation in an Atomic Gas with Attractive Interactions. *Phys. Rev. Lett.* 75 (1995), 1687–1690.
- [9] Edwards, M., and Burnett, K. Numerical solution of the nonlinear Schrödinger equation for small samples of trapped neutral atoms. *Phys. Rev. A* 51 (1995), 1382–1386.
- [10] Ruprecht, P. A., Holland, M. J., Burnett, K., and Edwards, M. Time-dependent solution of the nonlinear Schrödinger equation for Bose-condensed trapped neutral atoms. *Phys. Rev. A* 51 (1995), 4704–4711.
- [11] Dalfovo, F., and Stringari, S. Bosons in anisotropic traps: Ground state and vortices. *Phys. Rev. A* 53 (1996), 2477–2485.
- [12] Holland, M., and Cooper, J. Expansion of a Bose-Einstein condensate in a harmonic potential. *Phys. Rev. A* 53 (1996), R1954–R1957.

- 
- [13] Pethick, C. J., and Smith, H. *Bose-Einstein Condensation in Dilute Gases*. Cambridge University Press, 2008.
- [14] Pitaevskii, L. P., and Stringari, S. *Bose-Einstein Condensation*. Oxford University Press, 2003.
- [15] Snyman, J. *Practical Mathematical Optimization*. Springer, 2005.
- [16] Koch, T., Lahaye, T., Metz, J., Fröhlich, B., Griesmaier, A., and Pfau, T. Stabilization of a purely dipolar quantum gas against collapse. *Nat. Phys.* 4 (2008), 218–222.
- [17] Lu, M., Burdick, N. Q., Youn, S. H., and Lev, B. L. Strongly Dipolar Bose-Einstein Condensate of Dysprosium. *Phys. Rev. Lett.* 107 (2011), 190401.
- [18] Aikawa, K., Frisch, A., Mark, M., Baier, S., Rietzler, A., Grimm, R., and Ferlaino, F. Bose-Einstein Condensation of Erbium. *Phys. Rev. Lett.* 108 (2012), 210401.
- [19] Góral, K., and Santos, L. Ground state and elementary excitations of single and binary Bose-Einstein condensates of trapped dipolar gases. *Phys. Rev. A* 66 (2002), 023613.
- [20] Agrawal, G. P. *Nonlinear Fiber Optics*. Academic Press, 2012.
- [21] Antoine, X., Bao, W., and Besse, C. Computational methods for the dynamics of the nonlinear Schrödinger/Gross-Pitaevskii equations. *Comput. Phys. Commun.* 184 (2013), 2621–2633.
- [22] Muruganandam, P., and Adhikari, S. Fortran programs for the time-dependent Gross-Pitaevskii equation in a fully anisotropic trap. *Comput. Phys. Commun.* 180 (2009), 1888–1912.
- [23] Vudragović, D., Vidanović, I., Balaž, A., Muruganandam, P., and Adhikari, S. K. C programs for solving the time-dependent Gross-Pitaevskii equation in a fully anisotropic trap. *Comput. Phys. Commun.* 183 (2012), 2021–2025.
- [24] Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*, 3 ed. Cambridge University Press, New York, NY, USA, 2007.

- [25] Young-S., L. E., Vudragović, D., Muruganandam, P., Adhikari, S. K., and Balaž, A. OpenMP Fortran and C programs for solving the time-dependent Gross-Pitaevskii equation in an anisotropic trap. *Comput. Phys. Commun.* 204 (2016), 209–213.
- [26] Satarić, B., Slavnić, V., Belić, A., Balaž, A., Muruganandam, P., and Adhikari, S. K. Hybrid OpenMP/MPI programs for solving the time-dependent Gross-Pitaevskii equation in a fully anisotropic trap. *Comput. Phys. Commun.* 200 (2016), 411–417.
- [27] Frigo, M., and Johnson, S. G. The design and implementation of FFTW3. *Proc. IEEE* 93 (2005), 216–231.
- [28] Kumar, R. K., Young-S., L. E., Vudragović, D., Balaž, A., Muruganandam, P., and Adhikari, S. Fortran and C programs for the time-dependent dipolar Gross-Pitaevskii equation in an anisotropic trap. *Comput. Phys. Commun.* 195 (2015), 117–128.
- [29] Lončar, V., Balaž, A., Bogojević, A., Škrbić, S., Muruganandam, P., and Adhikari, S. K. CUDA programs for solving the time-dependent dipolar Gross-Pitaevskii equation in an anisotropic trap. *Comput. Phys. Commun.* 200 (2016), 406–410.
- [30] Lončar, V., Young-S., L. E., Škrbić, S., Muruganandam, P., Adhikari, S. K., and Balaž, A. OpenMP, OpenMP/MPI, and CUDA/MPI C programs for solving the time-dependent dipolar Gross-Pitaevskii equation. *Comput. Phys. Commun.* 209 (2016), 190–196.
- [31] Blelloch, G. Scans as primitive parallel operations. *IEEE Trans. Comput.* 38 (1989), 1526–1538.
- [32] Blelloch, G. E. *Prefix sums and their applications*. Tech. rep., School of Computer Science, Carnegie Mellon University, 1990.
- [33] NVIDIA Corporation. *cuFFT Library User's Guide*, 7.5 ed., 2017.
- [34] Lee, J., Samadi, M., Park, Y., and Mahlke, S. Transparent cpu-gpu collaboration for data-parallel kernels on heterogeneous systems. In *Proceedings of the 22Nd International Conference on Parallel Architectures and Compilation Techniques* (Piscataway, NJ, USA, 2013), PACT '13, IEEE Press, pp. 245–256.

- 
- [35] Papadrakakis, M., Stavroulakis, G., and Karatarakis, A. A new era in scientific computing: Domain decomposition methods in hybrid CPU–GPU architectures. *Comput. Meth. Appl. Mech. Eng.* 200 (2011), 1490–1508.
- [36] Li, Y., Diamond, J. R., Wang, X., Lin, H., Yang, Y., and Han, Z. Large-scale fast fourier transform on a heterogeneous multi-core system. *Int. J. High Perform. Comput. Appl.* 26 (2012), 148–158.
- [37] Ogata, Y., Endo, T., Maruyama, N., and Matsuoka, S. An efficient, model-based CPU-GPU heterogeneous FFT library. In *2008 IEEE International Symposium on Parallel and Distributed Processing* (2008), IEEE.
- [38] Chen, S., and Li, X. A hybrid GPU/CPU FFT library for large FFT problems. In *2013 IEEE 32nd International Performance Computing and Communications Conference (IPCCC)* (2013), IEEE.
- [39] Chu, E., and George, A. *Inside the FFT Black Box*. CRC Press, 1999.
- [40] Lončar, V., Young-S., L. E., Škrbić, S., Muruganandam, P., Adhikari, S. K., and Balaž, A. *DBEC-OMP-CUDA-MPI package*. <https://data.mendeley.com/datasets/j3z9z379m8/2>, 2016.
- [41] Lončar, V., Young-S., L. E., Škrbić, S., Muruganandam, P., Adhikari, S. K., and Balaž, A. *Repository of DBEC-GP family of programs*. <https://git.ipb.ac.rs/vloncar/DBEC-GP>, 2016. Accessed: 2017-04-10.
- [42] Lončar, V. Hybrid parallel algorithms for solving nonlinear Schrödinger equation. *PhD Thesis, University of Novi Sad* (2017).
- [43] Darken, C., Chang, J., and Moody, J. Learning rate schedules for faster stochastic gradient search. In *Neural Networks for Signal Processing II Proceedings of the 1992 IEEE Workshop* (1992), IEEE.
- [44] Schaul, T., Zhang, S., and LeCun, Y. No more pesky learning rates. In *Proceedings of the 30th International Conference on Machine Learning* (Atlanta, Georgia, USA, 2013), S. Dasgupta and D. McAllester, Eds., vol. 28 of *Proceedings of Machine Learning Research*, PMLR, pp. 343–351.

- [45] Spall, J. C. Implementation of the simultaneous perturbation algorithm for stochastic optimization. *IEEE Trans. Aerosp. Electron. Syst.* 34 (1998), 817–823.
- [46] Spall, J. C. Stochastic optimization and the simultaneous perturbation method. In *Proceedings of the 31st conference on Winter simulation Simulation—a bridge to the future - WSC 1999* (1999), ACM Press.
- [47] Holland, J. H. *Adaptation in Natural and Artificial Systems*. MIT University Press Group Ltd, 1992.
- [48] Mitchell, M. *An Introduction to Genetic Algorithms*. MIT Press, 1998.
- [49] Gen, M., and Cheng, R. *Genetic Algorithms and Engineering Optimization*. John Wiley & Sons Inc, 1999.
- [50] Kumar, A. Encoding schemes in genetic algorithm. *Int. J. Adv. Res. IT Eng.* 2 (2013), 1–7.
- [51] Eshelman, L. J., Caruana, R. A., and Schaffer, J. D. Biases in the crossover landscape. In *Proceedings of the Third International Conference on Genetic Algorithms* (San Francisco, CA, USA, 1989), Morgan Kaufmann Publishers Inc., pp. 10–19.
- [52] Spears, V. M., and Jong, K. A. D. On the virtues of parameterized uniform crossover. In *Proceedings of the Fourth International Conference on Genetic Algorithms* (1991), pp. 230–236.
- [53] Richter, J. N. *On mutation and crossover in the theory of evolutionary algorithms*. PhD thesis, Montana State University, 2010.
- [54] White, D. R., and Poulding, S. A rigorous evaluation of crossover and mutation in genetic programming. In *Lecture Notes in Computer Science*, vol. 5481. Springer Berlin Heidelberg, 2009, pp. 220–231.
- [55] da Silva, A. P. A., and Falcao, D. M. Fundamentals of genetic algorithms. In *Modern Heuristic Optimization Techniques: Theory and Applications to Power Systems*, K. Y. Lee and M. A. El-Sharkawi, Eds. John Wiley & Sons Inc, 2008, ch. 2, pp. 25–42.

- 
- [56] Luke, S., and Spector, L. A comparison of crossover and mutation in genetic programming. In *Genetic Programming 1997: Proceedings of the Second Annual Conference (1997)*, J. R. e. a. Koza, Ed., Morgan Kaufmann, pp. 240–248.
- [57] Luke, S., and Spector, L. A revised comparison of crossover and mutation in genetic programming. In *Genetic Programming 1998: Proceedings of the Third Annual Conference (1998)*, J. R. e. a. Koza, Ed., Morgan Kaufmann, pp. 208–213.
- [58] Amdahl, G. M. Validity of the single processor approach to achieving large scale computing capabilities. In *Proceedings of the April 18-20, 1967, Spring Joint Computer Conference - AFIPS '67 (Spring) (1967)*, ACM Press.
- [59] Open MPI project. *Faq: Running cuda-aware open mpi*. <https://www.open-mpi.org/faq/?category=runcuda>, 2016. Accessed: 2017-04-10.
- [60] Frigo, M., and Johnson, S. G. *FFTW 3 documentation*, version 3.3.6. [http://www.fftw.org/fftw3\\_doc/](http://www.fftw.org/fftw3_doc/), 2017. Accessed: 2017-04-10.
- [61] Edwards, M., Ruprecht, P. A., Burnett, K., Dodd, R. J., and Clark, C. W. Collective Excitations of Atomic Bose-Einstein Condensates. *Phys. Rev. Lett.* 77 (1996), 1671–1674.
- [62] Vidanović, I., Balaž, A., Al-Jibbouri, H., and Pelster, A. Nonlinear Bose-Einstein-condensate dynamics induced by a harmonic modulation of the s-wave scattering length. *Phys. Rev. A* 84 (2011), 013618.
- [63] Balaž, A., and Nicolin, A. I. Faraday waves in binary nonmiscible Bose-Einstein condensates. *Phys. Rev. A* 85 (2012), 023613.
- [64] Balaž, A., Paun, R., Nicolin, A. I., Balasubramanian, S., and Ramaswamy, R. Faraday waves in collisionally inhomogeneous Bose-Einstein condensates. *Phys. Rev. A* 89 (2014), 023609.
- [65] Sudharsan, J. B., Radha, R., Carina Raportaru, M., Nicolin, A. I., and Balaž, A. Faraday and resonant waves in binary collisionally-inhomogeneous Bose-Einstein condensates. *J. Phys. B-At. Mol. Opt. Phys.* 49 (2016), 165303.



- [66] Vidanović, I., van Druten, N. J., and Haque, M. Spin modulation instabilities and phase separation dynamics in trapped two-component Bose condensates. *New J. Phys.* 15 (2013), 035008.
- [67] Al-Jibbouri, H., Vidanović, I., Balaž, A., and Pelster, A. Geometric resonances in Bose-Einstein condensates with two- and three-body interactions. *J. Phys. B-At. Mol. Opt. Phys.* 46 (2013), 065303.
- [68] Sakhel, R. R., Sakhel, A. R., and Ghassib, H. B. Nonequilibrium Dynamics of a Bose-Einstein Condensate Excited by a Red Laser Inside a Power-Law Trap with Hard Walls. *J. Low Temp. Phys.* 173 (2013), 177–206.
- [69] Sakhel, R. R., Sakhel, A. R., and Ghassib, H. B. On the phase-correlation and phase-fluctuation dynamics of a strongly excited Bose gas. *Physica B* 478 (2015), 68–76.
- [70] Sakhel, R. R., Sakhel, A. R., Ghassib, H. B., and Balaz, A. Conditions for order and chaos in the dynamics of a trapped Bose-Einstein condensate in coordinate and energy space. *Eur. Phys. J. D* 70 (2016), 66.
- [71] Xi, K.-T., Li, J., and Shi, D.-N. Localization of a two-component Bose-Einstein condensate in a one-dimensional random potential. *Physica B* 459 (2015), 6–11.
- [72] Mithun, T., Porsezian, K., and Dey, B. Disorder-induced vortex lattice melting in a Bose-Einstein condensate. *Phys. Rev. A* 93 (2016), 013620.
- [73] Gautam, S., and Adhikari, S. K. Fractional-charge vortex in a spinor Bose-Einstein condensate. *Phys. Rev. A* 93 (2016), 013630.
- [74] Akram, J., and Pelster, A. Sculpting quasi-one-dimensional Bose-Einstein condensate to generate calibrated matter waves. *Phys. Rev. A* 93 (2016), 023606.
- [75] Akram, J., and Pelster, A. Statics and dynamics of quasi one-dimensional Bose-Einstein condensate in harmonic and dimple trap. *Laser Phys.* 26 (2016), 065501.
- [76] Sakhel, R. R., and Sakhel, A. R. Application of the Lagrangian variational method to a one-dimensional Bose gas in a dimple trap. *J. Phys. B-At. Mol. Opt. Phys.* 50 (2017), 105301.

- 
- [77] Akram, J., and Pelster, A. Numerical study of localized impurity in a Bose-Einstein condensate. *Phys. Rev. A* 93 (2016), 033610.
- [78] Manikandan, K., Muruganandam, P., Senthilvelan, M., and Lakshmanan, M. Manipulating localized matter waves in multicomponent Bose-Einstein condensates. *Phys. Rev E* 93 (2016), 032212.
- [79] Lai, C.-Y., and Chien, C.-C. Geometry-Induced Memory Effects in Isolated Quantum Systems: Cold-Atom Applications. *Phys. Rev. Applied* 5 (2016), 034001.
- [80] Akram, J., Girodias, B., and Pelster, A. Quasi one-dimensional Bose-Einstein condensate in a gravito-optical surface trap. *J. Phys. B-At. Mol. Opt. Phys.* 49 (2016), 075302.
- [81] Khellil, T., Balaž, A., and Pelster, A. Analytical and numerical study of dirty bosons in a quasi-one-dimensional harmonic trap. *New J. Phys.* 18 (2016), 063003.
- [82] Nakamura, K., Babajanov, D., Matrasulov, D., Kobayashi, M., and Muruganandam, P. Dynamics of trapped interacting vortices in Bose-Einstein condensates: a role of breathing degree of freedom. *J. Phys. A-Math. Theor.* 49 (2016), 315102.
- [83] Sakhel, R. R., and Sakhel, A. R. Elements of Vortex-Dipole Dynamics in a Nonuniform Bose-Einstein Condensate. *J. Low Temp. Phys.* 184 (2016), 1092–1113.
- [84] Adhikari, S. K. Statics and dynamics of a self-bound dipolar matter-wave droplet. *Laser Phys. Lett.* 14 (2017), 025501.
- [85] Adhikari, S. K. Statics and dynamics of a self-bound matter-wave quantum ball. *Phys. Rev. A* 95 (2017), 023606.
- [86] Adhikari, S. K. Elastic collision and molecule formation of spatiotemporal light bullets in a cubic-quintic nonlinear medium. *Phys. Rev E* 94 (2016), 032217.
- [87] Gautam, S., and Adhikari, S. K. Phase separation in a spin-orbit-coupled Bose-Einstein condensate. *Phys. Rev. A* 90 (2014), 043619.


- [88] Gautam, S., and Adhikari, S. K. Mobile vector soliton in a spin-orbit coupled spin-1 condensate. *Laser Phys. Lett.* 12 (2015), 045501.
- [89] Gautam, S., and Adhikari, S. K. Vector solitons in a spin-orbit-coupled spin-2 Bose-Einstein condensate. *Phys. Rev. A* 91 (2015), 063617.
- [90] Gautam, S., and Adhikari, S. K. Vortex-bright solitons in a spin-orbit-coupled spin-1 condensate. *Phys. Rev. A* 95 (2017), 013608.
- [91] Vinayagam, P. S., Radha, R., Bhuvaneswari, S., Ravisankar, R., and Muruganandam, P. Bright soliton dynamics in spin orbit-Rabi coupled Bose-Einstein condensates. *Commun. Nonlinear Sci. Numer. Simulat.* 50 (2017), 68–76.
- [92] Vasić, I., and Balaž, A. Excitation spectra of a Bose-Einstein condensate with an angular spin-orbit coupling. *Phys. Rev. A* 94 (2016), 033627.
- [93] Kishor Kumar, R., and Muruganandam, P. Vortex dynamics of rotating dipolar Bose-Einstein condensates. *J. Phys. B-At. Mol. Opt. Phys.* 45 (2012), 215301.
- [94] Kishor Kumar, R., and Muruganandam, P. Numerical studies on vortices in rotating dipolar Bose-Einstein condensates. In *J. Phys.: Conf. Ser.* (2014), vol. 497, p. 012036.
- [95] Kumar, R. K., and Muruganandam, P. Effect of optical lattice potentials on the vortices in rotating dipolar Bose-Einstein condensates. *Eur. Phys. J. D* 68 (2014), 289.
- [96] Kishor Kumar, R., Sriraman, T., Fabrelli, H., Muruganandam, P., and Gammal, A. Three-dimensional vortex structures in a rotating dipolar Bose-Einstein condensate. *J. Phys. B-At. Mol. Opt. Phys.* 49 (2016), 155301.
- [97] Young-S., L. E., and Adhikari, S. K. Mixing, demixing, and structure formation in a binary dipolar Bose-Einstein condensate. *Phys. Rev. A* 86 (2012), 063611.
- [98] Young-S., L. E., and Adhikari, S. K. Dipolar droplet bound in a trapped Bose-Einstein condensate. *Phys. Rev. A* 87 (2013), 013618.

- 
- [99] Adhikari, S. K. Stability of trapped degenerate dipolar Bose and Fermi gases. *J. Phys. B-At. Mol. Opt. Phys.* **46** (2013), 115301.
- [100] Adhikari, S. K. Stability and collapse of fermions in a binary dipolar boson-fermion  $^{164}\text{Dy}$ - $^{161}\text{Dy}$  mixture. *Phys. Rev. A* **88** (2013), 043603.
- [101] Kishor Kumar, R., Muruganandam, P., and Malomed, B. A. Vortical and fundamental solitons in dipolar Bose-Einstein condensates trapped in isotropic and anisotropic nonlinear potentials. *J. Phys. B-At. Mol. Opt. Phys.* **46** (2013), 175302.
- [102] Adhikari, S. K. Stable and mobile excited two-dimensional dipolar Bose-Einstein condensate solitons. *J. Phys. B-At. Mol. Opt. Phys.* **47** (2014), 225304.
- [103] Adhikari, S. K. Stable, mobile, dark-in-bright, dipolar Bose-Einstein-condensate solitons. *Phys. Rev. A* **89** (2014), 043615.
- [104] Adhikari, S. K. Demixing and symmetry breaking in binary dipolar Bose-Einstein-condensate solitons. *Phys. Rev. A* **89** (2014), 013630.
- [105] Adhikari, S. K. Bright dipolar Bose-Einstein-condensate soliton mobile in a direction perpendicular to polarization. *Phys. Rev. A* **90** (2014), 055601.
- [106] Adhikari, S. K. Two-dimensional bright and dark-in-bright dipolar Bose-Einstein condensate solitons on a one-dimensional optical lattice. *Laser Phys. Lett.* **13** (2016), 085501.
- [107] Adhikari, S. K. Stable and mobile two-dimensional dipolar ring-dark-in-bright Bose-Einstein condensate soliton. *Laser Phys. Lett.* **13** (2016), 035502.
- [108] Adhikari, S. K. Self-trapping of a dipolar Bose-Einstein condensate in a double well. *Phys. Rev. A* **89** (2014), 043609.
- [109] Ho, T.-L., and Shenoy, V. B. Binary Mixtures of Bose Condensates of Alkali Atoms. *Phys. Rev. Lett.* **77** (1996), 3276–3279.
- [110] Matthews, M. R., Hall, D. S., Jin, D. S., Ensher, J. R., Wieman, C. E., Cornell, E. A., Dalfovo, F., Minniti, C., and Stringari, S. Dynamical

- Response of a Bose-Einstein Condensate to a Discontinuous Change in Internal State. *Phys. Rev. Lett.* 81 (1998), 243–247.
- [111] Hall, D. S., Matthews, M. R., Ensher, J. R., Wieman, C. E., and Cornell, E. A. Dynamics of Component Separation in a Binary Mixture of Bose-Einstein Condensates. *Phys. Rev. Lett.* 81 (1998), 1539–1542.
- [112] Miesner, H.-J., Stamper-Kurn, D. M., Stenger, J., Inouye, S., Chikkatur, A. P., and Ketterle, W. Observation of Metastable States in Spinor Bose-Einstein Condensates. *Phys. Rev. Lett.* 82 (1999), 2228–2231.
- [113] Salasnich, L., Parola, A., and Reatto, L. Effective wave equations for the dynamics of cigar-shaped and disk-shaped Bose condensates. *Phys. Rev. A* 65 (2002), 043614.
- [114] Pu, H., and Bigelow, N. P. Collective Excitations, Metastability, and Nonlinear Response of a Trapped Two-Species Bose-Einstein Condensate. *Phys. Rev. Lett.* 80 (1998), 1134–1137.
- [115] Timmermans, E. Phase Separation of Bose-Einstein Condensates. *Phys. Rev. Lett.* 81 (1998), 5718–5721.
- [116] Kasamatsu, K., and Tsubota, M. Multiple Domain Formation Induced by Modulation Instability in Two-Component Bose-Einstein Condensates. *Phys. Rev. Lett.* 93 (2004), 100402.
- [117] Kasamatsu, K., and Tsubota, M. Modulation instability and solitary-wave formation in two-component Bose-Einstein condensates. *Phys. Rev. A* 74 (2006), 013617.
- [118] Navarro, R., Carretero-González, R., and Kevrekidis, P. G. Phase separation and dynamics of two-component Bose-Einstein condensates. *Phys. Rev. A* 80 (2009), 023613.
- [119] Lin, Y.-J., Jiménez-García, K., and Spielman, I. B. Spin-orbit-coupled Bose-Einstein condensates. *Nature (London)* 471 (2011), 83–86.
- [120] Beeler, M. C., Williams, R. A., Jiménez-García, K., Leblanc, L. J., Perry, A. R., and Spielman, I. B. The spin Hall effect in a quantum gas. *Nature (London)* 498 (2013), 201–204.

- 
- [121] Zhang, J.-Y., Ji, S.-C., Chen, Z., Zhang, L., Du, Z.-D., Yan, B., Pan, G.-S., Zhao, B., Deng, Y.-J., Zhai, H., Chen, S., and Pan, J.-W. Collective Dipole Oscillations of a Spin-Orbit Coupled Bose-Einstein Condensate. *Phys. Rev. Lett.* *109* (2012), 115301.
- [122] Kamehchi, M. A., Zhang, Y., Hamner, C., Busch, T., and Engels, P. Measurement of collective excitations in a spin-orbit-coupled Bose-Einstein condensate. *Phys. Rev. A* *90* (2014), 063624.
- [123] Cheuk, L. W., Sommer, A. T., Hadzibabic, Z., Yefsah, T., Bakr, W. S., and Zwierlein, M. W. Spin-injection spectroscopy of a spin-orbit coupled fermi gas. *Phys. Rev. Lett.* *109* (2012), 095302.
- [124] Li, J., Huang, W., Shteynas, B., Burchesky, S., Top, F. C., Su, E., Lee, J., Jamison, A. O., and Ketterle, W. Spin-Orbit Coupling and Spin Textures in Optical Superlattices. *Phys. Rev. Lett.* *117* (2016), 185301.
- [125] Li, J.-R., Lee, J., Huang, W., Burchesky, S., Shteynas, B., Top, F. C., Jamison, A. O., and Ketterle, W. A stripe phase with supersolid properties in spin-orbit-coupled Bose-Einstein condensates. *Nature (London)* *543* (2017), 91–94.
- [126] Galitski, V., and Spielman, I. B. Spin-orbit coupling in quantum gases. *Nature (London)* *494* (2013), 49–54.
- [127] Zhai, H. Degenerate quantum gases with spin-orbit coupling: a review. *Rep. Prog. Phys.* *78* (2015), 026001.
- [128] Hu, Y.-X., Miniatura, C., and Grémaud, B. Half-skyrmion and vortex-antivortex pairs in spinor condensates. *Phys. Rev. A* *92* (2015), 033615.
- [129] DeMarco, M., and Pu, H. Angular spin-orbit coupling in cold atoms. *Phys. Rev. A* *91* (2015), 033630.
- [130] Qu, C., Sun, K., and Zhang, C. Quantum phases of Bose-Einstein condensates with synthetic spin-orbital-angular-momentum coupling. *Phys. Rev. A* *91* (2015), 053630.
- [131] Sun, K., Qu, C., and Zhang, C. Spin-orbital-angular-momentum coupling in Bose-Einstein condensates. *Phys. Rev. A* *91* (2015), 063627.

- [132] Pitaevskii, L. P., and Rosch, A. Breathing modes and hidden symmetry of trapped atoms in two dimensions. *Phys. Rev. A* 55 (1997), R853–R856.
- [133] Kwon, W. J., Moon, G., Seo, S. W., and Shin, Y. Critical velocity for vortex shedding in a Bose-Einstein condensate. *Phys. Rev. A* 91 (2015), 053615.
- [134] Kadau, H., Schmitt, M., Wenzel, M., Wink, C., Maier, T., Ferrier-Barbut, I., and Pfau, T. Observing the Rosensweig instability of a quantum ferrofluid. *Nature (London)* 530 (2016), 194–197.



# CUDA programs for solving the time-dependent dipolar Gross–Pitaevskii equation in an anisotropic trap



Vladimir Lončar<sup>a,\*</sup>, Antun Balaž<sup>a</sup>, Aleksandar Bogojević<sup>a</sup>, Srdjan Škrbić<sup>b</sup>, Paulsamy Muruganandam<sup>c</sup>, Sadhan K. Adhikari<sup>d</sup>

<sup>a</sup> Scientific Computing Laboratory, Institute of Physics Belgrade, University of Belgrade, Pregrevica 118, 11080 Belgrade, Serbia

<sup>b</sup> Department of Mathematics and Informatics, Faculty of Sciences, University of Novi Sad, Trg Dositeja Obradovića 4, 21000 Novi Sad, Serbia

<sup>c</sup> School of Physics, Bharathidasan University, Palkalaiperur Campus, Tiruchirappalli – 620024, Tamil Nadu, India

<sup>d</sup> Instituto de Física Teórica, UNESP – Universidade Estadual Paulista, 01.140-70 São Paulo, São Paulo, Brazil

## ARTICLE INFO

### Article history:

Received 12 November 2015

Accepted 17 November 2015

Available online 17 December 2015

### Keywords:

Bose–Einstein condensate

Dipolar atoms

Gross–Pitaevskii equation

Split-step Crank–Nicolson scheme

Real- and imaginary-time propagation

C program

GPU

CUDA program

Partial differential equation

## ABSTRACT

In this paper we present new versions of previously published numerical programs for solving the dipolar Gross–Pitaevskii (GP) equation including the contact interaction in two and three spatial dimensions in imaginary and in real time, yielding both stationary and non-stationary solutions. New versions of programs were developed using CUDA toolkit and can make use of Nvidia GPU devices. The algorithm used is the same split-step semi-implicit Crank–Nicolson method as in the previous version (Kishor Kumar et al., 2015), which is here implemented as a series of CUDA kernels that compute the solution on the GPU. In addition, the Fast Fourier Transform (FFT) library used in the previous version is replaced by cuFFT library, which works on CUDA-enabled GPUs. We present speedup test results obtained using new versions of programs and demonstrate an average speedup of 12–25, depending on the program and input size.

### New version program summary

*Program title:* DBEC-GP-CUDA package, consisting of: (i) imag2dXY-cuda, (ii) imag2dXZ-cuda, (iii) imag3d-cuda, (iv) real2dXY-cuda, (v) real2dXZ-cuda, (vi) real3d-cuda.

*Catalogue identifier:* AEWL\_v2\_0

*Program Summary URL:* [http://cpc.cs.qub.ac.uk/summaries/AEWL\\_v2\\_0.html](http://cpc.cs.qub.ac.uk/summaries/AEWL_v2_0.html)

*Program obtainable from:* CPC Program Library, Queen's University of Belfast, N. Ireland.

*Licensing provisions:* Standard CPC licence, <http://cpc.cs.qub.ac.uk/licence/licence.html>

*No. of lines in distributed program, including test data, etc.:* 18297.

*No. of bytes in distributed program, including test data, etc.:* 128586.

*Distribution format:* tar.gz.

*Programming language:* CUDA C.

*Computer:* Any modern computer with Nvidia GPU with Compute Capability 2.0 or higher, with CUDA toolkit (compiler and runtime, with cuFFT library, minimum version 6.0) installed.

*Operating system:* Linux.

*RAM:* With provided example inputs, programs should run on a computer with 512 MB GPU RAM. There is no upper limit to amount of memory that can be used, as larger grid sizes require more memory, which scales as  $NX \times NY$  or  $NX \times NZ$  (in 2d) or  $NX \times NY \times NZ$  (in 3d). All programs require roughly the same amount of CPU and GPU RAM.

*Number of processors used:* One CPU core and one Nvidia GPU.

*Classification:* 2.9, 4.3, 4.12.

*External routines/libraries:* CUDA toolkit, version 6.0 or higher, with cuFFT library.

*Catalogue identifier of previous version:* AEWL\_v1\_0.

\* Corresponding author.

E-mail addresses: [vladimir.loncar@ipb.ac.rs](mailto:vladimir.loncar@ipb.ac.rs) (V. Lončar), [antun.balaz@ipb.ac.rs](mailto:antun.balaz@ipb.ac.rs) (A. Balaž), [aleksandar.bogojevic@ipb.ac.rs](mailto:aleksandar.bogojevic@ipb.ac.rs) (A. Bogojević), [srdjan.skrbic@dmi.uns.ac.rs](mailto:srdjan.skrbic@dmi.uns.ac.rs) (S. Škrbić), [anand@cnld.bdu.ac.in](mailto:anand@cnld.bdu.ac.in) (P. Muruganandam), [adhikari@ift.unesp.br](mailto:adhikari@ift.unesp.br) (S.K. Adhikari).



*Journal reference of previous version:* Comput. Phys. Commun. 195 (2015) 117.

*Does the new version supersede the previous version?:* No.

*Nature of problem:* These programs are designed to solve the time-dependent nonlinear partial differential Gross–Pitaevskii (GP) equation with contact and dipolar interactions in two or three spatial dimensions in a harmonic anisotropic trap. The GP equation describes the properties of a dilute trapped Bose–Einstein condensate.

*Solution method:* The time-dependent GP equation is solved by the split-step Crank–Nicolson method by discretizing in space and time. The discretized equation is then solved by propagation, in either imaginary or real time, over small time steps. The contribution of the dipolar interaction is evaluated by a Fourier transformation to momentum space using a convolution theorem. The method yields the solution of stationary and/or non-stationary problems.

*Reasons for the new version:* Previously published dipolar Fortran and C programs [1], based on earlier programs and algorithms for GP equation with the contact interaction [2], are already used within the ultra-cold atoms community [3]. However, they are sequential, and thus did not allow for use of the maximum computing performance modern computers can offer. For this reason we have explored possible ways to accelerate our programs. Detailed profiling revealed that the calculation of FFTs is the most computationally demanding part of our programs. Since using GPUs to compute FFTs with optimized libraries like the cuFFT can lead to much better performance, we have decided to parallelize our programs using Nvidia CUDA toolkit. Also, the massive parallelism offered by GPUs could be exploited to parallelize the nested loops our programs have. We have focused on 2d and 3d versions of our programs, as they perform enough computation to justify and require the use of massive parallelism.

*Summary of revisions:* Previous C programs in two or three spatial dimensions are parallelized using CUDA toolkit from Nvidia and named similarly, with “-cuda” suffix appended to their names. The structure of all programs is identical. Computationally most demanding functions performing time evolution (calcpsidd2, calcnu, calclux, calcluy, calcluz), normalization of the wave function (calcnorm), and calculation of physically relevant quantities (calcmuen, calcrms) were implemented as a series of CUDA kernels, which are executed on GPU. All kernels are implemented with grid-stride loops [4], which allow us to use the same kernel block sizes for all of our kernels. These block sizes can be changed in src/utills/cudautills.cuh, containing the optimal values for current Nvidia Tesla GPUs.

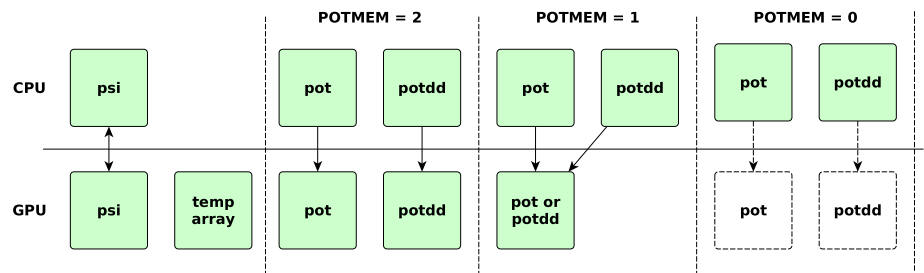
As before, CPU performs the initialization of variables and controls the flow of programs, offloading computation to GPU when needed. Because of the initialization, programs still require almost the same amount of CPU RAM as GPU RAM. Before any computation begins, relevant variables are copied to GPU, where they remain during computation, and only wave function array is returned back to CPU when it is required for writing output.

Parallelization with CUDA toolkit required some dynamically allocated arrays (tensors, matrices, or vectors) to become private for each GPU thread. This has caused an increase in the amount of used GPU memory, since the number of running threads on GPU is very large. Coupled with the fact that GPUs usually have smaller amount of RAM than CPU, this meant that our GPU versions of programs could be used for much smaller input in comparison to sequential versions. In order to fix this problem and reduce memory usage, our programs reuse temporary arrays as much as possible. Aside from allocation of complex tensor/matrix (for 3d or 2d case, respectively) in which we store wave function values, we allocate one complex tensor/matrix, and up to two double precision tensors/matrices, and reuse them for different purposes in computations. Allocated complex tensor/matrix is later also used as two double precision tensors/matrices, for other purposes. This required some reorganization of computation in several functions, mainly in calcmuen and calcpsidd2. In calcmuen we have reorganized computation to reuse temporary array and store partial derivatives in it, so instead of using three (in 3d) or two (in 2d) separate tensors/matrices for partial derivatives, we now use a single temporary tensor/matrix, which we also use for different purposes in other places in programs. In calcpsidd2 we have removed the use of additional temporary array that was only used in FFT computation, and also use real-to-complex and complex-to-real FFT transformations in place of complex-to-complex transformations of previous program versions. This change was possible because condensate density (input array for FFT) is purely real, and thus it exhibits Hermitian symmetry. Some FFT libraries, like the cuFFT used in these programs, can exploit this to reduce memory usage and provide better performance by calculating only non-redundant parts of the array. Additionally, programs can further reduce GPU RAM consumption by keeping the tensor/matrix used to store trap potential and dipolar potential in main RAM, configurable through POTMEM parameter in the input file. Setting value of POTMEM to 2 maximizes performance, and means that programs will allocate two separate tensors/matrices for storing trap potential and dipolar potential in GPU memory. This provides the best performance, but at the cost of a larger total memory consumption. If we set the value of POTMEM to 1, only one tensor/matrix will be allocated in GPU memory, to which trap potential and dipolar potential will be asynchronously copied from main memory when they are needed for computation. In this case, tensor/matrix will initially contain trap potential, which will be replaced with dipolar potential during execution of FFT in calcpsidd2, and replaced back with trap potential during inverse FFT. Finally, setting POTMEM to 0 will instruct the programs not to allocate any GPU memory for storing potentials and will instead use main memory, which GPU can access through slower PCI-Express bus. Figure 1 explains how memory is used and the possible values of POTMEM. We suggest using POTMEM value of 2 if memory permits, and using values of 1 or 0 if problem cannot fit into GPU memory. If POTMEM is not specified, programs will check if GPU memory is large enough to fit all variables and set POTMEM accordingly.

Time propagation functions calclux, calcluy, and calcluz have a recursive relation that makes them difficult to parallelize. In principle, recursive relations could be parallelized using a higher-order prefix

sum algorithm [5] (also known as scan algorithm), but implementation of this would require multiple CUDA kernels [6]. Since recursive relations are in the innermost loop, launching of all required kernels would create a sizeable overhead. Also, the number of grid points in each dimension is usually not large enough to compensate that overhead. Therefore, we have chosen an approach that, instead of parallelizing the inner loop which has the recursive relation, we parallelize the outer loops, and each GPU thread computes the whole innermost loop. Since each GPU thread now requires its own array for storing Crank–Nicolson coefficients  $cbeta$ , we reuse existing temporary tensor/matrix for storing these values. Similar pattern of parallelizing outer loops was also used in `calcnorm`, `calcrms`, and `calcmuen`.

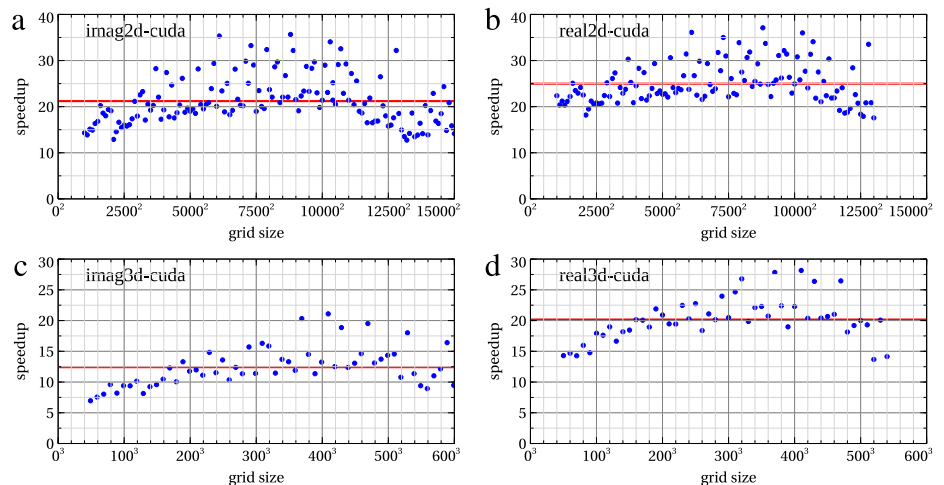
We tested our programs at the PARADOX supercomputing facility at the Scientific Computing Laboratory of the Institute of Physics Belgrade. Nodes used for testing had Intel Xeon E5-2670 CPUs with 32 GB of RAM and Nvidia Tesla M2090 GPU with 6 GB of RAM. Figure 2 shows the speedup obtained for six DBEC-GP-CUDA programs compared to their previous versions [1] executed on a single CPU core. Profiling reveals that the execution time is dominated by execution of FFTs and that the speedup varies significantly with changing of the grid size. This is due to FFT libraries used (FFTW in previous CPU version [1] and cuFFT in this version), which use different algorithms for different input array sizes. We thus conclude that the best performance can be achieved by experimenting with different grid sizes around the desired target.



**Fig. 1.** Illustration of placement of relevant variables in CPU and GPU memory. CPU initializes its own wave function tensor/matrix ( $psi$ ), trap potential ( $pot$ ) and dipolar potential ( $potdd$ ), which is copied to GPU memory. Depending on value of POTMEM variable, GPU will either allocate the same tensors/matrices for trap and dipolar potential (POTMEM = 2), allocate only one tensor/matrix and use it for different purposes (POTMEM = 1), or will map  $pot$  and  $potdd$  from CPU and not allocate extra memory on GPU (POTMEM = 0). Additionally, GPU allocates one complex tensor/matrix which is used for temporary data. This tensor/matrix is used either as a single complex tensor/matrix, or is divided into two double tensors/matrices which can then each contain the same number of elements as the complex tensor/matrix.

#### Restrictions:

Programs will only run on computers with Nvidia GPU card (Tesla or GeForce) with Compute Capability 2.0 or higher (Fermi architecture and newer) and with CUDA toolkit installed (version 6.0 or higher).



**Fig. 2.** Speedup in execution time of `imag2dXY-cuda` and `imag2dXZ-cuda` (top-left), `real2dXY-cuda` and `real2dXZ-cuda` (top-right), `imag3d-cuda` (bottom-left) and `real3d-cuda` (bottom-right) compared to the previous versions of programs [1] executed on a single CPU core. Solid red line represents average speedup obtained. We tested linear grid sizes starting from  $50^3$  in 3d and  $1000^2$  in 2d, up to the maximum that could fit in GPU memory, which was  $600^3$  for `imag3d-cuda`,  $540^3$  for `real3d-cuda`,  $15000^2$  for `imag2dXY-cuda` and `imag2dXZ-cuda`, and  $13000^2$  for `real2dXY-cuda` and `real2dXZ-cuda`. Note that the dispersion of data is due to the use of `FFT_W_ESTIMATE` flag in library calls to FFTW in the CPU programs.

#### Unusual features of all programs:

As part of the memory usage optimizations, programs may slightly increase the number of spatial grid points in each dimension ( $NX, NY, NZ$ ). This is due to FFT algorithms of cuFFT library that require additional memory to store temporary results. Our programs reuse already allocated memory to provide cuFFT with the temporary memory it requires, however, some problem sizes require much more memory, up to eight times more [7]. For instance, if the number of grid points in any dimension is a large prime number, cuFFT

uses an algorithm that requires eight times more memory than similarly sized power of two number. Adjustments of the number of grid points made in the programs ensure that cuFFT will not require such significantly increased additional memory. In case the programs perform the adjustments to grid size, this is reported in the output.

*Additional comments:*

This package consists of 6 programs, see Program title above. For the particular purpose of each program, please see descriptions below.

*Running time:*

Example inputs provided with the programs take less than one minute on Nvidia Tesla M2090 GPU.

Program summary (i)

*Program title:* imag2dXY-cuda.

*Title of electronic files:* imag2dXY-cuda.cu and imag2dXY-cuda.cuh.

*Maximum RAM memory:* No upper bound.

*Programming language used:* CUDA C.

*Typical running time:* Minutes on a medium PC.

*Nature of physical problem:* This program is designed to solve the time-dependent dipolar nonlinear partial differential GP equation in two space dimensions in an anisotropic harmonic trap. The GP equation describes the properties of a dilute trapped Bose–Einstein condensate.

*Method of solution:* The time-dependent GP equation is solved by the split-step Crank–Nicolson method by discretizing in space and time. The discretized equation is then solved by propagation in imaginary time over small time steps. The method yields the solution of stationary problems.

Program summary (ii)

*Program title:* imag2dXZ-cuda.

*Title of electronic files:* imag2dXZ-cuda.cu and imag2dXZ-cuda.cuh.

*Maximum RAM memory:* No upper bound.

*Programming language used:* CUDA C.

*Typical running time:* Minutes on a medium PC.

*Nature of physical problem:* This program is designed to solve the time-dependent dipolar nonlinear partial differential GP equation in two space dimensions in an anisotropic harmonic trap. The GP equation describes the properties of a dilute trapped Bose–Einstein condensate.

*Method of solution:* The time-dependent GP equation is solved by the split-step Crank–Nicolson method by discretizing in space and time. The discretized equation is then solved by propagation in imaginary time over small time steps. The method yields the solution of stationary problems.

Program summary (iii)

*Program title:* imag3d-cuda.

*Title of electronic files:* imag3d-cuda.cu and imag3d-cuda.cuh.

*Maximum RAM memory:* No upper bound.

*Programming language used:* CUDA C.

*Typical running time:* Tens of minutes on a medium PC.

*Nature of physical problem:* This program is designed to solve the time-dependent dipolar nonlinear partial differential GP equation in three space dimensions in an anisotropic harmonic trap. The GP equation describes the properties of a dilute trapped Bose–Einstein condensate.

*Method of solution:* The time-dependent GP equation is solved by the split-step Crank–Nicolson method by discretizing in space and time. The discretized equation is then solved by propagation in imaginary time over small time steps. The method yields the solution of stationary problems.

Program summary (iv)

*Program title:* real2dXY-cuda.

*Title of electronic files:* real2dXY-cuda.cu and real2dXY-cuda.cuh.

*Maximum RAM memory:* No upper bound.

*Programming language used:* CUDA C.

*Typical running time:* Tens of minutes on a good workstation.

*Unusual feature:* If NSTP = 0, the program requires and reads the file imag2dXY-den.txt, generated by executing imag2dXY-cuda with the same grid size parameters.

*Nature of physical problem:* This program is designed to solve the time-dependent dipolar nonlinear partial differential GP equation in two space dimensions in an anisotropic harmonic trap. The GP equation describes the properties of a dilute trapped Bose–Einstein condensate.

*Method of solution:* The time-dependent GP equation is solved by the split-step Crank–Nicolson method by discretizing in space and time. The discretized equation is then solved by propagation in real time over small time steps. The method yields the solution of dynamical problems.

Program summary (v)

*Program title:* real2dXZ-cuda.

*Title of electronic files:* real2dXZ-cuda.cu and real2dXZ-cuda.cuh.

*Maximum RAM memory:* No upper bound.

*Programming language used:* CUDA C.

*Typical running time:* Tens of minutes on a good workstation.

*Unusual feature:* If  $NSTP = 0$ , the program requires and reads the file `imag2dXZ-den.txt`, generated by executing `imag2dXZ-cuda` with the same grid size parameters.

*Nature of physical problem:* This program is designed to solve the time-dependent dipolar nonlinear partial differential GP equation in two space dimensions in an anisotropic harmonic trap. The GP equation describes the properties of a dilute trapped Bose–Einstein condensate.

*Method of solution:* The time-dependent GP equation is solved by the split-step Crank–Nicolson method by discretizing in space and time. The discretized equation is then solved by propagation in real time over small time steps. The method yields the solution of dynamical problems.

Program summary (vi)

*Program title:* `real3d-cuda`.

*Title of electronic files:* `real3d-cuda.cu` and `real3d-cuda.cuh`.

*Maximum RAM memory:* No upper bound.

*Programming language used:* CUDA C.

*Typical running time:* Tens of minutes on a good workstation.

*Unusual feature:* If  $NSTP = 0$ , the program requires and reads the file `imag3d-den.txt`, generated by executing `imag3d-cuda` with the same grid size parameters.

*Nature of physical problem:* This program is designed to solve the time-dependent dipolar nonlinear partial differential GP equation in three space dimensions in an anisotropic harmonic trap. The GP equation describes the properties of a dilute trapped Bose–Einstein condensate.


*Method of solution:* The time-dependent GP equation is solved by the split-step Crank–Nicolson method by discretizing in space and time. The discretized equation is then solved by propagation in real time over small time steps. The method yields the solution of dynamical problems.

#### Acknowledgments

V. L., A. B., A. B., and S. Š acknowledge support by the Ministry of Education, Science, and Technological Development of the Republic of Serbia under projects ON171017, III43007, ON174023, and IBEC, and by the DAAD - German Academic and Exchange Service under project IBEC. P.M. acknowledges support by the Science and Engineering Research Board, Department of Science and Technology, Government of India under project No. EMR/2014/000644. S.K.A. acknowledges support by the CNPq of Brazil under project 303280/2014-0, and by the FAPESP of Brazil under project 2012/00451-0. Numerical simulations were run on the PARADOX supercomputing facility at the Scientific Computing Laboratory of the Institute of Physics Belgrade, supported in part by the Ministry of Education, Science, and Technological Development of the Republic of Serbia under project ON171017.

#### References

- [1] R. Kishor Kumar, L. E. Young-S., D. Vudragović, A. Balaž, P. Muruganandam, and S. K. Adhikari, Fortran and C programs for the time-dependent dipolar Gross–Pitaevskii equation in an anisotropic trap, *Comput. Phys. Commun.* **195** (2015) 117.
- [2] P. Muruganandam and S. K. Adhikari, *Comput. Phys. Commun.* **180** (2009) 1888; D. Vudragović, I. Vidanović, A. Balaž, P. Muruganandam, and S. K. Adhikari, *Comput. Phys. Commun.* **183** (2012) 2021; P. Muruganandam and S. K. Adhikari, *J. Phys. B: At. Mol. Opt. Phys.* **36** (2003) 2501.
- [3] R. Kishor Kumar, P. Muruganandam, and B. A. Malomed, *J. Phys. B: At. Mol. Opt. Phys.* **46** (2013) 175302; S. K. Adhikari, Bright dipolar Bose–Einstein-condensate soliton mobile in a direction perpendicular to polarization, *Phys. Rev. A* **90** (2014) 055601; S. K. Adhikari, Stable matter-wave solitons in the vortex core of a uniform condensate, *J. Phys. B: At. Mol. Opt. Phys.* **48** (2015) 165303; S. K. Adhikari, Stable spatial and spatiotemporal optical soliton in the core of an optical vortex, *Phys. Rev. E* **92** (2015) 042926; T. Khellil, A. Balaž, and A. Pelster, Dirty bosons in a quasi-one-dimensional harmonic trap, e-print arXiv:1510.04985 (2015).
- [4] M. Harris, CUDA Pro Tip: Write Flexible Kernels with Grid-Stride Loops, Parallel Forall Blog, <http://devblogs.nvidia.com/parallelforall/cuda-pro-tip-write-flexible-kernels-grid-stride-loops/> (2013).
- [5] G. E. Blelloch, Prefix Sums and Their Applications, In J. H. Reif (Ed.), *Synthesis of Parallel Algorithms*, Morgan Kaufmann, San Francisco (1990).
- [6] M. Harris, Parallel Prefix Sum (Scan) with CUDA, EECS 570 Parallel Computer Architecture Course, University of Michigan, <http://www.eecs.umich.edu/courses/eecs570/hw/prefix.pdf> (2007).
- [7] cuFFT, CUDA API References, CUDA Toolkit Documentation v7.5, <http://docs.nvidia.com/cuda/cufft/> (2015).



# OpenMP, OpenMP/MPI, and CUDA/MPI C programs for solving the time-dependent dipolar Gross–Pitaevskii equation



Vladimir Lončar<sup>a,\*</sup>, Luis E. Young-S.<sup>b,c</sup>, Srdjan Škrbić<sup>d</sup>, Paulsamy Muruganandam<sup>e</sup>, Sadhan K. Adhikari<sup>c</sup>, Antun Balaž<sup>a</sup>

<sup>a</sup> Scientific Computing Laboratory, Center for the Study of Complex Systems, Institute of Physics Belgrade, University of Belgrade, Pregrevica 118, 11080 Belgrade, Serbia

<sup>b</sup> Departamento de Ciencias Básicas, Universidad Santo Tomás, 150001 Tunja, Boyacá, Colombia

<sup>c</sup> Instituto de Física Teórica, UNESP—Universidade Estadual Paulista, 01.140-70 São Paulo, São Paulo, Brazil

<sup>d</sup> Department of Mathematics and Informatics, Faculty of Sciences, University of Novi Sad, Trg Dositeja Obradovića 4, 21000 Novi Sad, Serbia

<sup>e</sup> School of Physics, Bharathidasan University, Palkalaiperur Campus, Tiruchirappalli – 620024, Tamil Nadu, India

## ARTICLE INFO

### Article history:

Received 13 July 2016

Accepted 15 July 2016

Available online 6 September 2016

### Keywords:

Bose–Einstein condensate

Dipolar atoms

Gross–Pitaevskii equation

Split-step Crank–Nicolson scheme

C program

OpenMP

GPU

CUDA program

MPI

## ABSTRACT

We present new versions of the previously published C and CUDA programs for solving the dipolar Gross–Pitaevskii equation in one, two, and three spatial dimensions, which calculate stationary and non-stationary solutions by propagation in imaginary or real time. Presented programs are improved and parallelized versions of previous programs, divided into three packages according to the type of parallelization. First package contains improved and threaded version of sequential C programs using OpenMP. Second package additionally parallelizes three-dimensional variants of the OpenMP programs using MPI, allowing them to be run on distributed-memory systems. Finally, previous three-dimensional CUDA-parallelized programs are further parallelized using MPI, similarly as the OpenMP programs. We also present speedup test results obtained using new versions of programs in comparison with the previous sequential C and parallel CUDA programs. The improvements to the sequential version yield a speedup of 1.1–1.9, depending on the program. OpenMP parallelization yields further speedup of 2–12 on a 16-core workstation, while OpenMP/MPI version demonstrates a speedup of 11.5–16.5 on a computer cluster with 32 nodes used. CUDA/MPI version shows a speedup of 9–10 on a computer cluster with 32 nodes.

### New version program summary

**Program Title:** DBEC-GP-OMP-CUDA-MPI: (1) DBEC-GP-OMP package: (i) imag1dX-th, (ii) imag1dZ-th, (iii) imag2dXY-th, (iv) imag2dXZ-th, (v) imag3d-th, (vi) real1dX-th, (vii) real1dZ-th, (viii) real2dXY-th, (ix) real2dXZ-th, (x) real3d-th; (2) DBEC-GP-MPI package: (i) imag3d-mpi, (ii) real3d-mpi; (3) DBEC-GP-MPI-CUDA package: (i) imag3d-mpicuda, (ii) real3d-mpicuda.

**Program Files doi:** <http://dx.doi.org/10.17632/j3z9z379m8.1>

**Licensing provisions:** Apache License 2.0

**Programming language:** OpenMP C; CUDA C.

**Computer:** DBEC-GP-OMP runs on any multi-core personal computer or workstation with an OpenMP-capable C compiler and FFTW3 library installed. MPI versions are intended for a computer cluster with a recent MPI implementation installed. Additionally, DBEC-GP-MPI-CUDA requires CUDA-aware MPI implementation installed, as well as that a computer or a cluster has Nvidia GPU with Compute Capability 2.0 or higher, with CUDA toolkit (minimum version 7.5) installed.

**Number of processors used:** All available CPU cores on the executing computer for OpenMP version, all available CPU cores across all cluster nodes used for OpenMP/MPI version, and all available Nvidia GPUs across all cluster nodes used for CUDA/MPI version.

**Journal reference of previous version:** Comput. Phys. Commun. **195** (2015) 117; *ibid.* **200** (2016) 406.

**Does the new version supersede the previous version?:** Not completely. OpenMP version does supersede previous AEWL\_v1\_0 version, while MPI versions do not supersede previous versions and are meant for execution on computer clusters and multi-GPU workstations.

\* Corresponding author.

E-mail addresses: [vladimir.loncar@ipb.ac.rs](mailto:vladimir.loncar@ipb.ac.rs) (V. Lončar), [luisevery@gmail.com](mailto:luisevery@gmail.com) (L.E. Young-S.), [srdjan.skrbic@dmi.uns.ac.rs](mailto:srdjan.skrbic@dmi.uns.ac.rs) (S. Škrbić), [anand@cnld.bdu.ac.in](mailto:anand@cnld.bdu.ac.in) (P. Muruganandam), [adhikari@ift.unesp.br](mailto:adhikari@ift.unesp.br) (S.K. Adhikari), [antun.balaz@ipb.ac.rs](mailto:antun.balaz@ipb.ac.rs) (A. Balaž).

*Nature of problem:* These programs are designed to solve the time-dependent nonlinear partial differential Gross–Pitaevskii (GP) equation with contact and dipolar interaction in a harmonic anisotropic trap. The GP equation describes the properties of a dilute trapped Bose–Einstein condensate. OpenMP package contains programs for solving the GP equation in one, two, and three spatial dimensions, while MPI packages contain only three-dimensional programs, which are computationally intensive or memory demanding enough to require such level of parallelization.

*Solution method:* The time-dependent GP equation is solved by the split-step Crank–Nicolson method by discretizing in space and time. The discretized equation is then solved by propagation, in either imaginary or real time, over small time steps. The contribution of the dipolar interaction is evaluated by a Fourier transformation to momentum space using a convolution theorem. MPI parallelization is done using the domain decomposition. The method yields the solution of stationary and/or non-stationary problems.

*Reasons for the new version:* Previously published C and Fortran programs [1] for solving the dipolar GP equation are sequential in nature and do not exploit the multiple cores or CPUs found in typical modern computers. A parallel implementation exists, using Nvidia CUDA [2], and both versions are already used within the ultra-cold atoms community [3]. However, CUDA version requires special hardware, which limits its usability. Furthermore, many researchers have access to high performance computer clusters, which could be used to either further speed up the computation, or to work with problems which cannot fit into a memory of a single computer. In light of these observations, we have parallelized all programs using OpenMP, and then extended the parallelization of three-dimensional programs using MPI to distributed-memory clusters. Since the CUDA implementation uses the same algorithm, and thus has the same structure and flow, we have applied the same data distribution scheme to provide the distributed-memory CUDA/MPI implementation of three-dimensional programs.


*Summary of revisions:*

**Package DBEC-GP-OMP:** Previous serial C programs [1] are here improved and then parallelized using OpenMP (package DBEC-GP-OMP). The main improvement consists of switching to real-to-complex (R2C) Fourier transform, which is possible due to the fact that input of the transform is purely real. In this case the result of the transform has Hermitian symmetry, where one half of the values are complex conjugates of the other half. The fast Fourier transformation (FFT) libraries we use can exploit this to compute the result faster, using half the memory.

To parallelize the programs, we have used OpenMP with the same approach as described in [4], and extended the parallelization routines to include the computation of the dipolar term. The FFT, used in computation of the dipolar term, was also parallelized in a straightforward manner, by using the built-in support for OpenMP in FFTW3 library [5]. With the introduction of multiple threads memory usage has increased, driven by the need to have some variables private to each thread. To reduce the memory consumed, we resorted to using techniques similar to the ones used in our CUDA implementation [2], i.e., we have reduced the memory required for FFT by exploiting the aforementioned R2C FFT, and reused the memory with pointer aliases whenever possible.

**Package DBEC-GP-MPI:** Next step in the parallelization (package DBEC-GP-MPI) was to extend the programs to run on distributed-memory systems, i.e., on computer clusters using domain decomposition with MPI programming paradigm. We chose to use the newly-implemented threaded versions of the programs as the starting point. Alternatively, we could have used serial versions, and attempt a pure MPI parallelization, however we have found that OpenMP-parallelized routines better exploit the data locality and thus outperform the pure MPI implementation. Therefore, our OpenMP/MPI-parallelized programs are intended to run one MPI process per cluster node, and each process would spawn the OpenMP threads as needed on its cluster node. Note that this is not a requirement, and users may run more than one MPI process per node, but we advise against it due to performance reasons. With the suggested execution strategy (one MPI process per cluster node, each spawning as many threads as CPU cores available), OpenMP threads perform most of the computation, and MPI is used for data exchanges between processes.

There are numerous ways to distribute the data between MPI processes, and we decided to use a simple one-dimensional data distribution, also known as slab decomposition. Data is distributed along the first (slowest changing) dimension, which corresponds to NX spatial dimension in our programs (see Fig. 1). Each process is assigned a different portion of the NX dimension, and contains the entire NY and NZ spatial dimensions locally. This allows each process to perform computation on those two dimensions in the same way as before, without any data exchanges. In case the computation requires whole NX dimension to be local to each process, we transpose the data, and after the computation, we transpose the data back.




**Fig. 1.** Illustration of data distribution between MPI processes. On the left, the data are distributed along the NX dimension, while on the right the same data are redistributed along the NY dimension.

Transpose routine can be implemented in many ways using MPI, most commonly using `MPI_Alltoall` function, or using transpose routines from external libraries, like FFTW3 [5] or 2DECOMP&FFT [6]. Since we already rely on FFTW3 library for FFT, we have utilized its dedicated



transpose interface to perform the necessary transformations. To speed up transpose operation, we do not perform full transposition of data, but rather leave it *locally* transposed. That is, we transform from  $\text{local\_NX} \times \text{NY} \times \text{NZ}$ , stored in row-major order, to  $\text{NX} \times \text{local\_NY} \times \text{NZ}$  in row-major order (where  $\text{local\_NX} = \text{NX} / \text{number\_of\_processes}$ , and equivalently for  $\text{local\_NY}$ ). This approach has an additional benefit that we do not have to make significant changes in the way array elements are processed, and in most cases we only have to adjust the loop limit of the non-local dimension.


**Package DBEC-GP-MPI-CUDA:** The aforementioned data distribution scheme can be also applied to the CUDA version of programs [2]. However, there is no support for CUDA in FFTW3, and cuFFT (used in CUDA programs for FFT) does not provide equivalent MPI or transpose interface. Instead, we developed our own transpose routines, and used them in FFT computation. One example of manual implementation of transpose routines is shown in Ref. [7], and while we could readily use the same code, we wanted to have the same result as when using FFTW3. To achieve this, we use the same basic principle as in Ref. [7], first we create a custom MPI data type that maps to portions of the data to be exchanged, followed by an all-to-all communication to exchange the data between processes, see Fig. 2 for details.



**Fig. 2.** Example of a transpose routine of a  $4 \times 4 \times 4$  data between four MPI processes. Initially, all processes have  $1/4$  of the  $\text{NX}$  dimension, and whole  $\text{NY}$  and  $\text{NZ}$  dimensions. After transposing, each process has full  $\text{NX}$  and  $\text{NZ}$  dimensions, and  $1/4$  of the  $\text{NY}$  dimension.

The implemented transpose routines are also used to compute a distributed-memory FFT, performed over all MPI processes. To divide the computation of a multidimensional FFT, in our case three-dimensional, we use a well-known row-column algorithm. The basic idea of the algorithm is perhaps best explained on a two-dimensional FFT of  $N \times M$  data, stored in row-major order, illustrated in Fig. 3. First the  $N$  one-dimensional FFTs of length  $M$  are performed (along the row of data), followed by a transpose, after which data are stored as  $M \times N$  in row-major format. Now  $M$  FFTs of length  $N$  can be performed along what used to be a column of original data, but are stored as rows after transposing. Finally, an optional transpose can be performed to return the data in their original  $N \times M$  form. In three dimensions, we can perform a two-dimensional FFT, transpose the data, and perform the FFT along the third dimension. This algorithm can be easily adapted for distributed memory systems. We use advanced cuFFT interface for local computation of FFT, and use our transpose routine to redistribute the data.

Note that DBEC-GP-MPI-CUDA programs can be easily modified to work on a single workstation with multiple GPU cards, or a computer cluster with multiple GPU cards per node. In that case, for each GPU card a separate MPI process should be launched and the programs should be modified to assign a separate GPU card for processes on the same cluster node.



**Fig. 3.** Illustration of four stages of row-column FFT algorithm. The last transpose operation may be omitted, and often yields better performance.

**MPI output format:** Given that the distributed memory versions of the programs can be used for much larger grid sizes, the output they produce (i.e., the density profiles) can be much larger and difficult to handle. To alleviate this problem somewhat, we have switched to a binary output instead of the textual. This allowed us to reduce the size of files, while still retaining precision. All MPI processes will write the output to the same file, at the corresponding offset, relieving the user of the task of combining the files. The binary output can be subsequently converted to textual, for example by using `hexdump` command on UNIX-like systems. We have developed a simple script which converts the output from binary to textual format and included it in the software package.

**Testing results:** We have tested all programs on the PARADOX supercomputing facility at the Scientific Computing Laboratory of the Institute of Physics Belgrade. Nodes used for testing had two Intel Xeon E5-2670 CPUs (with a total of  $2 \times 8 = 16$  CPU cores) with 32 GB of RAM and one Nvidia Tesla M2090 GPU with 6 GB of RAM, each connected by Infiniband QDR interconnect. The presented results are obtained for arbitrary grid sizes, which are not tailored to maximize performance of the programs. We also stress that execution times and speedups reported here are calculated for critical parallelized parts of the programs

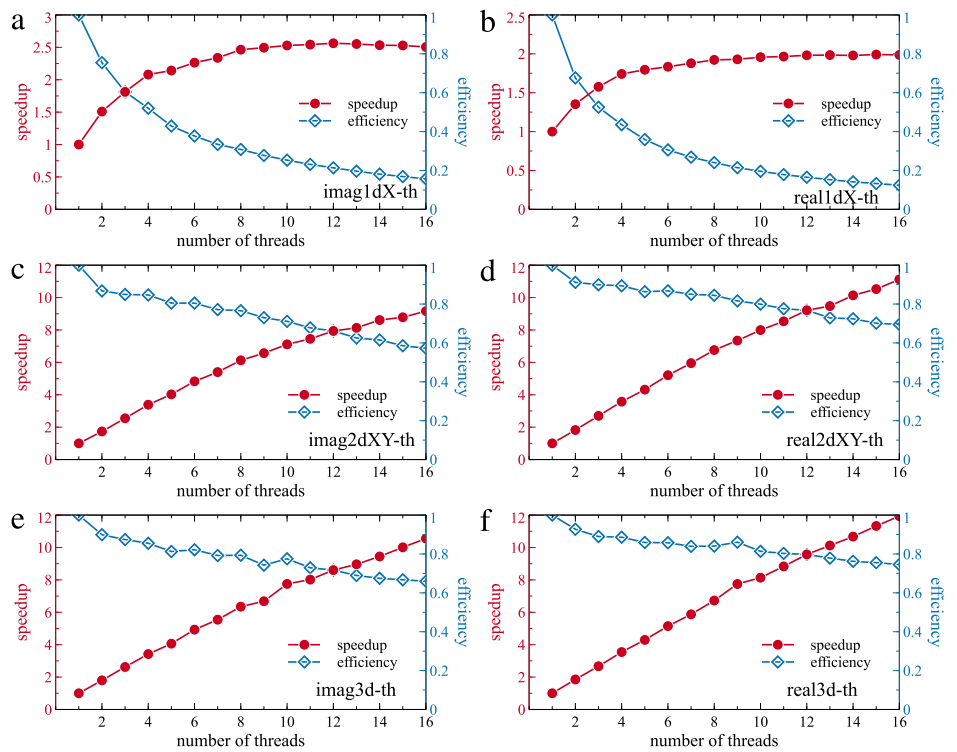
performing iterations over imaginary or real time steps, and they exclude time spent on initialization (threads initialization, MPI environment, allocation/deallocation of memory, creating/destroying FFTW plans, I/O operations). As a part of its output, each program separately prints initialization time and time spent on iterations for GP propagation. The latter time is used to calculate a speedup, as a speedup obtained this way does not depend on the number of iterations and is more useful for large numbers of iterations.

The testing of OpenMP versions of programs DBEC-GP-OMP was performed with the number of threads varying from 1 to 16. Table 1 and Fig. 4 show the obtained absolute wall-clock times, speedups, and scaling efficiencies, as well as comparison with the previous serial version of programs [1]. As we can see from the table, improvements in the FFT routine used already yield a speedup of 1.3 to 1.9 for single-threaded ( $T = 1$ ) 2d and 3d programs compared to the previous serial programs, and somewhat smaller speedup for 1d programs, 1.1 to 1.3. The use of additional threads brings about further speedup of 2 to 2.5 for 1d programs, and 9 to 12 for 2d and 3d programs. From Fig. 4 we see that for 1d programs, although speedup increases with the number of threads used, the efficiency decreases due to insufficient size of the problem, and one can achieve almost maximal value of speedup already with  $T = 4$  threads, while still keeping the efficiency around 50%. We also see, as expected, that speedup and efficiency of 2d and 3d programs behave quite well as we increase the numbers of threads. In particular, we note that the efficiency is always above 60%, making the use of all available CPU cores worthwhile.

**Table 1**

Wall-clock execution times of DBEC-GP-OMP programs compiled with Intel's icc compiler, compared to the execution times of previously published serial versions. The execution times given here are for 1000 iterations (in seconds, excluding initialization and input/output operations, as reported by each program) with grid sizes:  $10^5$  for 1d programs,  $10^4 \times 10^4$  for 2d programs, and  $480 \times 480 \times 480$  for 3d programs. Columns  $T = 1$ ,  $T = 2$ ,  $T = 4$ ,  $T = 8$ , and  $T = 16$  correspond to the number of threads used, while the last column shows the obtained speedup with 16 OpenMP threads ( $T = 16$ ) compared to one OpenMP thread ( $T = 1$ ). Note that the reduction in the execution time is not solely due to the introduction of multiple threads, as the improvements in the FFT routine used also have noticeable impact. This is most evident when comparing execution times of serial versions to OpenMP versions with one thread. Execution times and speedups of imag1dZ-th, real1dZ-th, imag2dXZ-th, and real2dXZ-th (not reported here) are similar to those of imag1dX-th, real1dX-th, imag2dXY-th, and real2dXY-th, respectively.

	Serial [1]	$T = 1$	$T = 2$	$T = 4$	$T = 8$	$T = 16$	Speedup
imag1dX-th	9.1	7.1	4.7	3.4	2.9	2.8	2.5
real1dX-th	15.2	14.2	10.5	8.2	7.3	7.2	2.0
imag2dXY-th	13657	7314	4215	2159	1193	798	9.2
real2dXY-th	17281	11700	6417	3271	1730	1052	11.1
imag3d-th	16064	9353	5201	2734	1473	888	10.5
real3d-th	22611	17496	9434	4935	2602	1466	11.9



**Fig. 4.** Speedup in the execution time and scaling efficiency of DBEC-GP-OMP programs compared to single-threaded runs: (a) imag1dX-th, (b) real1dX-th, (c) imag2dXY-th, (d) real2dXY-th, (e) imag3d-th, (f) real3d-th. Scaling efficiency is calculated as a fraction of the obtained speedup compared to a theoretical maximum. Grid sizes used for testing are the same as in Table 1. Speedups and efficiencies of imag1dZ-th, real1dZ-th, imag2dXZ-th, and real2dXZ-th (not reported here) are similar to those of imag1dX-th, real1dX-th, imag2dXY-th, and real2dXY-th, respectively.




For testing of MPI versions we have used a similar methodology to measure the strong scaling performance. For OpenMP/MPI programs DBEC-GP-MPI, the obtained wall-clock times are shown in Table 2, together with the corresponding wall-clock times for the OpenMP programs DBEC-GP-OMP that served as a baseline to calculate speedups. The testing was done for varying number of cluster nodes, from 4 to 32, and the measured speedup ranged from 11 to 16.5. The corresponding graphs of speedups and efficiencies are shown in Fig. 5, where we can see that the speedup grows linearly with the number of nodes used, while the efficiency remains mostly constant in the range between 40% and 60%, thus making the use of OpenMP/MPI programs highly advantageous for problems with large grid sizes.

**Table 2**

Wall-clock execution times of DBEC-GP-MPI programs compiled with mpicc compiler from OpenMPI implementation of MPI, backed by Intel's icc compiler, compared to the execution times of OpenMP (DBEC-GP-OMP) versions on a single-node ( $T = 16$ ,  $N = 1$ ). The execution times given here are for 1000 iterations (in seconds, excluding initialization and input/output operations, as reported by each program) with the grid size  $480 \times 480 \times 500$ . Columns  $N = 4$ ,  $N = 8$ ,  $N = 16$ ,  $N = 24$ , and  $N = 32$  correspond to the number of cluster nodes used (each with  $T = 16$  threads), while the last column shows the obtained speedup with  $N = 32$  nodes compared to single-node runs.

	OpenMP	$N = 4$	$N = 8$	$N = 16$	$N = 24$	$N = 32$	Speedup
imag3d-mpi	1124	653	352	167	128	96	11.5
real3d-mpi	2140	979	513	277	220	129	16.5




**Fig. 5.** Speedup in the execution time and scaling efficiency of DBEC-GP-MPI programs compared to single-node OpenMP runs: (a) imag3d-mpi, (b) real3d-mpi. Scaling efficiency is calculated as a fraction of the obtained speedup compared to a theoretical maximum. Grid size used for testing is the same as in Table 2.

For CUDA/MPI programs DBEC-GP-MPI-CUDA we observe similar behavior in Table 3 and in Fig. 6. The obtained speedup with  $N = 32$  nodes here ranges from 9 to 10, with the efficiency between 30% and 40%. While the efficiency is slightly lower than in the case of OpenMP/MPI programs, which could be expected due to a more complex memory hierarchy when dealing with the multi-GPU system distributed over many cluster nodes, the speedup still grows linearly and makes CUDA/MPI programs ideal choice for use on GPU-enabled computer clusters. Additional benefit of using these programs is their low CPU usage (up to one CPU core), allowing for the possibility that same cluster nodes are used for other CPU-intensive simulations.

**Table 3**

Wall-clock execution times of DBEC-GP-MPI-CUDA programs compiled with Nvidia's nvcc compiler, with CUDA-aware OpenMPI implementation of MPI, backed by Intel's icc compiler, compared to the execution times of previous CUDA [2] versions on a single-node with one GPU card ( $N = 1$ ). The execution times given here are for 1000 iterations (in seconds, excluding initialization and input/output operations, as reported by each program) with the grid size  $480 \times 480 \times 250$ . Columns  $N = 4$ ,  $N = 8$ ,  $N = 16$ ,  $N = 24$ , and  $N = 32$  correspond to the number of cluster nodes used (each with one GPU card), while the last column shows the obtained speedup with  $N = 32$  nodes compared to single-node runs.

	CUDA [2]	$N = 4$	$N = 8$	$N = 16$	$N = 24$	$N = 32$	Speedup
imag3d-mpicuda	579	447	212	103	71	61	9.5
real3d-mpicuda	800	619	295	142	96	80	9.9



**Fig. 6.** Speedup in the execution time and scaling efficiency of DBEC-GP-MPI-CUDA programs compared to single-node runs of previous CUDA programs [2]: (a) imag3d-mpicuda, (b) real3d-mpicuda. Scaling efficiency is calculated as a fraction of the obtained speedup compared to a theoretical maximum. Grid size used for testing is the same as in Table 3.

The introduction of distributed transposes of data creates some overhead, which negatively impacts scaling efficiency. This is more evident in the CUDA/MPI version, as the transpose algorithm is inferior to the one provided by FFTW3. In our tests, both MPI versions of programs failed to achieve speedup on

less than 4 nodes, due to the introduction of the transpose routines. We therefore recommend using MPI versions only on 4 or more cluster nodes.

The MPI versions are highly dependent not only on the configuration of the cluster, mainly on the speed of interconnect, but also on the distribution of processes and threads, NUMA configuration, etc. We recommend that users experiment with several different configurations to achieve the best performance. The results presented are obtained without extensive tuning, with the aim to show the base performance.

Finally, we note that the best performance can be achieved by evenly distributing the workload among the MPI processes and OpenMP threads, and by using grid sizes which are optimal for FFT. In particular, the programs in DBEC-GP-OMP package have the best performance if NX, NY, and NZ are divisible by the number of OpenMP threads used. Similarly, for DBEC-GP-MPI programs the best performance is achieved if NX and NY are divisible by a product of the number of MPI processes and the number of OpenMP threads used. For DBEC-GP-MPI-CUDA programs, the best performance is achieved if NX and NY are divisible by a product of the number of MPI processes and the number of Streaming Multiprocessors (SM) in the GPU used. For all three packages, the best FFT performance is obtained if NX, NY and NZ can be expressed as  $2^a 3^b 5^c 7^d 11^e 13^f$ , where  $e$  and  $f$  are either 0 or 1, and the other exponents are non-negative integer numbers [8].

*Additional comments, restrictions, and unusual features:* MPI programs require that grid size (controlled by input parameters NX, NY and NZ) can be evenly distributed between the processes, i.e., that NX and NY are divisible by the number of MPI processes. Since the data is never distributed along the NZ dimension, there is no such requirement on NZ. Programs will test if these conditions are met, and inform the user if not (by reporting an error). Additionally, MPI versions of CUDA programs require CUDA-aware MPI implementation. This allows the MPI runtime to directly access GPU memory pointers and avoid having to copy the data to main RAM. List of CUDA-aware MPI implementations can be found in Ref. [9].

#### Acknowledgments

V.L., S.Š. and A.B. acknowledge support by the Ministry of Education, Science, and Technological Development of the Republic of Serbia under projects ON171017, OI1611005, and III43007, as well as SCOPEs project IZ74Z0-160453. L.E. Y.-S. acknowledges support by the FAPESP of Brazil under project 2012/21871-7 and 2014/16363-8. P.M. acknowledges support by the Science and Engineering Research Board, Department of Science and Technology, Government of India under project No. EMR/2014/000644. S.K.A. acknowledges support by the CNPq of Brazil under project 303280/2014-0, and by the FAPESP of Brazil under project 2012/00451-0.

#### References

- [1] R. Kishor Kumar, L. E. Young-S., D. Vudragović, A. Balaž, P. Muruganandam, and S. K. Adhikari, *Comput. Phys. Commun.* **195** (2015) 117.
- [2] V. Lončar, A. Balaž, A. Bogojević, S. Škrbić, P. Muruganandam, S. K. Adhikari, *Comput. Phys. Commun.* **200** (2016) 406.
- [3] R. Kishor Kumar, P. Muruganandam, and B. A. Malomed, *J. Phys. B: At. Mol. Opt. Phys.* **46** (2013) 175302;  
H. Al-Jibbouri, I. Vidanović, A. Balaž, and A. Pelster, *J. Phys. B: At. Mol. Opt. Phys.* **46** (2013) 065303;  
R. R. Sakhel, A. R. Sakhel, and H. B. Ghassib, *J. Low Temp. Phys.* **173** (2013) 177;  
B. Nikolić, A. Balaž, and A. Pelster, *Phys. Rev. A* **88** (2013) 013624;  
X. Antoine and R. Duboscq, *Comput. Phys. Commun.* **185** (2014) 2969;  
J. Luo, *Commun. Nonlinear Sci. Numer. Simul.* **19** (2014) 3591;  
K.-T. Xi, J. Li, and D.-N. Shi, *Physica B* **436** (2014) 149;  
S. K. Adhikari, *Phys. Rev. A* **90** (2014) 055601;  
M. C. Raportaru, J. Jovanovski, B. Jakimovski, D. Jakimovski, and A. Mishev, *Rom. J. Phys.* **59** (2014) 677;  
A. I. Nicolin, A. Balaž, J. B. Sudharsan, and R. Radha, *Rom. J. Phys.* **59** (2014) 204;  
A. Balaž, R. Paun, A. I. Nicolin, S. Balasubramanian, and R. Ramaswamy, *Phys. Rev. A* **89** (2014) 023609;  
A. I. Nicolin and I. Rata, High-Performance Computing Infrastructure for South East Europe's Research Communities: Results of the HP-SEE User Forum 2012, in *Springer Series: Modeling and Optimization in Science and Technologies* **2** (2014) 15;  
S. K. Adhikari, *J. Phys. B: At. Mol. Opt. Phys.* **48** (2015) 165303;  
S. K. Adhikari, *Phys. Rev. E* **92** (2015) 042926;  
T. Khellil and A. Pelster, arXiv:1512.04870 (2015);  
H. L. C. Couto and W. B. Cardoso, *J. Phys. B: At. Mol. Opt. Phys.* **48** (2015) 025301;  
R. R. Sakhel, A. R. Sakhel, and H. B. Ghassib, *Physica B* **478** (2015) 68;  
L. Salasnich and S. K. Adhikari, *Acta Phys. Pol. A* **128** (2015) 979;  
X. Antoine and R. Duboscq, *Lecture Notes Math.* **2146** (2015) 49;  
E. Chiquillo, *J. Phys. A: Math. Theor.* **48** (2015) 475001;  
S. Sabari, C. P. Jisha, K. Porsezian, and V. A. Brazhnyi, *Phys. Rev. E* **92** (2015) 032905;  
W. Wen, T. K. Shui, Y. F. Shan, and C. P. Zhu, *J. Phys. B: At. Mol. Opt. Phys.* **48** (2015) 175301;  
P. Das and P. K. Panigrahi, *Laser Phys.* **25** (2015) 125501;  
Y. S. Wang, S. T. Ji, Y. E. Luo, and Z. Y. Li, *J. Korean. Phys. Soc.* **67** (2015) L1504;  
A. I. Nicolin, M. C. Raportaru, and A. Balaž, *Rom. Rep. Phys.* **67** (2015) 143;  
V. S. Bagnato, D. J. Frantzeskakis, P. G. Kevrekidis, B. A. Malomed, and D. Mihalache, *Rom. Rep. Phys.* **67** (2015) 5;  
J. B. Sudharsan, R. Radha, H. Fabrelli, A. Gammal, and B. A. Malomed, *Phys. Rev. A* **92** (2015) 053601;  
K.-T. Xi, J. Li, and D.-N. Shi, *Physica B* **459** (2015) 6;  
E. J. M. Madarassy and V. T. Toth, *Phys. Rev. D* **91** (2015) 044041;  
F. I. Moxley III, T. Byrnes, B. Ma, Y. Yan, and W. Dai, *J. Comput. Phys.* **282** (2015) 303;  
D. Novoa, D. Tommasini, J. A. Nóvoa-López, *Phys. Rev. E* **91** (2015) 012904;

- Y. H. Wang, A. Kumar, F. Jendrzejewski, R. M. Wilson, M. Edwards, S. Eckel, G. K. Campbell, and C. W. Clark, *New J. Phys.* **17** (2015) 125012;  
T. Khellil, A. Balaž, and A. Pelster, *New J. Phys.* **18** (2016) 063003;  
T. Khellil and A. Pelster, *J. Stat. Mech.-Theory Exp.* (2016) 063301;  
J. Akram and A. Pelster, *Phys. Rev. A* **93** (2016) 023606;  
S. K. Adhikari, *Laser Phys. Lett.* **13** (2016) 035502;  
J. Akram and A. Pelster, *Phys. Rev. A* **93** (2016) 033610;  
J. Akram, B. Girodias, and A. Pelster, *J. Phys. B: At. Mol. Opt. Phys.* **49** (2016) 075302;  
S. K. Adhikari and S. Gautam, *Phys. Rev. A* **93** (2016) 013630;  
Ž. Marojević, E. Göklü, and C. Lämmerzahl, *Comput. Phys. Commun.* **202** (2016) 216;  
A. Paredes and H. Michninel, *Phys. Dark Universe* **12** (2016) 50;  
J. Akram and A. Pelster, *Laser Phys.* **26** (2016) 065501;  
T. Mithun, K. Porsezian, and B. Dey, *Phys. Rev. A* **93** (2016) 013620;  
C.-Y. Lai and C.-C. Chien, *Phys. Rev. Appl.* **5** (2016) 034001;  
S. K. Adhikari, *Laser Phys. Lett.* **13** (2016) 085501;  
K. Manikandan, P. Muruganandam, M. Senthilvelan, and M. Lakshmanan, *Phys. Rev. E* **93** (2016) 032212;  
R. R. Sakhel, A. R. Sakhel, H. B. Ghassib, and A. Balaž, *Eur. Phys. J. D* **70** (2016) 66;  
W. Bao, Q. Tang, and Y. Zhang, *Commun. Comput. Phys.* **19** (2016) 1141;  
R. Kishor Kumar, T. Sriraman, H. Fabrelli, P. Muruganandam, and A. Gammal, *J. Phys. B: At. Mol. Opt. Phys.* **49** (2016) 155301;  
A. Bogojević, A. Balaž, and A. Belić, *Phys. Rev. E* **72** (2005) 036128;  
A. Bogojević, I. Vidanović, A. Balaž, and A. Belić, *Phys. Lett. A* **372** (2008) 3341;  
I. Vidanović, A. Bogojević, A. Balaž, and A. Belić, *Phys. Rev. E* **80** (2009) 066706;  
A. Balaž, A. Bogojević, I. Vidanović, and A. Pelster, *Phys. Rev. E* **79** (2009) 036701;  
A. Balaž, I. Vidanović, A. Bogojević, and A. Pelster, *Phys. Lett. A* **374** (2010) 1539;  
A. I. Nicolin, *Physica A* **391** (2012) 1062;  
I. Vasić and A. Balaž, arXiv:1602.03538 (2016);  
O. Voronych, A. Buraczewski, M. Matuszewski, and M. Stobińska, arXiv:1603.02570 (2016);  
A. M. Martin, N. G. Marchant, D. H. J. O'Dell, and N. G. Parker, arXiv:1606.07107 (2016).  
[4] D. Vudragović, I. Vidanović, A. Balaž, P. Muruganandam, and S. K. Adhikari, *Comput. Phys. Commun.* **183** (2012) 2021.  
[5] FFTW3 library, <http://www.fftw.org/> (2016).  
[6] 2DECOMP&FFT library, <http://www.2decomp.org/> (2016).  
[7] B. Satarić, V. Slavnić, A. Belić, A. Balaž, P. Muruganandam, S. K. Adhikari, *Comput. Phys. Commun.* **200** (2016) 411.  
[8] Real-data DFTs with FFTW3, [http://www.fftw.org/fftw3\\_doc/Real\\_002ddata-DFTs.html](http://www.fftw.org/fftw3_doc/Real_002ddata-DFTs.html) (2014);  
Nvidia's cuFFT accuracy and performance, <http://docs.nvidia.com/cuda/cufft/#accuracy-and-performance> (2015).  
[9] Nvidia's MPI Solutions for GPUs, <https://developer.nvidia.com/mpi-solutions-gpus> (2016).



# OpenMP GNU and Intel Fortran programs for solving the time-dependent Gross–Pitaevskii equation

Luis E. Young-S.<sup>a</sup>, Paulsamy Muruganandam<sup>b</sup>, Sadhan K. Adhikari<sup>c</sup>, Vladimir Lončar<sup>d</sup>, Dušan Vudragović<sup>d</sup>, Antun Balaž<sup>d,\*</sup>

<sup>a</sup> Departamento de Ciencias Básicas, Universidad Santo Tomás, 150001 Tunja, Boyacá, Colombia

<sup>b</sup> Department of Physics, Bharathidasan University, Palkalaiperur Campus, Tiruchirappalli–620024, Tamil Nadu, India

<sup>c</sup> Instituto de Física Teórica, UNESP - Universidade Estadual Paulista, 01.140-70 São Paulo, São Paulo, Brazil

<sup>d</sup> Scientific Computing Laboratory, Center for the Study of Complex Systems, Institute of Physics Belgrade, University of Belgrade, Serbia

## ARTICLE INFO

### Article history:

Received 11 July 2017

Accepted 14 July 2017

Available online 10 August 2017

### Keywords:

Bose–Einstein condensate

Gross–Pitaevskii equation

Split-step Crank–Nicolson scheme

Intel and GNU Fortran programs

Open Multi-Processing

OpenMP

Partial differential equation

## ABSTRACT

We present Open Multi-Processing (OpenMP) version of Fortran 90 programs for solving the Gross–Pitaevskii (GP) equation for a Bose–Einstein condensate in one, two, and three spatial dimensions, optimized for use with GNU and Intel compilers. We use the split-step Crank–Nicolson algorithm for imaginary- and real-time propagation, which enables efficient calculation of stationary and non-stationary solutions, respectively. The present OpenMP programs are designed for computers with multi-core processors and optimized for compiling with both commercially-licensed Intel Fortran and popular free open-source GNU Fortran compiler. The programs are easy to use and are elaborated with helpful comments for the users. All input parameters are listed at the beginning of each program. Different output files provide physical quantities such as energy, chemical potential, root-mean-square sizes, densities, etc. We also present speedup test results for new versions of the programs.

### New version program summary

**Program title:** BEC-GP-OMP-FOR software package, consisting of: (i) imag1d-th, (ii) imag2d-th, (iii) imag3d-th, (iv) imagaxi-th, (v) imagcir-th, (vi) imagsph-th, (vii) real1d-th, (viii) real2d-th, (ix) real3d-th, (x) realaxi-th, (xi) realcir-th, (xii) realsph-th.

**Program files doi:** <http://dx.doi.org/10.17632/y8zk3jgn84.2>

**Licensing provisions:** Apache License 2.0

**Programming language:** OpenMP GNU and Intel Fortran 90.

**Computer:** Any multi-core personal computer or workstation with the appropriate OpenMP-capable Fortran compiler installed.

**Number of processors used:** All available CPU cores on the executing computer.

**Journal reference of previous version:** Comput. Phys. Commun. **180** (2009) 1888; *ibid.* **204** (2016) 209.

**Does the new version supersede the previous version?:** Not completely. It does supersede previous Fortran programs from both references above, but not OpenMP C programs from Comput. Phys. Commun. **204** (2016) 209.

**Nature of problem:** The present Open Multi-Processing (OpenMP) Fortran programs, optimized for use with commercially-licensed Intel Fortran and free open-source GNU Fortran compilers, solve the time-dependent nonlinear partial differential (GP) equation for a trapped Bose–Einstein condensate in one (1d), two (2d), and three (3d) spatial dimensions for six different trap symmetries: axially and radially symmetric traps in 3d, circularly symmetric traps in 2d, fully isotropic (spherically symmetric) and fully anisotropic traps in 2d and 3d, as well as 1d traps, where no spatial symmetry is considered.

**Solution method:** We employ the split-step Crank–Nicolson algorithm to discretize the time-dependent GP equation in space and time. The discretized equation is then solved by imaginary- or real-time propagation, employing adequately small space and time steps, to yield the solution of stationary and non-stationary problems, respectively.

\* Corresponding author.

E-mail addresses: [luis.young@usantoto.edu.co](mailto:luis.young@usantoto.edu.co) (L.E. Young-S), [anand@cncld.bdu.ac.in](mailto:anand@cncld.bdu.ac.in) (P. Muruganandam), [adhikari@ift.unesp.br](mailto:adhikari@ift.unesp.br) (S.K. Adhikari), [vladimir.loncar@ipb.ac.rs](mailto:vladimir.loncar@ipb.ac.rs) (V. Lončar), [dusan.vudragovic@ipb.ac.rs](mailto:dusan.vudragovic@ipb.ac.rs) (D. Vudragović), [antun.balaz@ipb.ac.rs](mailto:antun.balaz@ipb.ac.rs) (A. Balaž).

*Reasons for the new version:* Previously published Fortran programs [1,2] have now become popular tools [3] for solving the GP equation. These programs have been translated to the C programming language [4] and later extended to the more complex scenario of dipolar atoms [5]. Now virtually all computers have multi-core processors and some have motherboards with more than one physical computer processing unit (CPU), which may increase the number of available CPU cores on a single computer to several tens. The C programs have been adopted to be very fast on such multi-core modern computers using general-purpose graphic processing units (GPGPU) with Nvidia CUDA and computer clusters using Message Passing Interface (MPI) [6]. Nevertheless, previously developed Fortran programs are also commonly used for scientific computation and most of them use a single CPU core at a time in modern multi-core laptops, desktops, and workstations. Unless the Fortran programs are made aware and capable of making efficient use of the available CPU cores, the solution of even a realistic dynamical 1d problem, not to mention the more complicated 2d and 3d problems, could be time consuming using the Fortran programs. Previously, we published auto-parallel Fortran programs [2] suitable for Intel (but not GNU) compiler for solving the GP equation. Hence, a need for the full OpenMP version of the Fortran programs to reduce the execution time cannot be overemphasized. To address this issue, we provide here such OpenMP Fortran programs, optimized for both Intel and GNU Fortran compilers and capable of using all available CPU cores, which can significantly reduce the execution time.

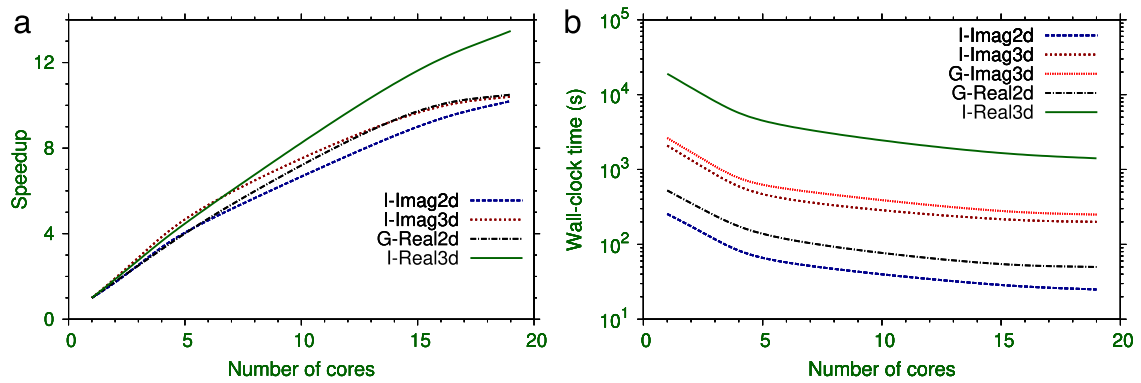
*Summary of revisions:* Previous Fortran programs [1] for solving the time-dependent GP equation in 1d, 2d, and 3d with different trap symmetries have been parallelized using the OpenMP interface to reduce the execution time on multi-core processors. There are six different trap symmetries considered, resulting in six programs for imaginary-time propagation and six for real-time propagation, totaling to 12 programs included in BEC-GP-OMP-FOR software package.

All input data (number of atoms, scattering length, harmonic oscillator trap length, trap anisotropy, etc.) are conveniently placed at the beginning of each program, as before [2]. Present programs introduce a new input parameter, which is designated by `Number_of_Threads` and defines the number of CPU cores of the processor to be used in the calculation. If one sets the value 0 for this parameter, all available CPU cores will be used. For the most efficient calculation it is advisable to leave one CPU core unused for the background system's jobs. For example, on a machine with 20 CPU cores such that we used for testing, it is advisable to use up to 19 CPU cores. However, the total number of used CPU cores can be divided into more than one job. For instance, one can run three simulations simultaneously using 10, 4, and 5 CPU cores, respectively, thus totaling to 19 used CPU cores on a 20-core computer.


The Fortran source programs are located in the directory `src`, and can be compiled by the `make` command using the `makefile` in the root directory `BEC-GP-OMP-FOR` of the software package. The examples of produced output files can be found in the directory `output`, although some large density files are omitted, to save space. The programs calculate the values of actually used dimensionless nonlinearities from the physical input parameters, where the input parameters correspond to the identical nonlinearity values as in the previously published programs [1], so that the output files of the old and new programs can be directly compared. The output files are conveniently named such that their contents can be easily identified, following the naming convention introduced in Ref. [2]. For example, a file named `<code>-out.txt`, where `<code>` is a name of the individual program, represents the general output file containing input data, time and space steps, nonlinearity, energy and chemical potential, and was named `fort.7` in the old Fortran version of programs [1]. A file named `<code>-den.txt` is the output file with the condensate density, which had the names `fort.3` and `fort.4` in the old Fortran version [1] for imaginary- and real-time propagation programs, respectively. Other possible density outputs, such as the initial density, are commented out in the programs to have a simpler set of output files, but users can uncomment and re-enable them, if needed. In addition, there are output files for reduced (integrated) 1d and 2d densities for different programs. In the real-time programs there is also an output file reporting the dynamics of evolution of root-mean-square sizes after a perturbation is introduced. The supplied real-time programs solve the stationary GP equation, and then calculate the dynamics. As the imaginary-time programs are more accurate than the real-time programs for the solution of a stationary problem, one can first solve the stationary problem using the imaginary-time programs, adapt the real-time programs to read the pre-calculated wave function and then study the dynamics. In that case the parameter `NSTP` in the real-time programs should be set to zero and the space mesh and nonlinearity parameters should be identical in both programs. The reader is advised to consult our previous publication where a complete description of the output files is given [2]. A `readme.txt` file, included in the root directory, explains the procedure to compile and run the programs.

We tested our programs on a workstation with two 10-core Intel Xeon E5-2650 v3 CPUs. The parameters used for testing are given in sample input files, provided in the corresponding directory together with the programs. In Table 1 we present wall-clock execution times for runs on 1, 6, and 19 CPU cores for programs compiled using Intel and GNU Fortran compilers. The corresponding columns "Intel speedup" and "GNU speedup" give the ratio of wall-clock execution times of runs on 1 and 19 CPU cores, and denote the actual measured speedup for 19 CPU cores. In all cases and for all numbers of CPU cores, although the GNU Fortran compiler gives excellent results, the Intel Fortran compiler turns out to be slightly faster. Note that during these tests we always ran only a single simulation on a workstation at a time, to avoid any possible interference issues. Therefore, the obtained wall-clock times are more reliable than the ones that could be measured with two or more jobs running simultaneously. We also studied the speedup of the programs as a function of the number of CPU cores used. The performance of the Intel and GNU Fortran compilers is illustrated in Fig. 1, where we plot the speedup and actual wall-clock times as functions of the number of CPU cores for 2d and 3d programs. We see that the speedup increases monotonically with the number of CPU cores in all cases and has large values (between 10 and 14 for 3d

programs) for the maximal number of cores. This fully justifies the development of OpenMP programs, which enable much faster and more efficient solving of the GP equation. However, a slow saturation in the speedup with the further increase in the number of CPU cores is observed in all cases, as expected.



**Fig. 1.** (a) Speedup for 2d and 3d programs compiled with the Intel (I) and GNU (G) Fortran compilers as a function of the number of CPU cores, measured on a workstation with two Intel Xeon E5-2650 v3 CPUs. (b) Wall-clock execution time (in seconds) of 2d and 3d programs compiled with the Intel (I) and GNU (G) Fortran compilers as a function of the number of CPU cores.



**Fig. 2.** Speedup of real2d-th program, compiled with the Intel Fortran 90 compiler and executed on 19 CPU cores on a workstation with two Intel Xeon E5-2650 v3 CPUs, as a function of the number of spatial discretization points  $NX=NY$ .

**Table 1**

Wall-clock execution times (in seconds) for runs with 1, 6, and 19 CPU cores of different programs using the Intel Fortran (`ifort`) and GNU Fortran (`gfortran`) compilers on a workstation with two Intel Xeon E5-2650 v3 CPUs, with a total of 20 CPU cores, and the obtained speedups for 19 CPU cores.

# of cores	1	1	6	6	19	19	19	19
Fortran	Intel time	GNU time	Intel time	GNU time	Intel time	GNU time	Intel speedup	GNU speedup
imag1d	52	60	22	22	20	22	2.6	2.7
imagcir	22	30	14	15	14	15	1.6	2.0
imagsph	24	30	12	15	12	14	2.4	2.1
real1d	205	345	76	108	62	86	3.3	4.0
realcir	145	220	55	73	48	59	3.0	3.7
realsph	155	250	57	76	46	61	3.4	2.7
imag2d	255	415	52	84	27	40	9.4	10.4
imagaxi	260	435	62	105	30	55	8.7	7.9
real2d	325	525	74	107	32	50	10.1	10.5
realaxi	160	265	35	49	16	24	10.0	11.0
imag3d	2080	2630	370	550	200	250	10.4	10.5
real3d	19500	26000	3650	5600	1410	2250	13.8	11.6



The speedup tends to increase for programs in higher dimensions, as they become more complex and have to process more data. This is why the speedups of the supplied 2d and 3d programs are larger than those of 1d programs. Also, for a single program the speedup increases with the size of the spatial grid, i.e., with the number of spatial discretization points, since this increases the amount of calculations performed by the program. To demonstrate this, we tested the supplied real2d-th program and varied the number of spatial discretization points  $NX=NY$  from 20 to 1000. The measured speedup obtained when running this program on 19 CPU cores as a function of the number of discretization points is shown in Fig. 2. The speedup first increases rapidly with the number of discretization points and eventually saturates.

*Additional comments:* Example inputs provided with the programs take less than 30 minutes to run on a workstation with two Intel Xeon E5-2650 v3 processors (2 QPI links, 10 CPU cores, 25 MB cache, 2.3 GHz).  
© 2017 Elsevier B.V. All rights reserved.

## Acknowledgments

V.L., D.V., and A.B. acknowledge support by the Ministry of Education, Science, and Technological Development of the Republic of Serbia under projects ON171017 and III43007. P.M. acknowledges support by the Science and Engineering Research Board, Department of Science and Technology, Government of India under project no. EMR/2014/000644. S.K.A. acknowledges support by the CNPq of Brazil under project 303280/2014-0, and by the FAPESP of Brazil under project 2012/00451-0. Numerical tests were partially carried out on the PARADOX supercomputing facility at the Scientific Computing Laboratory of the Institute of Physics Belgrade.

## References

- [1] P. Muruganandam, S.K. Adhikari, *Comput. Phys. Comm.* 180 (2009) 1888.
- [2] L.E. Young-S., D. Vudragović, P. Muruganandam, S.K. Adhikari, A. Balaž, *Comput. Phys. Comm.* 204 (2016) 209.
- [3] H. Fabrelli, et al., *J. Opt.* 19 (2017) 075501;  
S.K. Adhikari, *Laser Phys. Lett.* 14 (2017) 065402;  
A.N. Malmi-Kakkada, O.T. Valls, C. Dasgupta, *Phys. Rev. B* 95 (2017) 134512;  
P.S. Vinayagam, R. Radha, S. Bhuvaneshwari, R. Ravisankar, P. Muruganandam, *Commun. Nonlinear Sci. Numer. Simul.* 50 (2017) 68;  
O. Voronich, et al., *Comput. Phys. Comm.* 215 (2017) 246;  
V. Veljić, A. Balaž, A. Pelster, *Phys. Rev. A* 95 (2017) 053635;  
A.M. Martin, et al., *J. Phys.: Condens. Matter.* 29 (2017) 103004;  
R.R. Sakhel, A.R. Sakhel, *J. Phys. B: At. Mol. Opt. Phys.* 50 (2017) 105301;  
E. Chiquillo, *J. Phys. A* 50 (2017) 105001;  
G.A. Sekh, *Phys. Lett. A* 381 (2017) 852;  
W. Wen, B. Chen, X. Zhang, *J. Phys. B: At. Mol. Opt. Phys.* 50 (2017) 035301;  
S.K. Adhikari, *Phys. Rev. A* 95 (2017) 023606;  
S. Gautam, S.K. Adhikari, *Phys. Rev. A* 95 (2017) 013608;  
S.K. Adhikari, *Laser Phys. Lett.* 14 (2017) 025501;  
D. Mihalache, *Rom. Rep. Phys.* 69 (2017) 403;  
X.-F. Zhang, et al., *Ann. Phys.* 375 (2016) 368;  
G. Vergez, et al., *Comput. Phys. Comm.* 209 (2016) 144;  
S. Bhuvaneshwari, et al., *J. Phys. B: At. Mol. Opt. Phys.* 49 (2016) 245301;  
C.-Y. Lai, C.-C. Chien, *Sci. Rep.* 6 (2016) 37256;  
C.-Y. Lai, C.-C. Chien, *Phys. Rev. Appl.* 5 (2016) 034001;  
H. Gargoubi, et al., *Phys. Rev. E* 94 (2016) 043310;  
S.K. Adhikari, *Phys. Rev. E* 94 (2016) 032217;  
I. Vasić, A. Balaž, *Phys. Rev. A* 94 (2016) 033627;  
R.R. Sakhel, A.R. Sakhel, *J. Low Temp. Phys.* 184 (2016) 1092;  
J.B. Sudharsan, et al., *J. Phys. B: At. Mol. Opt. Phys.* 49 (2016) 165303;  
A. Li, et al., *Phys. Rev. A* 94 (2016) 023626;  
R.K. Kumar, et al., *J. Phys. B: At. Mol. Opt. Phys.* 49 (2016) 155301;  
K. Nakamura, et al., *J. Phys. A* 49 (2016) 315102;  
S.K. Adhikari, *Laser Phys. Lett.* 13 (2016) 085501;  
A. Paredes, H. Michinel, *Phys. Dark Universe* 12 (2016) 50;  
W. Bao, Q. Tang, Y. Zhang, *Commun. Comput. Phys.* 19 (2016) 1141;  
A.R. Sakhel, *Physica B* 493 (2016) 72;  
J. Akram, B. Girodias, A. Pelster, *J. Phys. B: At. Mol. Opt. Phys.* 49 (2016) 075302;  
J. Akram, A. Pelster, *Phys. Rev. A* 93 (2016) 033610;  
T. Khellil, A. Balaž, A. Pelster, *New J. Phys.* 18 (2016) 063003;  
D. Hocker, J. Yan, H. Rabitz, *Phys. Rev. A* 93 (2016) 053612;  
J. Akram, A. Pelster, *Phys. Rev. A* 93 (2016) 023606;  
S. Subramanian, *Eur. Phys. J. D* 70 (2016) 109;  
Z. Marojevic, E. Goeklue, C. Laemmerzahl, *Comput. Phys. Comm.* 202 (2016) 216;  
R.R. Sakhel, et al., *Eur. Phys. J. D* 70 (2016) 66;  
K. Manikandan, et al., *Phys. Rev. E* 93 (2016) 032212;  
S.K. Adhikari, *Laser Phys. Lett.* 13 (2016) 035502;  
S. Gautam, S.K. Adhikari, *Phys. Rev. A* 93 (2016) 013630;  
T. Mithun, K. Porsezian, B. Dey, *Phys. Rev. A* 93 (2016) 013620;  
D.-S. Wang, Y. Xue, Z. Zhang, *Romanian J. Phys.* 61 (2016) 827;  
S. Sabari, K. Porsezian, P. Muruganandam, *Romanian Rep. Phys.* 68 (2016) 990;  
J. Akram, A. Pelster, *Laser Phys.* 26 (2016) 065501;  
R.R. Sakhel, A.R. Sakhel, H.B. Ghassib, *Physica B* 478 (2015) 68;  
J.B. Sudharsan, et al., *Phys. Rev. A* 92 (2015) 053601.
- [4] D. Vudragović, I. Vidanović, A. Balaž, P. Muruganandam, S.K. Adhikari, *Comput. Phys. Comm.* 183 (2012) 2021.
- [5] R. Kishor Kumar, L.E. Young-S., A. Vudragović, P. Balaž, D. Muruganandam, S.K. Adhikari, *Comput. Phys. Comm.* 195 (2015) 117.
- [6] V. Lončar, A. Balaž, A. Bogojević, S. Škrbić, P. Muruganandam, S.K. Adhikari, *Comput. Phys. Comm.* 200 (2016) 406;  
V. Lončar, L.E. Young-S., S. Škrbić, P. Muruganandam, S.K. Adhikari, A. Balaž, *Comput. Phys. Comm.* 209 (2016) 190;  
B. Satarić, V. Slavnić, A. Belić, A. Balaž, P. Muruganandam, S.K. Adhikari, *Comput. Phys. Comm.* 200 (2016) 411.

# Parallelization of Minimum Spanning Tree Algorithms Using Distributed Memory Architectures

Vladimir Lončar, Srdjan Škrbić and Antun Balaž

**Abstract** Finding a minimum spanning tree of a graph is a well known problem in graph theory with many practical applications. We study serial variants of Prim's and Kruskal's algorithm and present their parallelization targeting message passing parallel machine with distributed memory. We consider large graphs that can not fit into memory of one process. Experimental results show that Prim's algorithm is a good choice for dense graphs while Kruskal's algorithm is better for sparse ones. Poor scalability of Prim's algorithm comes from its high communication cost while Kruskal's algorithm showed much better scaling to larger number of processes.

**Keywords** Distributed memory · Kruskal · MPI · MST · Paralellization · Prim

## 1 Introduction

A minimum spanning tree (MST) of a weighted graph  $G = (V, E)$  is a subset of  $E$  that forms a spanning tree of  $G$  with minimum total weight. MST problem has many applications in computer and communication network design, as well as indirect applications in fields such as computer vision and cluster analysis [12].

---

V. Lončar · S. Škrbić (✉)

Faculty of Science, University of Novi Sad, Trg Dositeja Obradovica 4, Novi Sad, Serbia  
e-mail: srdjan.skrbic@dmi.uns.ac.rs

V. Lončar

e-mail: vladimir.loncar@dmi.uns.ac.rs

A. Balaž

Scientific Computing Laboratory, Institute of Physics Belgrade, University of Belgrade,  
Pregrevica 118, Belgrade, Serbia  
e-mail: antun.balaz@scl.rs



In this paper we implement two parallel algorithms for finding MST of a graph, based on classical algorithms of Prim [23] and Kruskal [18], building upon our previous work in [19]. Algorithms target message passing parallel machine with distributed memory. Primary characteristic of this architecture is that the cost of inter-process communication is high in comparison to cost of computation. Our goal was to develop algorithms which minimize communication, and to measure the impact of communication on the performance of algorithms. Our primary interest were graphs which have significantly larger number of vertices than processors involved in computation. Since graphs of this size cannot fit into the memory of a single process, we use a partitioning scheme to divide the input graph among processes. We consider both sparse and dense graphs.

First algorithm is a parallelization of Prim's serial algorithm. Each process is assigned a subset of vertices and in each step of computation, every process finds a candidate minimum-weight edge connecting one of its vertices to MST. The root process collects those candidates and selects one with minimum weight which it adds to MST and broadcasts result to other processes. This step is repeated until every vertex is in MST.

Second algorithm is based on Kruskal's approach. Processes get a subset of  $G$  in the same way as in first algorithm, and then find local minimum spanning tree (or forest). Next, processes merge their MST edges until only one process remains, which holds edges that form MST of  $G$ .

Implementations of these algorithms are done using C programming language and MPI (Message Passing Interface) and tested on a parallel cluster PARADOX using up to 256 cores and 256 GB of distributed memory.

Section 2 contains references to the most important related papers. In Sect. 3 we continue with the description and analysis of algorithms—both serial and parallel versions, and their implementation. In the last section we describe experimental results, analyze them and draw our conclusions.

## 2 Related Work

Algorithms for MST problem have mostly been based on one of three approaches, that of Boruvka [3], Prim [23] and Kruskal [18], however, a number of new algorithms has been developed. Gallager et al. [10] presented an algorithm where processor exists at each node of the graph (thus  $n = p$ ), useful in computer network design. Katriel and Sanders designed an algorithm exploiting cycle property of a graph targeting dense graph, [17], while Ahrabian and Nowzari-Dalini's algorithm relies on depth first search of the graph [1].

Due to its parallel nature, Boruvka's algorithm (also known as Sollin's algorithm) has been the subject to most research related to parallel MST algorithms. Examples of algorithms based on Boruvka's approach include Chung and Condon [4], Wang and Gu [14] and Dehne and Götz [7].

Parallelization of Prim's algorithm has been presented by Deo and Yoo [8]. Their algorithm targets shared memory computers. Improved version of Prim's algorithm has been presented by Gonina and Kale [11]. Their algorithm adds multiple vertices per iteration, thus achieving significant speedups. Another approach targeting shared memory computers presented by Setia et al. [24] uses the cut property of a graph to grow multiple trees in parallel. Hybrid approach, combining both Boruvka's and Prim's approaches has been developed by Bader and Cong [2].

Examples of parallel implementation of Kruskal's algorithm can be found in work of Jin and Baker [16], and Osipov et al. [21]. Osipov et al. proposes a modification to Kruskal's algorithm to avoid edges which certainly are not in a graph. Their algorithm runs in near linear time if graph is not too sparse.

Bulk of the research into parallel MST algorithms has targeted shared memory computers like PRAM, i.e. computers where entire graph can fit into memory. Our algorithms target distributed memory computers and use partitioning scheme to divide the input graph evenly among processors. Because no process contains info about partition of other processes, we designed our algorithms to use predictable communication patterns, and not depend on the properties of input graph.

### 3 The Algorithms

Let us assume that graph  $G = (V, E)$ , with vertex set  $V$  and edge set  $E$  is connected and undirected. Without loss of generality, it can be assumed that each weight is distinct, thus  $G$  is guaranteed to have only one MST. This assumption simplifies implementation, otherwise a numbering scheme can be applied to edges with same weight, at the cost of additional implementation complexity.


Let  $n$  be the number of vertices,  $m$  the number of edges ( $|V| = n$ ,  $|E| = m$ ), and  $p$  the number of processes involved in computation of MST. Let  $w(v, u)$  denote weight of edge connecting vertices  $v$  and  $u$ . Input graph  $G$  is represented as  $n \times n$  adjacency matrix  $A = (a_{i,j})$  defined as:

$$a_{i,j} = \begin{cases} w(v_i, v_j) & \text{if } (v_i, v_j) \in E \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

#### 3.1 Prim's Algorithm

Prim's algorithm starts from an arbitrary vertex and then grows the MST by choosing a new vertex and adding it to MST in each iteration. Vertex with an edge with lightest weight incident on the vertices already in MST is added in every iteration. The algorithm continues until all the vertices have been added to the MST. This algorithm requires  $O(n^2)$  time. Implementations of Prim's algorithm commonly use auxiliary array  $d$  of length  $n$  to store distances (weight) from each

**Fig. 1** Partitioning of adjacency matrix among  $p$  processes



vertex to MST. In every iteration a lightest weight edge in  $d$  is added to MST and  $d$  is updated to reflect changes.

Parallelizing the main loop of Prim's algorithm is difficult [13], since after adding a vertex to MST lightest edges incident on MST change. Only two steps can be parallelized: selection of the minimum-weight edge connecting a vertex not in MST to a vertex in MST, and updating array  $d$  after a vertex is added to MST. Thus, parallelization can be achieved in the following way:

1. Partition the input set  $V$  into  $p$  subsets, such that each subset contains  $n/p$  consecutive vertices and their edges, and assign each process a different subset. Each process also contains part of array  $d$  for vertices in its partition. Let  $V_i$  be the subset assigned to process  $p_i$ , and  $d_i$  part of array  $d$  which  $p_i$  maintains. Partitioning of adjacency matrix is illustrated in Fig. 1.
2. Every process  $p_i$  finds minimum-weight edge  $e_i$  (*candidate*) connecting MST with a vertex in  $V_i$ .
3. Every process  $p_i$  sends its  $e_i$  edge to the root process using all-to-one reduction.
4. From the received edges, the root process selects one with a minimum weight (called global minimum-weight edge  $e_{min}$ ), adds it to MST and broadcasts it to all other processes.
5. Processes mark vertices connected by  $e_{min}$  as belonging to MST and update their part of array  $d$ .
6. Repeat steps 2–5 until every vertex is in MST.

Finding a minimum-weight edge and updating of  $d_i$  during each iteration costs  $O(n/p)$ . Each step also adds a communication cost of all-to-one reduction and all-to-one broadcast. These operations complete in  $O(\log p)$ . Combined, cost of one iteration is  $O(n/p + \log p)$ . Since there are  $n$  iterations, total parallel time this algorithm runs in is:

$$T_p = O\left(\frac{n^2}{p}\right) + O(n \log p) \quad (2)$$

Prim's algorithm is better suited for dense graphs and works best for complete graphs. This also applies to its parallel formulation presented here. Ineffectiveness of the algorithm on sparse graphs stems from the fact that Prim's algorithm runs in  $O(n^2)$ , regardless of the number of edges. A well-known modification [5] of Prim's algorithm is to use binary heap data structure and adjacency list representation of a graph to reduce the run time to  $O(m \log n)$ . Furthermore, using Fibonacci heap asymptotic running time of Prim's algorithm can be improved to  $O(m + n \log n)$ . Since we use adjacency matrix representation, investigating alternative approaches for Prim's algorithm was out of the scope of this paper.

### 3.2 Kruskal's Algorithm

Unlike Prim's algorithm which grows a single tree, Kruskal's algorithm grows multiple trees in parallel. Algorithm first creates a forest  $F$ , where each vertex in the graph is a separate tree. Next step is to sort all edges in  $E$  based on their weight. Algorithm then chooses minimum-weight edge  $e_{min}$  (i.e. first edge in sorted set). If  $e_{min}$  connects two different trees in  $F$ , it is added to the forest and two trees are combined into a single tree, otherwise  $e_{min}$  is discarded. Algorithm loops until either all edges have been selected, or  $F$  contains only one tree, which is the MST of  $G$ . This algorithm is commonly implemented using Union-Find algorithm [22]. *Find* operation is used to determine which tree a particular vertex is in, while *Union* operation is used to merge two trees. Kruskal's algorithm runs in  $O(m \log n)$  time, but can be made even more efficient by using more sophisticated Union-Find data structure, which uses *union by rank* and *path compression* [9]. If the edges are already sorted, using improved Union-Find data structure Kruskal's algorithm runs in  $O(m\alpha(n))$ , where  $\alpha(n)$  is the inverse of the Ackerman function.

Our parallel implementation of Kruskal's algorithm uses the same partitioning scheme of adjacency matrix as in Prim's approach and is thus bounded by  $O(n^2)$  time to find all edges in matrix. Having that in mind, our parallel algorithm proceeds through the following steps:

1. Every process  $p_i$  first sorts edges contained in its partition  $V_i$ .
2. Every process  $p_i$  finds a local minimum spanning tree (or forest, MSF)  $F_i$  using edges in its partition  $V_i$  applying the Kruskal's algorithm.
3. Processes merge their local MST's (or MSF's). Merging is performed in the following manner. Let  $a$  and  $b$  denote two processes which are to merge their local trees (or forests), and let  $F_a$  and  $F_b$  denote their respective set of local MST edges. Process  $a$  sends set  $F_a$  to  $b$ , which forms a new local MST (or MSF) from  $F_a \cup F_b$ . After merging, process  $a$  is no longer involved in computation and can terminate.
4. Merging continues until only one process remains. Its MST is the end result.

Creating a new local MSF during merge step can be performed in a number of different ways. Our approach is to perform Kruskal's algorithm again on  $F_a \cup F_b$ .

Computing the local MST takes  $O(n^2/p)$ . There is a total of  $\log p$  merging stages, each costing  $O(n^2 \log p)$ . During one merge step one process transmits maximum of  $O(n)$  edges for a total parallel time of:

$$T_p = O(n^2/p) + O(n^2 \log p) \quad (3)$$

Based on speedup and efficiency metrics, it can be shown that this parallel formulation is efficient for  $p = O(n/\log n)$ , same as the first algorithm.

### 3.3 Implementation

Described algorithms were implemented using ANSI C and Message Passing Interface (MPI). Fixed communication patterns in parallel formulation of the algorithms map directly to MPI operations. Complete source code can be found in [25].

## 4 Experimental Results

Implementations of algorithms were tested on a cluster of up to 32 computing nodes. Each computer in the cluster had two Intel Xeon E5345 2.33 GHz quad-core CPUs and 8 GB of memory, with Scientific Linux 6 operating system installed. We used OpenMPI v1.6 implementation of the MPI standard. The cluster nodes are connected to the network with a throughput of 1 Gbit/s. Both implementations were compiled using GCC 4.4 compiler. This cluster has enabled testing algorithms with up to 256 processes as shown in Table 1.

We tested graphs with densities of 1, 5, 10, 15 and 20 % with number of vertices ranging from 10,000 to 100,000, and number of edges from 500,000 to 1,000,000,000. Distribution of edges in graphs was uniformly random, and all edge weights were unique. Due to the high memory requirements of large graphs, not every input graph could be partitioned in a small number of cluster nodes, as can be seen in Table 1.

### 4.1 Results

Due to the large amount of obtained test results, we only present the most important ones here. Complete set of results can be found in [25].

In the Table 2 we show the behavior of algorithms with increasing number of processes on input graph of 50,000 vertices and density of 10 %.

Results show poor scalability of Prim's algorithm, due to its high communication cost. Otherwise, computation phase of Prim's algorithm is faster than that of

**Table 1** Testing parameters

Processes	Nodes	Processes per node	No. of vertices (k)
4	4	1	10–50
8	8	1	10–60
16	16	1	10–80
32	32	1	10–100
64	32	2	10–100
128	32	4	10–100
256	32	8	10–100

**Table 2** CPU time (in seconds) for algorithms with increasing number of processes

	4	8	16	32	64	128	256
Kruskal	38.468	19.94	10.608	5.342	2.958	1.796	1.382
Prim	16.703	15.479	25.201	30.382	30.824	32.661	39.737

**Table 3** CPU time (in seconds) for algorithms with increasing density

	1 %	5 %	10 %	15 %	20 %
Kruskal	0.607	2.603	5.342	8.164	10.663
Prim	30.189	30.007	30.382	30.518	30.589

Kruskal's. Due to the usage of adjacency matrix graph representation, Prim's algorithm performs almost the same regardless of the density of the input graph. This can be seen from the results of input graph with 50,000 vertices and 32 processes with varying density shown at Table 3.

On the other hand, Kruskal's algorithm shows degradation of performance with increasing density. Results of Kruskal's algorithm show that majority of local computation time is spent sorting the edges of input graph, which grows with larger density. Increasing the number of processes makes local partitions smaller and faster to process, thus allowing this algorithm to achieve good scalability. If the edges of input graph were already sorted, Kruskal's algorithm would be significantly faster than other MST algorithms.

## 4.2 Impact of Communication Overhead

Cost of communication is much greater than the cost of computation, so it is important to analyse the time spent in communication routines. During tests we measured the time spent waiting for the completion of the communication operations. In case of Prim's algorithm, we measured the time that the root process spends waiting for the completion of MPI\_Reduce and MPI\_Bcast operations. Communication in Kruskal's algorithm is measured as total time spent waiting for messages received over MPI\_Recv operation in the last active process (which will

**Table 4** Communication versus computation time (in seconds)

Processes	4	8	16	32	64	128	256
Prim's algorithm							
Total	16.703	15.479	25.201	30.382	30.824	32.661	39.737
Communication	8.188	11.183	23.009	29.248	30.237	32.322	39.467
Kruskal's algorithm							
Total	38.468	19.94	10.608	5.342	2.958	1.796	1.382
Communication	0.171	0.356	0.371	0.288	0.317	0.253	0.256

contain the MST after last iteration of the merge operation). This gives us a good insight into the duration of communication routines because the last active process will have to wait the most.

The Table 4 shows communication times of processing input graph of 50,000 vertices with 10 % density.

When comparing communication time with a total computation time it can be noted that the Prim's algorithm spends most of time in communication operations, and by increasing number of processes almost all the running time of the algorithm is spent on communication operations. A bottleneck in Prim's algorithm is the cost of MPI\_Reduce and MPI\_Bcast communication operations. These operations require communication between all processes, and are much more expensive than local computation within each process, because all processes must wait until the operation is completed, or until the data are transmitted over the network. This prevents Prim's algorithm from achieving substantial speedup of running time with increasing number of processes. Therefore, this algorithm is most efficient on the fewest number of processes that the partitioned input graph can fit.


On the other hand Kruskal algorithm spends much less time in communication operations, but instead spends most of the time in local computation. These differences are illustrated in Figs. 2 and 3. The diagrams show that communication in Prim's algorithm rises sharply with increasing number of processes, while execution time slowly reduces. In Kruskal's algorithm, the situation is reversed.

### 4.3 Analysis of Results


The experimental results confirmed some of the assumptions made during the development and analysis of algorithms, but also made a couple of unexpected results. Results of these experiments gave us directions for further improvement of the described algorithms.

Prim's algorithm has shown excellent performance in computational part of the algorithm, but a surprisingly high cost of communication operations spoils its final score. Finding candidate edges for inclusion in MST can be further improved by using techniques described in [5], but it will not significantly improve the total





**Fig. 2** Communication in Kruskal's algorithm



**Fig. 3** Communication in Prim's algorithm

time of the algorithm, as communication routines will remain the same. Unfortunately, the communication can not be further improved by changing the algorithm. The only way to reduce the cost of communication is to use a cluster that has a better quality network, or to rely on the semantics of the implementation of the MPI operation `MPI_Allreduce`.

Kruskal's algorithm has shown good performance, especially for sparse graphs, while the performance degrades with increasing density. It is important to note that many real-world graphs have density much smaller than 1 % (for example, graph of roads as edges and junctions as vertices has a density much smaller than 1 %). Also, this algorithm showed much better scaling to larger number of processes than Prim's algorithm. Cost of communication in Kruskal's algorithm is much smaller than in Prim's algorithm, but the local computation is slower. This can be improved by using more efficient Union-Find algorithms [9], or by improving merging of local trees between processes. Kruskal's algorithm does not use a lot of



slow messages like Prim's algorithm, but can send very large messages depending on the number of processes and the size of the graph. This can be improved by introducing techniques for compressing messages, or changing the structure of the message.

## 5 Alternate Parallelization Approaches

In this section we will give a brief overview of two other parallelization approaches we considered using for implementation of these algorithms. One approach would be using graphics processing unit (GPU) technologies like Nvidia CUDA or OpenCL. Another would be using shared-memory parallelization API like OpenMP to utilize multi-core processors on cluster nodes. We will go over advantages and disadvantages of both approaches.

With the introduction of CUDA and OpenCL programming models, using GPU for general-purpose computing (GPGPU) has become a powerful alternative to traditional CPU programming models. Nowadays GPUs can be found in most high-ranking supercomputers and even ordinary clusters. GPUs have their own RAM, which is separate from main RAM of a computer and was not accessible for distributed-memory technologies like MPI. This made writing multi-GPU programs more difficult, since it required expensive copy operations between GPU memory and host (CPU) memory which MPI could access. However, recent developments in MPI implementations have alleviated this problem, and newer versions of popular MPI implementations like OpenMPI and MVAPICH can access GPU memory directly. This unfortunately still doesn't make GPU the perfect platform for implementations of our algorithms. GPUs still have much smaller amount of RAM when compared to main memory (recently released models like Tesla K10 have up to 8 GB of memory [20]). This means that GPU solution could only be used on much smaller graphs. Alternatively, a different graph representation (like adjacency lists) would allow graphs with greater number of vertices, but would still be only useful for sparser graphs. Primary part of Prim's algorithm which could be accelerated by GPU is finding local (and then global) vertex with the smallest distance to the tree. This could be achieved by slightly modifying well-known parallel reduction algorithm for GPU [15]. Communication pattern between nodes would remain the same. Kruskal's algorithm is more complex to implement on GPU due to Union-Find data structure. Other important portions of Kruskal's algorithm, like sorting of input could be done using various GPU libraries.

Unlike the relatively new technology that is GPGPU, OpenMP has been successfully used to parallelize serial code since the late 90s. In some cases, OpenMP allows developers to parallelize their programs with minimal effort, using compiler directives around loops, often with good performance [6]. This technique could be used in parallelization of Prim's algorithm for finding local (and later

global) vertex with the smallest distance to the tree. Graph would be partitioned in such a way that each node in cluster receives an equal part, then each node would use all its processors and cores with OpenMP to find local minimum, and use MPI for communication between nodes. Kruskal's algorithm can be parallelized in similar way, although it would require a slightly greater effort for implementation of sorting and Union-Find data structure.

**Acknowledgements** Authors are partially supported by Ministry of Education, Science, and Technological Development of the Republic of Serbia, through projects no. ON174023: "Intelligent techniques and their integration into wide-spectrum decision support", and ON171017: "Modeling and numerical simulations of complex many-body systems", as well as European Commission through FP7 projects PRACE-2IP and PRACE-3IP.

## References

1. H. Ahrabian, A. Nowzari-Dalini, Parallel algorithms for minimum spanning tree problem. *Int. J. Comput. Math.* **79**(4), 441–448 (2002)
2. D.A. Bader, G. Cong, Fast shared-memory algorithms for computing the minimum spanning forest of sparse graphs. *J. Parallel Distrib. Comput.* **66**(11), 1366–1378 (2006)
3. O. Boruvka, O Jistém Problému Minimálnm (about a certain minimal problem) (in Czech, German summary). *Práce Mor. Prrodoved. Spol. v Brne III*, vol. 3 (1926)
4. S. Chung, A. Condon, Parallel implementation of borvka's minimum spanning tree algorithm. in *Proceedings of the 10th International Parallel Processing Symposium, IPPS '96* (IEEE Computer Society, Washington, DC, 1996), pp. 302–308
5. T.H. Cormen, C. Stein, R.L. Rivest, C.E. Leiserson, *Introduction to Algorithms*, 2nd edn. (McGraw-Hill Higher Education, Boston, 2001)
6. M. Curtis-Maury, X. Ding, C.D. Antonopoulos, D.S. Nikolopoulos, *An Evaluation of Openmp on Current and Emerging Multithreaded/Multicore Processors*, ed. by M.S. Mueller, B.M. Chapman, B.R. Supinski, A.D. Malony, M. Voss. *OpenMP Shared Memory Parallel Programming*, vol 4315 (*Lecture Notes in Computer Science*, Springer Berlin Heidelberg, 2008), pp. 133–144
7. F. Dehne, S. Gtz, Practical Parallel Algorithms for Minimum Spanning Trees, in *Workshop on Advances in Parallel and Distributed Systems* (1998), pp. 366–371
8. N. Deo, Y.B. Yoo, Parallel algorithms for the minimum spanning tree problem, in *Proceedings of the International Conference on Parallel Processing* (1981), pp. 188–189
9. Z. Galil, G.F. Italiano, Data structures and algorithms for disjoint set union problems. *ACM Comput. Surv.* **23**(3), 319–344 (1991)
10. R.G. Gallager, P.A. Humblet, P.M. Spira, A distributed algorithm for minimum-weight spanning trees. *ACM Trans. Program. Lang. Syst.* **5**(1), 66–77 (1983)
11. E. Gonina, L.V. Kale, Parallel prim's algorithm on dense graphs with a novel extension, in *PPL Technical Report*, Oct 2007
12. R.L. Graham, P. Hell, On the history of the minimum spanning tree problem. *IEEE Ann. Hist. Comput.* **7**(1), 43–57 (1985)
13. A. Grama, G. Karypis, V. Kumar, A. Gupta, *Introduction to Parallel Computing*, 2nd edn. (Addison Wesley, Reading, 2003)
14. W. Guang-rong, G. Nai-jie, An efficient parallel minimum spanning tree algorithm on message passing parallel machine. *J. Softw.* **11**(7), 889–898 (2000)
15. M. Harris, Optimizing parallel reduction in CUDA. *CUDA tips and tricks*

16. M. Jin, J.W. Baker, Two graph algorithms on an associative computing model, in *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications, PDPTA 2007*, vol 1, Las Vegas, Nevada, 25–28 June 2007, pp. 271–277
17. I. Katriel, P. Sanders, J.L. Trff, J.L. Tra, A practical minimum spanning tree algorithm using the cycle property, in *11th European Symposium on Algorithms (ESA)*, vol. 2832 in *LNCS* (Springer, New York, 2003), pp. 679–690
18. J.B. Kruskal, On the shortest spanning subtree of a graph and the traveling salesman problem. *Proc. Am. Math. Soc.* **7**(1), 48–50 (1956)
19. V. Lončar, S. Škrbić, A. Balaž, Distributed memory parallel algorithms for minimum spanning trees, in *Lecture Notes in Engineering and Computer Science: Proceedings of the World Congress on Engineering 2013, WCE 2013*, London, 3–5 July 2013, pp. 1271–1275
20. Nvidia, Nvidia tesla GPU accelerators. Nvidia Tesla Product Datasheet (2012)
21. V. Osipov, P. Sanders, J. Singler, The filter-kruskal minimum spanning tree algorithm, in *ALENEX'09* (2009), pp. 52–61
22. D.-Z. Pan, Z.-B. Liu, X.-F. Ding, Q. Zheng, The application of union-find sets in kruskal algorithm, in *Proceedings of the 2009 International Conference on Artificial Intelligence and Computational Intelligence (AICI '09)*, vol 2 (IEEE Computer Society, Washington, DC, 2009), pp. 159–162
23. R.C. Prim, Shortest connection networks and some generalizations. *Bell Syst. Technol. J.* **36**, 1389–1401 (1957)
24. R. Setia, A. Nedunchezian, S. Balachandran, A new parallel algorithm for minimum spanning tree problem, in *Proceedings of the International Conference on High Performance Computing (HiPC)* (2009), pp. 1–5
25. S. Škrbić, Scientific Computing Seminar (2013)

# Parallel implementation of minimum spanning tree algorithms using MPI

Vladimir Lončar\* and Srdjan Škrbic\*\*

Faculty of Science, Department for Mathematics and Informatics, University of Novi Sad, Serbia

\* vlada.loncar@gmail.com

\*\* shkrba@uns.ac.rs

**Abstract**—In this paper we study parallel algorithms for finding minimum spanning tree of a graph. We present two algorithms, based on sequential algorithms of Prim and Kruskal, targeting message passing parallel machine with distributed memory. First algorithm runs in  $O(n^2/p + n \log p)$  and second algorithm runs in  $O(n^2/p + n^2 \log p)$ .

**Index Terms**—Minimum spanning tree, parallel algorithms, message passing, distributed memory computer.

## I. INTRODUCTION

A minimum spanning tree (MST) of a graph  $G = (V, E)$  is a subset of  $E$  that forms a spanning tree of  $G$  with minimum total weight. MST problem has many applications in computer and communication network design, as well as indirect applications in fields such as computer vision and cluster analysis [1].

In this paper we implement two parallel algorithms for finding MST of a graph, based on classical algorithms of Prim and Kruskal. Algorithms target message passing parallel machine with distributed memory. Primary characteristic of this architecture is that the cost of inter-process communication is high in comparison to cost of computation. Our goal was to develop algorithms which minimize communication, and to measure the impact of communication on the performance of algorithms. Our primary interest were graphs which have significantly larger number of vertices than processors involved in computation. Since graphs of this size cannot fit into a memory of single process, we use simple partitioning scheme to divide the input graph among processes. We considered both sparse and dense graphs.

First algorithm is a parallelization of Prim's sequential algorithm. Each process is assigned a subset of vertices and in each step of computation, every process finds a *candidate* minimum-weight edge connecting one of its vertices to MST. Leader process collects those candidates and selects one with minimum weight which it adds to MST, and broadcasts result to other processes. This step is repeated until every vertex is in MST.

Second algorithm is based on Kruskal's approach. Processes get a subset of  $G$  in the same way as in first algorithm, and then find local minimum spanning tree (or forest). Next, processes merge their MST edges until only one process remains, which holds edges that form MST of  $G$ .

Algorithms we present are both easy to understand and implement, and since they use fixed communication patterns, their performance can easily be predicted.

## II. RELATED WORK

Algorithms for MST problem have mostly been based on one of three approaches, that of Boruvka [2], Prim [3] and Kruskal [4], however, a number of new algorithms has been developed. Gallager et al. presented an algorithm where processor exists at each node of the graph (thus  $n = p$ ), useful in computer network design [5]. Katriel and Sanders designed an algorithm exploiting cycle property of a graph targeting dense graph, [6], while Ahrabian and Nowzari-Dalini's algorithm relies on depth first search of the graph [7].

Due to its parallel nature, Boruvka's algorithm (also known as Sollin's algorithm) has seen the most research. Examples of algorithms based on Boruvka's approach include Chung and Condon [8], Wang and Gu [9] and Dehne and Götz [10].

Parallelization of Prim's algorithm has been presented by Deo and Yoo [11]. Their algorithm targets shared-memory computers. Improved version of Prim's algorithm has been presented by Gonina and Kale [12]. Their algorithm adds multiple vertices per iteration, thus achieving significant speedups. Another approach targeting shared-memory computers presented by Setia et al. [13] uses the cut property of a graph to grow multiple trees in parallel. Hybrid approach, combining both Boruvka's and Prim's approaches has been developed by Bader and Cong [14].

Examples of parallel implementation of Kruskal's algorithm can be found in work of Jin and Baker [15], and Osipov et al [16]. Osipov et al. proposes a modification to Kruskal's algorithm to avoid edges which certainly are not in a graph. Their algorithm runs in near linear time if graph is not too sparse.

Bulk of the research into parallel MST algorithms has targeted shared-memory computers like PRAM, i.e. computers where entire graph can fit into memory. Our algorithms target distributed-memory computers and use partitioning scheme to divide the input graph evenly among processors. Because no process contains info about partition of other processes, we designed our algorithms to use predictable communication patterns, and not depend on the properties of input graph.

### III. THE ALGORITHMS

In the remainder of this paper, we will assume that graph  $G = (V, E)$  is connected and undirected. Without loss of generality, it can be assumed that each weight is distinct, thus  $G$  is guaranteed to have only one MST. This assumption simplifies implementation, otherwise a numbering scheme can be applied to edges with same weight, at the cost of additional implementation complexity.

Let  $n$  be the number of vertices,  $m$  the number of edges ( $|V| = n$ ,  $|E| = m$ ), and  $p$  the number of processes involved in computation of MST. Let  $w(v, u)$  denote weight of edge connecting vertices  $v$  and  $u$ . Input graph  $G$  is represented as  $n \times n$  adjacency matrix  $A = (a_{i,j})$  defined as:

$$a_{i,j} = \begin{cases} w(v_i, v_j) & \text{if } (v_i, v_j) \in E \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

#### A. Prim's Algorithm

Prim's algorithm starts from an arbitrary vertex and then grows the MST by choosing a new vertex and adding it to MST in each iteration. Vertex with an edge with lightest weight incident on the vertices already in MST is added in every iteration. The algorithm continues until all the vertices have been added to the MST. This algorithm requires  $O(n^2)$  time. Implementations of Prim's algorithm commonly use auxiliary array  $d$  of length  $n$  to store distances (weight) from each vertex to MST. In every iteration a lightest weight edge in  $d$  is added to MST and  $d$  is updated to reflect changes.

Parallelizing the main loop of Prim's algorithm is difficult [17], since after adding a vertex to MST lightest edges incident on MST change. Only two steps can be parallelized: selection of the minimum-weight edge connecting a vertex not in MST to a vertex in MST, and updating array  $d$  after a vertex is added to MST. Thus, parallelization can be achieved in the following way:

- 1) Partition the input set  $V$  into  $p$  subsets, such that each subset contains  $n/p$  consecutive vertices and their edges, and assign each process a different subset. Each process also contains part of array  $d$  for vertices in its partition. Let  $V_i$  be the subset assigned to process  $p_i$ , and  $d_i$  part of array  $d$  which  $p_i$  maintains. Partitioning of adjacency matrix is illustrated in Fig. 1.
- 2) Every process  $p_i$  finds minimum-weight edge  $e_i$  (candidate) connecting MST with a vertex in  $V_i$ .
- 3) Every process  $p_i$  sends its  $e_i$  edge to leader process using all-to-one reduction.
- 4) From the received edges, leader process selects one with a minimum weight (called global minimum-weight edge  $e_{min}$ ), adds it to MST and broadcasts it to all other processes.
- 5) Processes mark vertices connected by  $e_{min}$  as belonging to MST and update their part of array  $d$ .
- 6) Repeat steps 2-5 until every vertex is in MST.




Fig. 1. Partitioning of adjacency matrix among  $p$  processes

Finding a minimum-weight edge and updating of  $d_i$  during each iteration costs  $O(n/p)$ . Each step also adds a communication cost of all-to-one reduction and all-to-one broadcast. These operations complete in  $O(\log p)$ . Combined, cost of one iteration is  $O(n/p + \log p)$ . Since there are  $n$  iterations, total parallel time this algorithm runs in is:

$$T_p = O\left(\frac{n^2}{p}\right) + O(n \log p) \quad (2)$$

In comparison to sequential algorithm, this algorithm achieves a speedup and efficiency of:

$$S = \frac{O(n^2)}{O(n^2/p + n \log p)} \quad (3)$$

$$E = \frac{1}{1 + O((p \log p)/n)} \quad (4)$$

From equations 3 and 4 we conclude that this formulation of Prim's algorithm is efficient only for  $p = O(n/\log n)$  processes.

Prim's algorithm is better suited for dense graphs and works best for complete graphs. This also applies to its parallel formulation presented here. Ineffectiveness of the algorithm on sparse graphs stems from the fact that Prim's algorithm runs in  $O(n^2)$ , regardless of the number of edges. A well-known modification [18] of Prim's algorithm is to use binary heap data structure and adjacency list representation of a graph to reduce the run time to  $O(m \log n)$ . Furthermore, using Fibonacci heap asymptotic running time of Prim's algorithm can be improved to  $O(m + n \log n)$ . Since we use adjacency matrix representation, investigating alternative approaches for Prim's algorithm was out of the scope of this paper.

#### B. Kruskal's Algorithm

Unlike Prim's algorithm which grows a single tree, Kruskal's algorithm grows multiple trees in parallel. Algorithm first creates a forest  $F$ , where each vertex in the graph is a separate tree. Next step is to sort all edges in  $E$  based on their weight. Algorithm then loops the sorted set and chooses minimum-weight edge  $e_{min}$  (i.e. first edge in sorted set). If  $e_{min}$  connects two different

trees in  $F$ , add it to the forest and combine two trees into a single tree, otherwise discard  $e_{min}$ . Algorithm loops until either all edges have been selected, or  $F$  contains only one tree, which is the MST of  $G$ . This algorithm is commonly implemented using Union-Find algorithm [19]. *Find* operation is used to determine which tree a particular vertex is in, while *Union* operation is used to merge two trees. Kruskal's algorithm runs in  $O(m \log n)$  time, but can be made even more efficient by using more sophisticated Union-Find data structure, which uses *union by rank* and *path compression* [20]. If the edges are already sorted, using improved Union-Find data structure Kruskal's algorithm runs in  $O(m\alpha(n))$ , where  $\alpha(n)$  is the inverse of an Ackerman function.

Our parallel implementation of Kruskal's algorithm uses the same partitioning scheme of adjacency matrix as in Prim's approach and is thus bounded by  $O(n^2)$  time to find all edges in matrix. Every process first sorts edges contained in its partition. From edges in partition  $V_i$ , every process  $p_i$  finds a local minimum spanning tree (or forest, MSF)  $T_i$  using Kruskal's algorithm. At the end of this step, local MSTs are merged. Merging is performed in the following manner. Let  $a$  and  $b$  denote two processes which are to merge their local trees (or forests), and let  $A$  and  $B$  denote their respective set of local MST edges. Process  $a$  sends set  $A$  to  $b$ , which forms a new local MST (or MSF) from  $A \cup B$ . After merging, process  $a$  is no longer involved in computation and can terminate. Merging continues until only one process remains, which will contain MST of  $G$ .

Example of parallel Kruskal's algorithm is illustrated in Fig. 2. Input graph in (a) is divided among processes  $p_1$  and  $p_2$  which compute local MST based on edges incident on vertices assigned to them ((b) and (c)). Next, processes merge their local MST-s to form a MST of input graph. The dashed lines represent edges which are in local MST of a process, but are removed after merging.

Creating a new local MSF during merge step can be performed in a number of different ways. One approach is to perform Kruskal's algorithm again on  $A \cup B$ . Alternatively, a modified *depth-first search* (DFS) can be used. For every edge in  $A$ , it is first determined if it is already in the same tree of  $B$  (using *find* operation). If it is not, it is added in MSF and *union* operation is called. Merging two trees can produce a cycle, so a modified DFS is run to eliminate edge with a heaviest weight.

Computing the local MST takes  $O(n^2/p)$ . There is a total of  $\log p$  merging stages, each costing  $O(n^2 \log p)$ . During one merge step one process transmits maximum of  $O(n)$  edges for a total parallel time of:

$$T_p = O(n^2/p) + O(n^2 \log p) \quad (5)$$

Based on speedup and efficiency metrics, it can be shown that this parallel formulation is efficient for  $p = O(n/\log n)$ , same as first algorithm.

#### IV. IMPLEMENTATION

Algorithms were simple to implement using ANSI C and MPI. Simplicity is the result of fixed communication

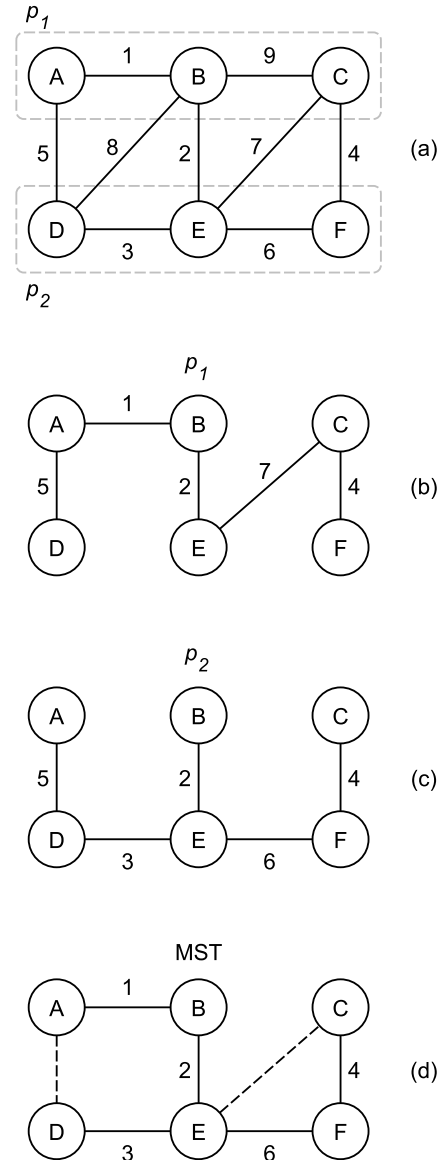


Fig. 2. Example of merge step for two processes

patterns which directly map to MPI operations. During implementation we explored alternative communication patterns in order to grow multiple trees in parallel, similar to approaches of parallelization of Boruvka's algorithm. We have found that using data-dependant communication paths results in imbalanced computation due to arbitrary communication between processes. Also, implementation of arbitrary communication can be difficult with MPI, since number of messages each process sends or receives is not known in advance for every input. Overcoming this obstacle often requires adding additional communication complexity, at the cost of overall performance.

Communication pattern in Prim's algorithm can be improved by using MPI *MPI\_Allreduce* operation instead of the standard combination of *MPI\_Reduce* and *MPI\_Bcast*. This optimization does not necessarily result in better performance, since MPI implementations can implement

*MPI\_Allreduce* operation as a simple all-to-one reduce, followed by a broadcast, without any performance improvements [21].

Main performance bottleneck of Prim's algorithm is communication overhead of all-to-one reduce operation. Reduce operation is costly in comparison to local computation, and all other processes are idle while waiting for reduce to complete. This prevents Prim's algorithm to achieve significant speedups on a larger number of processors. Therefore, Prim's algorithm is best used on a smallest number of processes on which partitioned input graph can fit.

Unlike Prim's algorithm, Kruskal's algorithm doesn't use collective communication operations during which all processes except one are idle. Performance-wise, critical part Kruskal's algorithm is merging of local MST-s. Merge part of Kruskal's algorithm is only fully efficient in case when  $p$  is a power of 2. Since merging is pairwise operation, in other cases at least one merge step will have a process without a pair. This process will be idle until a merge partner is available. In our implementation, idling can span multiple merge steps, thus causing a considerable efficiency degradation. For example, if there are 9 processes in computation, one process will be idle until the very last merge step. One approach to solving this issue would be introduction of a special 3-way merge (or in general a  $d$ -ary merge, where  $d = 2, 3, 4, \dots$ ) along with a load balancing logic to minimize or remove the effect on performance of algorithm.

## V. CONCLUSION

We presented two parallel implementations of algorithms for finding minimum spanning tree of a graph. Our algorithms are parallelizations of classical sequential algorithms of Prim and Kruskal. Parallel processes work on a subset of input graph, and communicate using fixed communication pattern. First algorithm takes  $O(n^2/p + n \log p)$  time, while second takes  $O(n^2/p + n^2 \log p)$  time.

Our analysis has identified several bottlenecks in our implementations, and further work in this area would include minimizing communication cost by reducing the number and size of messages passed, as well as improving merge step of the second algorithm. Also, further experimental work would give us information about practical limitations of our algorithms for wider array of input graphs and uncover new areas for improvement.

## ACKNOWLEDGMENT

Authors are partially supported by Ministry of Education and Science of the Republic of Serbia, through project no. III47003: "Infrastructure for technology enhanced learning in Serbia".

## REFERENCES

[1] R. L. Graham and P. Hell, "On the history of the minimum spanning tree problem," *IEEE Ann. Hist. Comput.*, vol. 7, no. 1, pp. 43–57, Jan. 1985. [Online]. Available: <http://dx.doi.org/10.1109/MAHC.1985.10011>

[2] O. Boruvka, "O Jistém Problému Minimálním (About a Certain Minimal Problem) (in Czech, German summary)," *Práce Mor. Přírodoved. Spol. v Brně III*, vol. 3, 1926.

[3] R. C. Prim, "Shortest connection networks and some generalizations," *Bell System Technology Journal*, vol. 36, pp. 1389–1401, 1957.

[4] J. B. Kruskal, "On the Shortest Spanning Subtree of a Graph and the Traveling Salesman Problem," *Proceedings of the American Mathematical Society*, vol. 7, no. 1, pp. 48–50, Feb. 1956. [Online]. Available: <http://www.jstor.org/stable/2033241>

[5] R. G. Gallager, P. A. Humblet, and P. M. Spira, "A distributed algorithm for minimum-weight spanning trees," *ACM Trans. Program. Lang. Syst.*, vol. 5, no. 1, pp. 66–77, Jan. 1983. [Online]. Available: <http://doi.acm.org/10.1145/357195.357200>

[6] I. Katriel, P. Sanders, J. L. Träff, and J. L. Tra, "A practical minimum spanning tree algorithm using the cycle property," in *In 11th European Symposium on Algorithms (ESA), number 2832 in LNCS*. Springer, 2003, pp. 679–690.

[7] H. Ahrabian and A. Nowzari-Dalini, "Parallel algorithms for minimum spanning tree problem," *International Journal of Computer Mathematics*, vol. 79, no. 4, pp. 441–448, 2002.

[8] S. Chung and A. Condon, "Parallel implementation of boruvka's minimum spanning tree algorithm," in *Proceedings of the 10th International Parallel Processing Symposium*, ser. IPPS '96. Washington, DC, USA: IEEE Computer Society, 1996, pp. 302–308. [Online]. Available: <http://dl.acm.org/citation.cfm?id=645606.661036>

[9] W. Guang-rong and G. Nai-jie, "An efficient parallel minimum spanning tree algorithm on message passing parallel machine," *Journal of Software*, vol. 11, no. 7, pp. 889–898, 2000.

[10] F. Dehne and S. Götz, "Practical parallel algorithms for minimum spanning trees," in *In Workshop on Advances in Parallel and Distributed Systems*, 1998, pp. 366–371.

[11] Y. Y. B. Deo, Narsingh, "Parallel algorithms for the minimum spanning tree problem," in *Proceedings of the International Conference on Parallel Processing*, 1981, pp. 188–189.

[12] E. Gonina and L. V. Kale, "Parallel prim's algorithm on dense graphs with a novel extension," *PPL Technical Report*, October 2007.

[13] A. N. R. Setia and S. Balachandran, "A new parallel algorithm for minimum spanning tree problem," in *Proc. International Conference on High Performance Computing (HiPC)*, 2009, pp. 1–5.

[14] D. A. Bader and G. Cong, "Fast shared-memory algorithms for computing the minimum spanning forest of sparse graphs," *J. Parallel Distrib. Comput.*, vol. 66, no. 11, pp. 1366–1378, Nov. 2006. [Online]. Available: <http://dx.doi.org/10.1016/j.jpdc.2006.06.001>

[15] M. Jin and J. W. Baker, "Two graph algorithms on an associative computing model," in *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications, PDPTA 2007, Las Vegas, Nevada, USA, June 25-28, 2007, Volume 1*, 2007, pp. 271–277.

[16] V. Osipov, P. Sanders, and J. Singler, "The filter-kruskal minimum spanning tree algorithm," in *ALENEX'09*, 2009, pp. 52–61.

[17] A. Grama, G. Karypis, V. Kumar, and A. Gupta, *Introduction to Parallel Computing (2nd Edition)*, 2nd ed. Addison Wesley, Jan. 2003. [Online]. Available: <http://www.worldcat.org/isbn/0201648652>

[18] T. H. Cormen, C. Stein, R. L. Rivest, and C. E. Leiserson, *Introduction to Algorithms*, 2nd ed. McGraw-Hill Higher Education, 2001.

[19] D.-Z. Pan, Z.-B. Liu, X.-F. Ding, and Q. Zheng, "The application of union-find sets in kruskal algorithm," in *Proceedings of the 2009 International Conference on Artificial Intelligence and Computational Intelligence - Volume 02*, ser. AICI '09. Washington, DC, USA: IEEE Computer Society, 2009, pp. 159–162. [Online]. Available: <http://dx.doi.org/10.1109/AICI.2009.155>

[20] Z. Galil and G. F. Italiano, "Data structures and algorithms for disjoint set union problems," *ACM Comput. Surv.*, vol. 23, no. 3, pp. 319–344, Sep. 1991. [Online]. Available: <http://doi.acm.org/10.1145/116873.116878>

[21] P. Patarasuk and X. Yuan, "Bandwidth optimal all-reduce algorithms for clusters of workstations," *J. Parallel Distrib. Comput.*, vol. 69, no. 2, pp. 117–124, Feb. 2009. [Online]. Available: <http://dx.doi.org/10.1016/j.jpdc.2008.09.002>

## Quench Dynamics for Trapped Dipolar Fermi Gases

V. Veljić<sup>1</sup>, A. Balaž<sup>1</sup> and A. Pelster<sup>2</sup>

<sup>1</sup>*Scientific Computing Laboratory, Institute of Physics Belgrade,  
University of Belgrade, Pregrevica 118, 11080 Belgrade, Serbia*

<sup>2</sup>*Physics Department and Research center OPTIMAS,  
Technical University of Kaiserslautern, 67663 Kaiserslautern, Germany*  
e-mail: vveljic@ipb.ac.rs

A recent time-of-flight expansion experiment for polarized fermionic erbium atoms managed to detect a Fermi surface deformation which is due to the dipolar interaction [1]. Here we perform a systematic study of quench dynamics of trapped dipolar Fermi gases at zero temperature, which are induced by a sudden change of the magnetic field, which enforces the polarization of the magnetic moments of the erbium atoms. As this modifies the equilibrium configuration, oscillations of the fermionic erbium cloud emerge around the new equilibrium, which are characteristic for the presence of the dipole-dipole interaction. In order to analyze the emergent dynamics we follow Ref. [2] and solve analytically the underlying Boltzmann-Vlasov equation within the relaxation approximation in the vicinity of the new equilibrium configuration by using a suitable rescaling of the equilibrium distribution [3]. The resulting ordinary differential equations of motion for the scaling parameters are solved numerically for experimentally relevant parameters all the way from the collisionless to the hydrodynamic regime. A comparison with a corresponding linear stability analysis reveals that the resulting quench dynamics can be understood in terms of the low-lying collective modes due to the smallness of the dipolar interaction strength. All our theoretical and numerical calculations can be tested in current experiments with ultracold dipolar fermionic atoms.

### REFERENCES

- [1] K. Aikawa et al., *Science* 345, 1484 (2014).
- [2] F. Wächtler, A. R. P. Lima, A. Pelster, arXiv: 1311.5100 (2013).
- [3] P. Pedri, D. Guery-Odelin, S. Stringari, *Phys. Rev. A* 68, 043608 (2003).

## Trapped Bose-Einstein Condensates with Strong Disorder

V. Lončar<sup>1</sup>, A. Balaž<sup>1</sup> and A. Pelster<sup>2</sup>

<sup>1</sup>*Scientific Computing Laboratory, Institute of Physics Belgrade,  
University of Belgrade, Pregrevica 118, 11080 Belgrade, Serbia*

<sup>2</sup>*Physics Department and Research center OPTIMAS,  
Technical University of Kaiserslautern, 67663 Kaiserslautern, Germany*  
e-mail: vloncar@ipb.ac.rs

We work out a non-perturbative approach towards the dirty boson problem at zero temperature that is based on a Gaussian approximation for correlation functions of the disorder problem and the condensate wave function solving the Gross-Pitaevskii problem. For harmonically trapped Bose-Einstein condensates we apply, in addition, the



semiclassical approximation and derive with this self-consistency equations between the disorder ensemble-averages of particle density and condensate density. Invoking, furthermore, the Thomas-Fermi approximation we obtain results that reproduce for weak disorder the seminal results of a Bogoliubov theory of dirty bosons [1-3], but do not yield for strong disorder a Bose-glass phase. Afterwards, we go beyond the Thomas-Fermi approximation and perform a full numerical treatment of the self-consistency equations based on the Crank-Nicolson split-step semi-implicit imaginary-time propagation [4], which yields a quantum phase transition to a Bose-glass phase for strong disorder [5].

## REFERENCES

- [1] K. Huang, H.-F. Meng, Phys. Rev. Lett. 69, 644 (1992).
- [2] G. M. Falco, A. Pelster, R. Graham, Phys. Rev. A 75, 063619 (2007).
- [3] G. M. Falco, A. Pelster, R. Graham, Phys. Rev. A 76, 013624 (2007).
- [4] D. Vudragović et al., Comput. Phys. Comm.. 183, 2021 (2012).
- [5] P. Navez, A. Pelster, R. Graham, Appl. Phys. B 86, 395 (2007).

## Faraday Waves in Dipolar Bose-Einstein Condensates

D. Vudragović and A. Balaž

*Scientific Computing Laboratory, Institute of Physics Belgrade,  
University of Belgrade, Pregrevica 118, 11080 Belgrade, Serbia*  
e-mail: dusan@ipb.ac.rs

We present the emergence of Faraday waves in cigar-shaped  $^{52}\text{Cr}$  and  $^{164}\text{Dy}$  Bose-Einstein condensates. These density waves are induced by periodic modulation of the frequency of the trapping potential. We study through extensive numerical simulations and detailed variational treatment the effects of the strong dipolar interaction on the spatial and time-period of the Faraday waves. Unlike in the case of homogeneous [1] or inhomogeneous contact interactions [2], the emergence of Faraday waves is found to further destabilize the condensate in the presence of strong dipolar interaction. The interesting effect of spatial period variation of generated density patterns is observed numerically and studied within the Gaussian variational approach.

## REFERENCES

- [1] A. Balaž, A. I. Nicolin, Phys. Rev. A 85, 023613 (2012).
- [2] A. Balaž et al., Phys. Rev. A 89, 023609 (2014).

## Q 17: Quantum Gases: Bosons I

Time: Tuesday 11:00–13:00

Location: e001

## Group Report

Q 17.1 Tue 11:00 e001

**Rosensweig instability and solitary waves in a dipolar Bose-Einstein condensate** — ●MATTHIAS WENZEL, HOLGER KADAU, MATTHIAS SCHMITT, IGOR FERRIER-BARBUT, and TILMAN PFAU — 5. Physikalisches Institut and Center for Integrated Quantum Science and Technology, Universität Stuttgart, Pfaffenwaldring 57, 70569 Stuttgart, Germany

Ferrofluids show unusual hydrodynamic effects due to the magnetic nature of their constituents. For increasing magnetization a classical ferrofluid undergoes a Rosensweig instability and creates self-organized ordered surface structures or droplet crystals.

In the experiment we observe a similar behavior in a sample of ultracold dysprosium atoms, a quantum ferrofluid. By controlling the short-range interaction with a Feshbach resonance we can induce a finite-wavelength instability due to the dipolar interaction.

Subsequently, we observe the spontaneous transition from an unstructured superfluid to an ordered arrangement of droplets by in situ imaging. These patterns are surprisingly long-lived and show hysteretic behavior. When transferring the sample to a waveguide we observe mutually interacting solitary waves. Time-of-flight measurements allow us to show the existence of an equilibrium between dipolar attraction and short-range repulsion. In addition we observe interference between droplets.

In conclusion, our system shows both superfluidity and translational symmetry breaking. This novel state of matter is thus a possible candidate for a supersolid ground state.

Q 17.2 Tue 11:30 e001

**Rosensweig instability due to three-body interaction or quantum fluctuations?** — VLADIMIR LONČAR<sup>1</sup>, DUŠAN VUDRAGOVIĆ<sup>1</sup>, ●ANTUN BALAŽ<sup>1</sup>, and AXEL PELSTER<sup>2</sup> — <sup>1</sup>Scientific Computing Laboratory, Institute of Physics Belgrade, University of Belgrade, Serbia — <sup>2</sup>Physics Department and Research Center OPTIMAS, Technical University of Kaiserslautern, Germany

In the recent experiment [1], the Rosensweig instability was observed in a <sup>164</sup>Dy Bose-Einstein condensate, which represents a quantum ferrofluid due to the large atomic magnetic dipole moments. After a sudden reduction of the scattering length, which is realized by tuning the external magnetic field far away from a Feshbach resonance, the dipolar quantum gas creates self-ordered surface structures in form of droplet crystals. As the underlying Gross-Pitaevskii equation is not able to explain the emergence of that Rosensweig instability, we extend it by both three-body interactions [2-4] and quantum fluctuations [5]. We then use extensive numerical simulations in order to study the interplay of three-body interactions as well as quantum fluctuations on the emergence of the Rosensweig instability.

[1] H. Kadau, M. Schmitt, et al., [arXiv:1508.05007v2](https://arxiv.org/abs/1508.05007v2) (2015).

[2] H. Al-Jibbouri, I. Vidanović, A. Balaž, and A. Pelster, *J. Phys. B* **46**, 065303 (2013).

[3] R. N. Bisset and P. B. Blakie, [arXiv:1510.09013](https://arxiv.org/abs/1510.09013) (2015).

[4] K.-T. Xi and H. Saito, [arXiv:1510.07842](https://arxiv.org/abs/1510.07842) (2015).

[5] A. R. P. Lima and A. Pelster, *Phys. Rev. A* **84**, 041604(R) (2011); *Phys. Rev. A* **86**, 063609 (2012).

Q 17.3 Tue 11:45 e001

**Phonon to roton crossover and droplet formation in trapped dipolar Bose-Einstein condensates** — ●FALK WÄCHTLER and LUIS SANTOS — Institut für Theoretische Physik, Leibniz Universität Hannover, Hannover, Germany

The stability, elementary excitations, and instability dynamics of dipolar Bose-Einstein condensates depend crucially on the trap geometry. In particular, dipolar condensates in a pancake trap with its main plane orthogonal to the dipole orientation are expected to present under proper conditions a roton-like dispersion minimum, which if softening induces the so-called roton instability. On the contrary, cigar-shape traps are expected to present no dispersion minimum, and to undergo phonon (global) instability if destabilized. In this talk we investigate by means of numerical simulations of the non-local non-linear Schrödinger equation and the corresponding Bogoliubov-de Gennes equations the stability threshold as a function of the trap aspect ratio, mapping the crossover between phonon and roton instability. We will discuss in particular how this crossover may be observed in destabilization ex-

periments to reveal rotonization.

In a second part, motivated by recent experiments on droplet formation in Stuttgart, we introduce large conservative three-body interactions, and study how these forces affect the destabilization dynamics. We will discuss the ground-state physics of the individual droplets, and the crucial role that is played by the interplay between internal droplet energy, external center-of mass energy of the droplets, and energy dissipation in the nucleation of droplets observed in experiments.

Q 17.4 Tue 12:00 e001

**Lattice Physics with Ultracold Magnetic Erbium** — ●SIMON BAIER<sup>1</sup>, MANFRED J. MARK<sup>1,2</sup>, DANIEL PETER<sup>1</sup>, KIYOTAKA AIKAWA<sup>1</sup>, LAURIANE CHOMAZ<sup>1,2</sup>, ZI CAI<sup>2</sup>, MIKHAIL BARANOV<sup>2</sup>, PETER ZOLLER<sup>2,3</sup>, and FRANCESCA FERLAINO<sup>1,2</sup> — <sup>1</sup>Institut für Experimentalphysik, Universität Innsbruck, Technikerstraße 25, 6020 Innsbruck, Austria — <sup>2</sup>Institut für Quantenoptik und Quanteninformatik, Österreichische Akademie der Wissenschaften, 6020 Innsbruck, Austria — <sup>3</sup>Institut für Theoretische Physik, Universität Innsbruck, Technikerstraße 21A, 6020 Innsbruck, Austria

Strongly magnetic atoms are an ideal systems to study many-body quantum phenomena with anisotropic and long-range interactions. Here, we report on the first observation of the manifestation of magnetic dipolar interaction in extended Bose-Hubbard (eBH) dynamics by studying an ultracold gas of Er atoms in a three-dimensional optical lattice. We drive the superfluid-to-Mott-insulator (SF-to-MI) quantum phase transition and demonstrate that the dipolar interaction can favor the SF or the MI phase depending on the orientation of the atomic dipoles. The system is well described by the individual terms of the eBH Hamiltonian. This includes the onsite interaction, which, additional to the isotropic contact interaction, can be tuned with the dipole-dipole interaction by changing the dipole orientation and the shape of the onsite Wannier functions. We find for the first time the presence of the nearest-neighbor interaction between two adjacent particles. Future work will investigate dipolar effects with erbium molecules and fermions as well as spin physics in our lattice system.

Q 17.5 Tue 12:15 e001

**Strong-wave-turbulence character of non-thermal fixed points in Bose gases** — ●ISARA CHANTESANA<sup>1,2,3</sup> and THOMAS GASENZER<sup>2,3</sup> — <sup>1</sup>Institut für Theoretische Physik, Ruprecht-Karls-Universität Heidelberg, Philosophenweg 16, 69120 Heidelberg, Germany — <sup>2</sup>Kirchhoff Institut für Physik, INF 227, 69120 Heidelberg, Germany — <sup>3</sup>ExtreMe Matter Institute EMMI, GSI Helmholtzzentrum für Schwerionenforschung GmbH, Planckstraße 1, 64291 Darmstadt, Germany

Far-from equilibrium dynamics of a dilute Bose gas is studied by means of the two-particle irreducible effective action formalism. We investigate the properties of non-thermal fixed points predicted previously, which are related to non-perturbative strong wave turbulence solutions of the many-body dynamic equations. Instead of using a scaling analysis, we study the Boltzmann equation of the scattering integral by means of direct integration equation for sound waves. In this way we obtain a direct prediction of the scaling behaviour of the possible fixed-point solutions in the context of sound-wave turbulence. Implication for the real-time dynamics of the non-equilibrium system are discussed.

Q 17.6 Tue 12:30 e001

**Evidence of Non-Thermal Fixed Points in one-dimensional Bose gases** — ●SEBASTIAN ERNE<sup>1,2,4</sup>, ROBERT BÜCKER<sup>4</sup>, WOLFGANG ROHRINGER<sup>4</sup>, THOMAS GASENZER<sup>1,2,3</sup>, and JÖRG SCHMIEDMAYER<sup>4</sup> — <sup>1</sup>Institut für Theoretische Physik, Ruprecht-Karls-Universität Heidelberg, Philosophenweg 16, 69120 Heidelberg, Germany — <sup>2</sup>ExtreMe Matter Institute EMMI, GSI Helmholtzzentrum für Schwerionenforschung GmbH, Planckstraße 1, 64291 Darmstadt, Germany — <sup>3</sup>Kirchhoff-Institut für Physik, INF 227, 69120 Heidelberg, Germany — <sup>4</sup>Vienna Center for Quantum Science and Technology (VCQ), Atominstytut, TU Wien, Vienna, Austria

This work investigates the rapid cooling quench over the dimensional- and quasicondensate-crossover. Analyzing experiments performed at the Atominstytut, we study the relaxation of such a far-from equilibrium system. The early stage of condensate formation is dominated

by solitonic excitations, leading to a characteristic momentum distribution in agreement with a model of randomly distributed defects. The number of solitons increases with the quenchrates giving rise to an incompressible condensate. The isolated system follows a self-similar evolution governed by a universal time-independent nonthermal fixed point distribution. The dynamic universality classes of these nonequilibrium attractor solutions are relevant for a wide variety of physical systems ranging from relativistic high-energy physics to cold quantum gases. At later times of the evolution the system fully equilibrates leading to deviations from the self-similar evolution. Our results show a new way of condensation in far from equilibrium 1d Bose gases.

Q 17.7 Tue 12:45 e001

**Spin phonon dynamics with classical statistical methods**

— •ASIER PIÑEIRO ORIOLI<sup>1,2</sup>, ARGHAVAN SAFAVI-NAINI<sup>2</sup>, MICHAEL WALL<sup>2</sup>, and JOHANNES SCHACHENMAYER<sup>2</sup> — <sup>1</sup>Institute for Theoretical Physics, Heidelberg, Germany — <sup>2</sup>JILA, NIST and University of Colorado, Boulder, Colorado, USA

Systems with both spin and phonon degrees of freedom are ubiquitous in physical fields ranging from condensed matter to biophysics. However, methods to compute the dynamics of such systems are scarce, especially in high dimensions. In this work, we combine the Truncated Wigner Approximation (TWA) for bosons with its recently developed discrete version (dTWA) for spins to describe the dynamics of coupled spin-phonon systems. We benchmark the method by comparing to exact results and discuss applications to trapped-ion and cavity experiments.

## Parallel solvers for dipolar Gross-Pitaevskii equation

V. Lončar<sup>1</sup>, D. Vudragović<sup>1</sup>, S. K. Adhikari<sup>2</sup>, and A. Balaž<sup>1</sup>

<sup>1</sup>*Scientific Computing Laboratory, Center for the Study of Complex Systems,  
Institute of Physics Belgrade, University of Belgrade, Serbia*

<sup>2</sup>*Instituto de Física Teórica, UNESP – Universidade Estadual Paulista, São Paulo, Brazil*  
e-mail:vladimir.loncar@ipb.ac.rs

We present serial and parallel semi-implicit split-step Crank-Nicolson algorithms for solving the dipolar Gross-Pitaevskii equation [1, 2], used for study of ultracold Bose systems with the dipole-dipole interaction. Six parallel algorithms will be presented: C implementation parallelized with OpenMP targeting single shared memory system [3], CUDA implementation targeting single Nvidia GPU [4], hybrid C/CUDA implementation combining the two previous approaches, and their parallelizations to distributed memory systems using MPI [5]. We first give an overview of the split-step Crank-Nicolson method and describe how the dipolar term is computed using FFT, which forms the basis of all presented algorithms. We then move on to describing the concepts used in each of the parallel implementations, and finally present a performance evaluation of each algorithm. In our tests OpenMP implementation demonstrates a speedup of 12 on a 16-core workstation, CUDA version has a speedup of up to 13, hybrid version has a speedup of up to 16, while the MPI parallelization yields a further speedup of 16 for the OpenMP/MPI version, speedup of 10 for the CUDA/MPI version, and speedup of 6 for the hybrid version.

### REFERENCES

- [1] P. Muruganandam, et. al., *Comput. Phys. Commun.* 180, 1888 (2009).
- [2] R. Kishor Kumar, et. al., *Comput. Phys. Commun.* 195, 117 (2015).
- [3] D. Vudragović, et. al., *Comput. Phys. Commun.* 183, 2021 (2012).
- [4] V. Lončar, et. al., *Comput. Phys. Commun.* 200, 406 (2016).
- [5] V. Lončar, et. al., *Comput. Phys. Commun.* 209, 190 (2016).



UNIVERSITY OF NOVI SAD  
FACULTY OF SCIENCES  
DEPARTMENT OF  
MATHEMATICS AND INFORMATICS



Vladimir Lončar

# Hybrid parallel algorithms for solving nonlinear Schrödinger equation

— PhD thesis —

Advisors:

Dr. Antun Balaž

Dr. Srđan Škrbić

Novi Sad, 2017

# Contents

<b>Acknowledgements</b>	<b>iii</b>
<b>Abstract</b>	<b>iv</b>
<b>Sažetak</b>	<b>v</b>
<b>1 Introduction</b>	<b>1</b>
<b>2 Methods for solving nonlinear Schrödinger equation</b>	<b>8</b>
2.1 Overview of solutions of NLSE and GPE . . . . .	8
2.2 Dipolar GPE in 3D, 2D and 1D . . . . .	11
2.3 Split-step Crank-Nicolson method for the GPE . . . . .	14
2.4 Calculation of relevant physical quantities . . . . .	18
<b>3 Algorithm for shared memory systems</b>	<b>21</b>
3.1 Description of the algorithm . . . . .	21
3.2 Serial implementation . . . . .	24
3.3 Parallelization on shared memory systems . . . . .	33
<b>4 Solving NLSE using GPU accelerators</b>	<b>39</b>
4.1 CUDA programming and execution model . . . . .	40
4.2 CUDA implementation of shared memory algorithm . . . . .	44
<b>5 Hybrid algorithm for heterogeneous computing systems</b>	<b>52</b>
5.1 Description of the hybrid algorithm . . . . .	53
5.2 Implementation of the hybrid algorithm . . . . .	57
5.3 Optimization of data transfers . . . . .	62
<b>6 Distributed memory algorithm</b>	<b>66</b>
6.1 Data distribution scheme . . . . .	67
6.2 Distributing the computation over multiple processes . . . . .	71
6.3 Improvements of input and output operations . . . . .	74
<b>7 Interacting with simulation</b>	<b>78</b>
7.1 Visualization of data . . . . .	79
7.2 Simulation steering . . . . .	86

<b>8 Performance evaluation and modeling</b>	<b>90</b>
8.1 Optimization of input parameters for hybrid implementations . . . . .	90
8.2 Testing methodology . . . . .	97
8.3 Performance test results and modeling of single node programs . . . . .	99
8.4 Performance test results and modeling of MPI programs . . . . .	107
8.5 Selecting optimal algorithm . . . . .	116
<b>9 Demonstration of usability of developed programs</b>	<b>118</b>
9.1 Formation of vortices in BEC . . . . .	118
9.2 Effects of dipolar interactions on vortex formation in BEC . . . . .	122
<b>10 Conclusions and future work</b>	<b>125</b>
<b>Bibliography</b>	<b>127</b>
<b>Prošireni sažetak</b>	<b>134</b>
<b>Short biography</b>	<b>148</b>
<b>Kratka biografija</b>	<b>149</b>
<b>Key Words Documentation</b>	<b>150</b>
<b>Ključna dokumentacijska informacija</b>	<b>153</b>

# Acknowledgements

First and foremost I offer my sincerest gratitude to my advisors, Dr. Antun Balaž and Dr. Srđan Škrbić, for supporting me during these past six years. Antun has set an example of excellence as a researcher, mentor, instructor, and role model. His patience, guidance, encouragement and advice have made my PhD experience more productive and stimulating, and I thank him for the great effort he put into training me in the scientific field. Srđan has been supportive throughout my thesis and has given me the freedom to pursue my goals without objection. I especially thank him for his support during my transition to Scientific Computing Laboratory.

I also have to thank the members of my PhD committee, Professors Dragan Mašulović, Nataša Krejić and Miljko Satarić for their helpful advice and invaluable feedback.

I would like to give special thanks to my colleagues from the Scientific Computing Laboratory, for all productive discussions we had, for all the laughs we shared over lunch, and for the great time I have spent with you.

This thesis is written in the Scientific Computing Laboratory, Center for the Study of Complex Systems of the Institute of Physics Belgrade and is supported by the Ministry of Education, Science, and Technological Development of the Republic of Serbia under project ON171017. Numerical simulations were run on the PARADOX supercomputing facility at the Scientific Computing Laboratory of the Institute of Physics Belgrade.

Lastly, I would like to thank my family and friends for all their love and encouragement. For my parents who raised me with a love of science and supported me in all my pursuits. For my friends for providing support and friendship that I needed. Thank you.



# Abstract

Numerical methods and algorithms for solving of partial differential equations, especially parallel algorithms, are an important research topic, given the very broad applicability range in all areas of science. Rapid advances of computer technology open up new possibilities for development of faster algorithms and numerical simulations of higher resolution. This is achieved through parallelization at different levels that practically all current computers support.

In this thesis we develop parallel algorithms for solving one kind of partial differential equations known as nonlinear Schrödinger equation (NLSE) with a convolution integral kernel. Equations of this type arise in many fields of physics such as nonlinear optics, plasma physics and physics of ultracold atoms, as well as economics and quantitative finance. We focus on a special type of NLSE, the dipolar Gross-Pitaevskii equation (GPE), which characterizes the behavior of ultracold atoms in the state of Bose-Einstein condensation.

We present novel parallel algorithms for numerically solving GPE for a wide range of modern parallel computing platforms, from shared memory systems and dedicated hardware accelerators in the form of graphics processing units (GPUs), to heterogeneous computer clusters. For shared memory systems, we provide an algorithm and implementation targeting multi-core processors using OpenMP. We also extend the algorithm to GPUs using CUDA toolkit and combine the OpenMP and CUDA approaches into a hybrid, heterogeneous algorithm that is capable of utilizing all available resources on a single computer.

Given the inherent memory limitation a single computer has, we develop a distributed memory algorithm based on Message Passing Interface (MPI) and previous shared memory approaches. To maximize the performance of hybrid implementations, we optimize the parameters governing the distribution of data and workload using a genetic algorithm. Visualization of the increased volume of output data, enabled by the efficiency of newly developed algorithms, represents a challenge in itself. To address this, we integrate the implementations with the state-of-the-art visualization tool (VisIt), and use it to study two use-cases which demonstrate how the developed programs can be applied to simulate real-world systems.

# Sažetak

Numerički metodi i algoritmi za rešavanje parcijalnih diferencijalnih jednačina, naročito paralelni algoritmi, predstavljaju izuzetno značajnu oblast istraživanja, uzimajući u obzir veoma široku primenljivost u svim oblastima nauke. Veliki napredak informacione tehnologije otvara nove mogućnosti za razvoj brzih algoritama i numeričkih simulacija visoke rezolucije. Ovo se ostvaruje kroz paralelizaciju na različitim nivoima koju poseduju praktično svi moderni računari.

U ovoj tezi razvijeni su paralelni algoritmi za rešavanje jedne vrste parcijalnih diferencijalnih jednačina poznate kao nelinearna Šredingerova jednačina sa integralnim konvolucionim kernelom. Jednačine ovog tipa se javljaju u raznim oblastima fizike poput nelinearne optike, fizike plazme i fizike ultrahladnih atoma, kao i u ekonomiji i kvantitativnim finansijama. Teza se bavi posebnim oblikom nelinearne Šredingerove jednačine, Gros-Pitaevski jednačinom sa dipol-dipol interakcionim članom, koja karakteriše ponašanje ultrahladnih atoma u stanju Boze-Ajnštajn kondenzacije.

U tezi su predstavljeni novi paralelni algoritmi za numeričko rešavanje Gros-Pitaevski jednačine za širok spektar modernih računarskih platformi, od sistema sa deljenom memorijom i specijalizovanih hardverskih akceleratora u obliku grafičkih procesora, do heterogenih računarskih klastera. Za sisteme sa deljenom memorijom, razvijen je algoritam i implementacija namenjena višejezgarnim centralnim procesorima korišćenjem OpenMP tehnologije. Ovaj algoritam je proširen tako da radi i u okruženju grafičkih procesora korišćenjem CUDA alata, a takođe je razvijen i predstavljen hibridni, heterogeni algoritam koji kombinuje OpenMP i CUDA pristupe i koji je u stanju da iskoristi sve raspoložive resurse jednog računara.

Imajući u vidu inherentna ograničenja raspoložive memorije koju pojedinačan računar poseduje, razvijen je i algoritam za sisteme sa distribuiranom memorijom zasnovan na Message Passing Interface tehnologiji i prethodnim algoritmima za sisteme sa deljenom memorijom. Da bi se maksimalizovale performanse razvijenih hibridnih implementacija, parametri koji određuju raspodelu podataka i računskog opterećenja su optimizovani korišćenjem genetskog algoritma. Poseban izazov je vizualizacija povećane količine izlaznih podataka, koji nastaju kao rezultat efikasnosti novorazvijenih algoritama. Ovo je u tezi rešeno kroz integraciju implementacija sa najsavremenijim alatom za vizualizaciju (VisIt), što je omogućilo proučavanje dva primera koji pokazuju kako razvijeni programi mogu da se iskoriste za simulacije realnih sistema.

# Chapter 1

## Introduction

The study of dynamical evolution of various physical systems is essential for their understanding and represents a key ingredient for possible applications. It is still debated how deterministic nature is, but the majority of models we use to describe systems in nature are deterministic. This means that we use first and second order differential equations to model their dynamic behavior. Even quantum systems, although notorious for their probabilistic behavior, are described in the deterministic mathematical framework and their evolution can be cast into the form of second order differential equations. Apart from very simple models, systems with multiple degrees of freedom are usually described by partial differential equations. At the most fundamental level, these equations are linear and the superposition principle applies. However, due to large numbers of degrees of freedom for realistic systems, the complexity of the corresponding systems of equations is prohibitively large to address either analytically or numerically. Therefore, the effective models with much smaller dimensionality are often used, albeit not without introducing additional complexities, such as nonlinearity of the system. This is true for many-body physics problems, especially physics of ultracold atoms, plasma physics and many others. There are however, systems where nonlinearity is inherent, such as nonlinear optics. Whatever the case may be, nonlinearity is always a source of interesting and novel phenomena and has to be carefully studied. As usual, this translates to mathematical complexity in solving of the corresponding nonlinear partial differential equations. Therefore, the development of mathematical methods and computer algorithms for solving of such kind of equations is of very broad interest, not only in physics, but also in applied physics and engineering, astrophysics, chemistry, quantitative finance, etc.

As far as analytical methods are concerned, our options are limited to a small number of exactly solvable models or application of approximative techniques. While the number of such techniques is considerable, many of them have the disadvantage of lacking the proper way to estimate errors associated with the obtained results. Even for methods where such an estimate is possible, only the lowest order calculations can be done within reason, and calculation of higher order corrections usually requires an exponentially growing number of symbolic operations, making the approach intractable. This is the reason why numerical simulations are so widely used today, and they represent the main tool for studying large number of systems in nature, as well as social systems.

Numerical simulations of realistic systems are often computationally very intensive, substantially more so in higher number of dimensions. Therefore, they cannot be performed efficiently on a single, serial computing resource, and instead require parallel processing. Parallel algorithms for solving

partial differential equations are an important research topic today, given the fact that practically all current computers support some form of multiprocessing. The development of parallel algorithms for solving a particular type of partial differential equations, nonlinear Schrödinger equation (NLSE) type, represents the main topic of this thesis.

Multi-core central processing units (CPU) form the first level of parallelism available in all modern computers. Individual cores of the CPU may perform tasks in parallel while sharing a single view of data in memory. By extension, several multi-core CPUs may be packaged within a single computer and given access to the shared memory to further increase a system's parallel processing capabilities. Developing parallel algorithms for shared memory systems is relatively easy, however maximizing the efficiency of their implementation is often hindered by the fact that many systems do not provide uniform memory access times, instead relying on non-uniform memory access (NUMA) design. Also, data coherence issues arise when caching is employed within separate cores. Numerous hardware and software technologies exist to enable better utilization and simpler control of shared memory systems, with OpenMP emerging as the most widespread application programming interface for such computing platforms.

Next level of parallelism in this hierarchy is provided by the dedicated hardware accelerators, such as graphics processing units (GPU) and early Intel Xeon Phi coprocessors. These accelerators feature a vastly different hardware architecture to the one found in CPUs, often use their own separate main memory, and come packaged as an add-on for the existing computer. The difference in architecture is reflected in the programming model, whose successful utilization usually requires a significant modification of the underlying shared memory algorithm. In this thesis, we focus on a specific type of accelerator, the Nvidia GPU, which is programmed using CUDA toolkit. Combining parallelism of the CPU with the one found in the GPU to take full advantage of the resources one computer has is challenging. While CPU and GPU have separate memories, they cannot easily be programmed as a distributed memory system, and can instead be considered a heterogeneous computing resource, which requires development of hybrid algorithms to be fully exploited. These algorithms must shuffle data between CPU and GPU efficiently, overlap computation, and ensure correctness of concurrent operations through synchronization.

Combining several computers into a computer cluster forms the final level of parallelism. Even with the rate at which current CPU and GPU technologies advance, a single computer may not have enough memory, processing power, or both, to address the computational requirements of numerical simulations of high resolution. Therefore, computer clusters, comprised of large number of individual computer nodes interconnected via high-speed networking infrastructure, and accompanied by software solutions that allow for the whole system to be used in a parallel manner (e.g., Message Passing Interface, MPI) are created to fill this gap and complete the hierarchy. These systems have distributed memory, meaning that computational tasks on each node of the cluster operate only on data that are local to that node, and access to remote data is obtained through communication with one or more remote computational tasks. Today, each individual cluster node itself represents a shared memory system and the algorithm may or may not exploit this to attain finer level of parallelism granulation. Distributed memory algorithms must take all this into account and provide a data distribution scheme that facilitates efficient communication between computing nodes and maximizes the computation on local data.

Efficient use of such diverse computing systems requires domain-specific knowledge and the

development of parallel algorithms that take into consideration the specifics of the hardware the system contains, as well as parallelization of numerical methods used to solve the given problem. We explore all levels of parallelism across the entire hierarchy of modern hardware resources in this thesis.

As mentioned above, parallel algorithms for solving NLSE are of great importance in the physics of ultracold atoms, which studies the properties of matter at very low temperatures, at the nanokelvin level. At these temperatures, rarified gases of alkali metal elements exhibit a phase transition known as Bose-Einstein condensation [1, 2, 3, 4, 5], which produces a new state of quantum matter, characterized by a coherent, collective behavior. Bose-Einstein condensation is a purely quantum effect and does not have a classical counterpart. It is one of macroscopic quantum phenomena, together with superfluidity, superconductivity and lasers. Bose-Einstein condensates (BEC) and ultracold atoms are one of hot research topics because they provide previously unimaginable control over parameters, properties and behavior patterns of the system. In particular, short-range contact interaction strength in such systems can be tuned over many orders of magnitude, and can even switch the sign, i.e., can be made either attractive or repulsive. Furthermore, the dimensionality of the system can also be tuned using the external trapping potential, making possible crossovers from three spatial dimensions (3D) to two (2D) or one (1D). This can be even performed dynamically, by changing the shape of the trapping potential from a spherical one to disc-shaped or cigar-shaped one. BECs have possible applications in quantum computing [6, 7], as well as in quantum simulations of other systems. As of today, they are considered the only feasible Feynman simulators [8].

Despite this versatility, contact interactions in BEC systems, realized through *s*-wave scattering of atoms or molecules, are extremely short-ranged and usually have very limited effect on the properties of the system. Therefore, in the last decade special attention was devoted to the study of systems in which, alongside the omnipresent contact interactions, atoms or molecules also exhibit long-range dipole-dipole interaction. The dipolar BEC was first realized in 2005 [9] with chromium atoms ( $^{52}\text{Cr}$ ), which possess a small permanent magnetic dipolar moment. Since then, new atomic and molecular species with larger permanent or induced magnetic ( $^{168}\text{Er}$ ,  $^{164}\text{Dy}$ ) or electric ( $^{39}\text{K}^{87}\text{Rb}$ ,  $^{87}\text{Rb}^{133}\text{Cs}$ ,  $^7\text{Li}^{133}\text{Cs}$ ) dipolar moments are broadly investigated, and some of them are successfully condensed. Heteronuclear molecular species with large permanent electric dipole moments, which should become experimentally accessible in the next few years, will have orders of magnitude larger dipole-dipole interactions and will enable full exploration of the strong interaction regime. The presence of long-range interactions can lead to interesting new phenomena, such as phases with nontrivial order (ferromagnetic, antiferromagnetic, striped, etc.), and therefore this field is a topic of intensive research efforts by many groups worldwide.

This motivates us to study numerical algorithms for solving of nonlinear partial differential equations used for effective description of dipolar BEC systems. On one hand, it is an interesting problem in physics in itself, with many possible applications and expected new breakthroughs in the near future. On the other hand, it is also a challenging computer science problem that requires development of novel algorithms, parallelization and the use of latest computer technology.

To begin the development of a (parallel) computer algorithm, we start from the mathematical formulation of the problem, which at its core is a many-body physics problem, involving a Hamiltonian with quantum fields, i.e., the formalism of the second quantization. BEC systems are usually

realized with dilute atomic or molecular gases and therefore, in first-order approximation, we can neglect quantum fluctuations. Since the temperature is of the order of nanokelvin, we can neglect thermal excitations as well, which leads to the well-known mean-field theory [10, 11], with the field operators replaced by the classical fields, i.e., the single-particle wave function of the system. This effective wave function  $\Psi(\mathbf{r}; t)$  depends on position  $\mathbf{r} = (x, y, z)$  and time  $t$ , and satisfies mean-field *Gross-Pitaevskii equation* (GPE), which for the contact interaction case reads

$$i\hbar \frac{\partial}{\partial t} \Psi(\mathbf{r}; t) = \left[ -\frac{\hbar^2}{2m} \nabla^2 + V(\mathbf{r}) + gN_{\text{at}} |\Psi(\mathbf{r}; t)|^2 \right] \Psi(\mathbf{r}; t), \quad (1.1)$$

where  $m$  is the mass of the atomic or molecular species,  $\hbar$  is the reduced Planck constant,  $N_{\text{at}}$  is the number of particles in the condensate and  $V(\mathbf{r})$  is the trapping potential that confines the condensate, usually a harmonic function. As already mentioned, Eq. (1.1) includes only contact interaction, represented by a nonlinear term with the interaction strength  $g$  determined by the  $s$ -wave scattering length. Without the nonlinear term, Eq. (1.1) represents the usual Schrödinger equation, where the first term on the right-hand side is the kinetic energy, and  $V(\mathbf{r})$  is the potential. The cubic term that arises from atomic interactions makes the equation nonlinear, and therefore it is usually called nonlinear Schrödinger equation. It is worth noting that similar equations also arise in other fields of physics and engineering, in different contexts. The stability and existence of solutions to GPE strongly depends on the sign of nonlinearity  $g$ . Namely, if  $g$  is positive, which corresponds to repulsive interactions, the system is stable and the above equation always has a physically meaningful solution. On the other hand, if  $g$  is negative, the system is stable only up to a critical number of particles and beyond this number it collapses.

For atomic and molecular species with dipolar moments we have to take into account the dipolar interaction term, which yields the dipolar GPE in the form

$$i\hbar \frac{\partial}{\partial t} \Psi(\mathbf{r}; t) = \left[ -\frac{\hbar^2}{2m} \nabla^2 + V(\mathbf{r}, t) + gN_{\text{at}} |\Psi(\mathbf{r}; t)|^2 + N_{\text{at}} \int U_{\text{dd}}(\mathbf{r} - \mathbf{r}') |\Psi(\mathbf{r}'; t)|^2 d\mathbf{r}' \right] \Psi(\mathbf{r}; t). \quad (1.2)$$

The dipolar interaction term is given by a convolution integral that contains the position-dependent dipolar potential  $U_{\text{dd}}(\mathbf{r} - \mathbf{r}')$ . While contact interaction is fixed for a given system, the dipolar interaction strongly depends on the relative orientation and distance of the dipoles, and can be both attractive and repulsive. This greatly affects stability of the system, similarly to the case of pure contact interaction: if the interaction is predominantly attractive, the system is unstable beyond a critical number of particles, while for predominantly repulsive interaction the system is unconditionally stable.

The dynamics of dipolar BEC system is described using the time-dependent GPE presented above. Typically, we start from a given state  $\Psi(\mathbf{r}, t = 0)$ , and propagate it in time to obtain the state of the system at some future time  $t$ . However, if the system is in the stationary, ground state, its time propagation reduces to just phase rotation, so that it can be written as  $\Psi_0(\mathbf{r})e^{-i\mu t/\hbar}$ , where  $\Psi_0(\mathbf{r})$  is a real-valued function which satisfies the time-independent GPE

$$\mu \Psi_0(\mathbf{r}) = \left[ -\frac{\hbar^2}{2m} \nabla^2 + V(\mathbf{r}) + gN_{\text{at}} \Psi_0(\mathbf{r})^2 + N_{\text{at}} \int U_{\text{dd}}(\mathbf{r} - \mathbf{r}') |\Psi_0(\mathbf{r}')|^2 d\mathbf{r}' \right] \Psi_0(\mathbf{r}), \quad (1.3)$$

where  $\mu$  is the chemical potential of the system which corresponds to the energy of the system in that state. By employing the Wick rotation, the solution of time-independent equation can be obtained

formally by turning to the imaginary time, starting from an arbitrary initial state and propagating the wave function until the convergence is reached. Therefore, the development of methods and algorithms for solving of GPE in both real and imaginary time is of practical interest and enables us to fully study the behavior and properties of BEC systems, from calculation of stationary states to dynamical evolution of a given state subject to a particular experimental protocol.

Numerical methods for solving partial differential equations require significant computing resources, in particular in presence of nonlinearities. The nonlinearities such as those in dipolar GPE further increase the complexity of the corresponding algorithms, insofar as BEC systems are often of multi-scale nature. For example, if the system dynamically develops features at different, independent length scales, such as vortices, density waves or collective oscillations, the smallest scale dictates the discretization spacing. Thus, it may turn out that the minimal mesh size, which is necessary for accurate modeling of the system, is too large to be handled by a single processing element. The necessity of employing fine-grained discretization schemes is so prevalent in many areas where equations of the NLSE type are used that the development of parallel algorithms and their implementation represents an important challenge.

Therefore it is not surprising that this challenge was already addressed, and that there are several approaches developed and tailored specifically for solving of GPE. Most of them are inherently serial and focus solely on the equation for the contact interaction only. Algorithms and software implementations for the dipolar GPE are practically missing, in particular parallel algorithms do not exist in public domain or in the literature. This thesis addresses the identified development gap and provides a significant step forward in the area of numerical methods and parallel algorithms for solving equations the NLSE type, of interest for the fields of ultracold atoms, nonlinear optics, and other areas.

Research results presented in this thesis contribute to the fields of scientific computing and parallel programming, as well as complexity modeling and optimization of hybrid algorithms. In particular, the main contribution of this thesis is the development of parallel algorithms for solving of GPE and related equations of the NLSE type. The thesis proposes and implements several parallel algorithms for a range of modern computing architectures:

- Shared memory systems consisting of one or more multi-core CPUs;
- Hardware accelerators in the form of GPUs;
- Heterogeneous systems combining the multi-core CPUs with the GPUs;
- Distributed memory systems - computer clusters with multi-core CPU nodes;
- Distributed memory systems - computer clusters with GPU-enabled nodes;
- Distributed memory systems - computer clusters with multi-core CPU & GPU-enabled nodes.

Development and implementation of the hybrid algorithms that fully overlap the computation on all resources required major research effort, however, tuning them to get the best performance on a specific hardware combination was even more challenging. This is addressed by automating the tuning of hybrid implementations via a genetic algorithm, where a general method was developed such that it can be applied to other heterogeneous algorithms. In addition to the development of the above parallel algorithms, the optimal data distribution scheme for hybrid algorithms targeting

heterogeneous systems is also presented. This scheme was selected during the research leading to this thesis and represents another contribution, as can be verified by its use for the hybrid Fourier transform algorithm, which was developed within the thesis as well. The hybrid Fourier transform algorithm is highly efficient as it enables simultaneous computation on both CPU and GPU, and can be reused independently for other purposes, thus representing an important contribution in itself. The thesis concludes its research by investigating and modeling the complexity of the developed parallel algorithms. The proposed models are tested and experimentally verified through a thorough measurement of the performance of developed programs. Finally, the thesis investigates the possibilities for visualization of large-scale simulations in real time and demonstrates its practical implementation on the example of study of vortex formation in dipolar BECs.

Following on from this introductory chapter, we discuss approaches to solving NLSE and specifically GPE along with the description of the numerical algorithm employed by all programs developed as part of this thesis, in Chapter 2. The chapter starts with an overview of both analytical and well as numerical methods for solving NLSE and GPE, accompanied by a survey of available solvers based on the numerical methods (Section 2.1). Next, we present the dimensionless dipolar GPEs in 3D, 2D and 1D (Section 2.2) which we numerically solve in this thesis, before moving on to the description of the semi-implicit split-step Crank-Nicolson numerical method which forms the basis of the algorithms used in the implementations (Section 2.3). In addition to the main method to solve the dipolar GPE, in Section 2.4 we give expressions for several relevant physical quantities that can be calculated from the wave function and used to study properties of BEC systems.

With the main numerical method described, we proceed to the description of the shared memory algorithm in Chapter 3. In Section 3.1 we describe how the numerical method can be used to construct a computer algorithm targeting shared memory systems, followed by a description of the serial implementation in Section 3.2 and parallel implementation in Section 3.3.

The shared memory algorithm can be extended to hardware accelerators like graphics processing units (GPUs), which we demonstrate in Chapter 4. To get a better understanding of how the algorithm can be implemented, we provide a brief overview of the main concepts of the programming on GPU in Section 4.1. This is followed by the description of the implementation of the shared memory algorithm using CUDA in Section 4.2.

Combining the algorithms from Chapter 3 and 4 to produce a new, hybrid, algorithm which simultaneously employs both CPU and GPU is discussed in Chapter 5. The hybrid algorithm is described in Section 5.1, followed by a naive implementation in Section 5.2, where we also note the inherent problems with such approach. The ways to improve it are discussed in Section 5.3.

Given that the single computer may not satisfy the requirements of large-scale simulations in terms of amount of processing power or available memory, in Chapter 6 we show how the algorithms from Chapters 3, 4 and 5 can be extended with MPI to work with distributed memory systems like computer clusters. To enable this, we demonstrate how data are distributed to separate processes in Section 6.1. Working with only a subset of data required changes to the way we perform computation in the main loop of our algorithms, which we describe in Section 6.2. In Section 6.3 we show how distributed memory processing can be used to improve the input and output (I/O) operations of the developed programs.

To get a better understanding of the results of the simulations created with our programs and enable study of the properties of BEC systems, in Chapter 7 we investigate how the programs can be



extended to provide easy and efficient visualization of produced data. We show how the programs can be integrated into a visualization tool (Section 7.1), and how the running simulations can be made interactive (Section 7.2).

Chapter 8 is dedicated to the detailed performance evaluation of all developed programs. To extract the best performance out of the hybrid implementations, we developed several optimization techniques which automatically tune the parameters governing the division of work between CPU and GPU, details of which are given in Section 8.1. This is followed by the description of the methodology used for the tests (Section 8.2), and results of the tests along with their modeling are presented in Sections 8.3 and 8.4.

An example of how the programs can be used for research into the behavior of BEC systems is given in Chapter 9. We demonstrate how the real-world experiments can be simulated in Section 9.1, which also serves as the ultimate verification not only of the algorithms and their implementations, but also of the physical model and mean-field approximation applied. In Section 9.2 we show how simulations of potential new phenomena can be performed using our programs on the example of study of the dipole-dipole interaction effects on the critical velocity for the emergence of vortices in a dipolar BEC.

We provide a summary of the work done as part of this thesis and an outlook of future research directions in Chapter 10.

## Chapter 2

# Methods for solving nonlinear Schrödinger equation

Nonlinear partial differential equations are known to be difficult to solve. The same is true for the GPE and NLSE in general, although their nonlinearity is not that severe. Numerous methods for solving such kind of equations have been developed, each with different advantages and setbacks. In this chapter we discuss relevant methods for solving one form of NLSE known as dipolar GPE and provide details on the numerical method underpinning the algorithms presented in this thesis.

In Section 2.1 we give an overview of well-known methods used to solve GPE and similar NLSE. We cover both analytic approaches and numerical methods, and go over the most popular solvers. The dimensionless form of dipolar GPE in 3D as well as the effective equations of reduced dimensionality are presented in Section 2.2. Our numerical method of choice, split-step semi-implicit Crank-Nicolson method, upon which all algorithms in this thesis are based, is the subject of Section 2.3. Relevant physical quantities which are usually calculated during numerical simulations of imaginary- and real-time propagation of the GPE are discussed in Section 2.4.

### 2.1 Overview of solutions of NLSE and GPE

While linear differential equations of different kinds can be generally always solved and the superposition principle allows for a construction of solutions satisfying any given initial (or boundary) conditions, this is not the case with nonlinear differential equations, including NLSE and GPE. In general, unless a special method exists for a particular type of nonlinear equations, one has to use approximative approaches, either analytical or numerical. Here we discuss the most relevant methods for NLSE-type of equations.

Perturbative approach, where the unknown function is expressed as a power series in an appropriate small parameter, is one of the most popular analytical methods. In the case of NLSE and GPE, it leads to a hierarchical system of coupled linear equations in the coefficients of the power series. These equations are coupled due to nonlinearity, and usually cannot be solved analytically. The hierarchy has to be cut at a given order of the smallness parameter, which brings an infinitely large system to a finite size and enables its solving. At the same time, cutting the hierarchy introduces an approximation and associated error to the obtained solution.

There are also physics-motivated approaches, where one neglects certain terms in NLSE/GPE if they are known to be sufficiently small. For example, when studying the ground state of a BEC with strong interactions, the kinetic energy is usually small compared to other terms, and can be neglected. This is done in the Thomas-Fermi approximation, which reduces the corresponding time-independent NLSE/GPE into an algebraic equation. This method can be also applied in real-time propagation if kinetic energy is guaranteed to always remain sufficiently small, however this cannot be known in advance and has to be independently verified. Thomas-Fermi approximation can be reliably used only for very strong interactions or in the limit of infinite interaction. In the other limit, when interaction is very weak, NLSE/GPE reduces to the usual Schrödinger equation, whose solutions can be obtained by the methods tailored for this type of equations, including known analytical solutions. One can even combine these methods with the perturbative expansion in the small interaction strength.

Another physics-motivated approach is the variational method, where one assumes certain mathematical form of the solution containing several unknown parameters (ansatz). These parameters are then optimized and their values are obtained (usually) by minimizing the corresponding functional (action or energy, in the language of physics). The approximative solution obtained in this way is the best in a given class of functions, however it is difficult to assess its quality or the associated error. Selection of a suitable ansatz depends on our knowledge of the behavior of the system, i.e., solutions for particular cases or limits, numerical solutions for some values of system parameters, or experimental data. The more we know about the properties of the system, we are able to build this knowledge into an ansatz that describes more reliably the dynamical behavior and possible solutions of the equations we seek to solve.

In addition to the analytic methods, various numerical methods also play an important role, as in many cases they are the only way to study the given physical system in detail. These methods can be roughly categorized in three groups:

- split-step methods,
- finite difference methods and
- spectral methods.

In split-step methods, propagation in time is separated into several substeps, such that each of them can be performed either analytically or numerically. In addition to errors associated with each individual substep, splitting the time propagation also introduces an overall error that has to be taken into account when estimating the total error of the solution. Finite difference methods rely on the discretization of the corresponding equation and solving it as a system of linear algebraic equations. Note that variations of these methods are usually combined with other approaches, in particular with split-step methods. Spectral methods assume that the solution is written in a complete basis of orthonormal special functions, and the equations for the corresponding coefficients are calculated numerically. The choice of a basis is made so that it leads to numerically tractable system of equations, with most popular examples being the Fourier or Laplace decomposition.

In the case of the dipolar GPE, the convolution integral representing the dipole-dipole interaction term is usually calculated via the discrete Fourier transformation (DFT), which has to be combined with the chosen method of solution.

There are several available solvers for NLSE and GPE and we briefly review the most popular ones: BEC-GP [12, 13, 14, 15] and DBEC-GP [16, 17, 18] families of solvers, TS-MPI [19], GPUE [20], NLSEmagic [21], GPELab [22, 23], GSGPEs [24], ATUS-PRO [25] and GPFEM [26]. All of these solvers are publicly available and published in journals such as Computer Physics Communications.

Published serial implementations have been in development since the 2000s, however they are largely superseded by the parallel implementations. One such notable and highly cited Fortran implementation by Muruganandam and Adhikari [12] combines split-step and finite difference methods, and utilizes Crank-Nicolson scheme to produce solvers for 1D, 2D and 3D GPE with contact interaction only. Parallelizations of these solvers have been developed in C with OpenMP [13] and MPI [14], and in Fortran with OpenMP [15]. New branch that includes the dipolar interaction was published in 2015 [16], and we have subsequently parallelized these solvers using CUDA [17], as well as OpenMP and MPI [18]. Parallelizations of the dipolar programs are presented in this thesis in detail. Note that only the solvers in this branch are capable of taking into account the dipolar interaction. As far as other software packages mentioned below are concerned, only one of them (GPELab) has this capability, while all others can solve GPE with the contact interaction only.

One of such parallel solvers is the Trotter-Suzuki-MPI library (TS-MPI) [19]. This library provides a parallel and distributed implementation based on the Trotter-Suzuki algorithm [27]. The library contains kernels for efficient computation on CPU and GPU, as well as a hybrid implementation, allowing for the simultaneous use of CPU and GPU. Additionally, TS-MPI exposes several external application programming interfaces (APIs), and can thus be used from Python or C++. Unfortunately, TS-MPI library can only be used for 1D and 2D problems, and there is no support for 3D systems.

Specialized GPU solvers also exist, an example of which are GPUE [20] and NLSEmagic [21]. GPUE is based on a time-splitting pseudospectral method, implemented in CUDA and targeting only 2D systems. Being implemented in CUDA, it relies on external Python routines to process the output. It supports only a single GPU, and there is no support for multiple GPUs within the same computer or any form of distributed computation. NLSEmagic can be used to simulate the NLSE in 1D, 2D and 3D. It utilizes a fully-explicit fourth-order Runge-Kutta scheme in time and both second- and fourth-order finite difference scheme in space. While implemented in C and CUDA, NLSEmagic is designed to interface with MATLAB. Like GPUE, it also uses only a single GPU.

There are also solvers implemented fully in MATLAB, like GPELab [22, 23] and GSGPEs [24]. GPELab is a toolbox able to solve stationary and dynamics problems in 1D, 2D and 3D, including the contact, dipole-dipole and user-provided interactions, as well as rotation of the system, stochastic effects and multi-components problems. It is based on a semi-implicit backward Euler scheme, a pseudo-spectral approximation and a Krylov subspace method [28]. While very popular due to its versatility, GPELab was shown to have very bad performance when compared to implementations based on compiled languages [29]. The other MATLAB-based implementation, GSGPEs, has fewer features and can only be used for the computation of the ground state of systems. It also relies on spectral decomposition method to perform this task.

Two solvers based on the finite element method are ATUS-PRO [25] and GPFEM [26]. ATUS-PRO calculates the solutions of the stationary and the time-dependent 1D, 2D and 3D GPE with contact interaction. Programs in this package are implemented in C++ by means of the deal.II

library, and support running in parallel on distributed memory systems using MPI. Numerical method employed in the solver is based on the fully implicit Crank-Nicolson method, while stationary states are obtained with a constrained Newton-like method. The other finite element-based solver, GPFEM, can be used to compute the stationary solutions of GPE with rotation, in 2D or 3D. It is implemented in FreeFem++, a free finite element toolbox with its own language. An interesting feature of this solver is that it offers two numerical methods for minimization of GP energy for the user to choose from: a steepest descent method based on Sobolev gradients and a minimization algorithm based on the optimization library Ipopt. This toolkit also provides several post-processing tools for tasks such as identification of quantized vortices, that can help in extracting values of physical quantities from the simulations.

As we can see from this overview, many of the previous implementations focus on a specific problem (e.g., computation of the stationary state), or on a specific dimensionality. Only DBEC-GP solver and GPELab are able to address the dipolar interaction, however the first one is serial and the other one is implemented in MATLAB, leading to their poor performance, in particular when treating 3D systems. One of the aims of this thesis is to overcome this problem and to develop and implement parallel algorithms, which are presented in later chapters.

## 2.2 Dipolar GPE in 3D, 2D and 1D

Many important properties of BEC systems at low temperatures can be successfully described using the mean-field theory, where quantum fluctuations are neglected and where we assume zero temperature, as already mentioned in Chapter 1. If we take into account scattering of the atoms or molecules through the nonlinear contact interaction term, and the dipole-dipole interaction assuming a fixed orientation of the dipoles along  $z$  direction, the corresponding mean-field Gross-Pitaevskii equation has a form

$$i\hbar \frac{\partial}{\partial t} \Psi(\mathbf{r}; t) = \left[ -\frac{\hbar^2}{2m} \nabla^2 + V(\mathbf{r}) + gN_{\text{at}} |\Psi(\mathbf{r}; t)|^2 + N_{\text{at}} \int U_{\text{dd}}(\mathbf{r} - \mathbf{r}') |\Psi(\mathbf{r}'; t)|^2 d\mathbf{r}' \right] \Psi(\mathbf{r}; t), \quad (2.1)$$

where the nonlinearity  $g = 4\pi\hbar^2 a_s/m$  is determined by the  $s$ -wave scattering length  $a_s$ , a quantity used in atomic physics for characterizing the interactions of atoms in the low-energy limit.  $N_{\text{at}}$  is the total number of condensed particles and  $V(\mathbf{r})$  is the trapping potential, usually anisotropic quadratic function of the form

$$V(\mathbf{r}) = \frac{1}{2} m(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2), \quad (2.2)$$

where  $\omega_x$ ,  $\omega_y$  and  $\omega_z$  are the corresponding frequencies. However, in general, it may be any arbitrary function depending on position and time  $V(\mathbf{r}, t)$ . The normalization condition of the wave function is given by

$$\int d\mathbf{r} |\Psi(\mathbf{r}; t)|^2 = 1, \quad (2.3)$$

which ensures that the number of particles in the system is fixed to  $N_{\text{at}}$  at all times. The boundary conditions the solution has to satisfy read

$$\lim_{\mathbf{r} \rightarrow \pm\infty} \Psi(\mathbf{r}; t) = 0. \quad (2.4)$$

The dipolar interaction potential is given by

$$U_{\text{dd}}(\mathbf{R}) = \frac{C_{\text{dd}}}{4\pi} \frac{1 - 3 \cos^2 \theta}{|\mathbf{R}|^3}, \quad (2.5)$$

where  $\mathbf{R} = \mathbf{r} - \mathbf{r}'$  determines relative position of the dipoles,  $\theta$  is the angle between  $\mathbf{R}$  and the orientation of dipoles, and  $C_{\text{dd}} = \mu_0 \bar{\mu}^2$  for magnetic dipoles and  $C_{\text{dd}} = d^2/\epsilon_0$  for electric dipoles. To compare contact and dipolar interactions, a new length scale is introduced, characterizing the dipole-dipole interaction:

$$a_{\text{dd}} = \frac{C_{\text{dd}} m}{12\pi \hbar^2}, \quad (2.6)$$

which also allows the dipolar GPE to be written in a more compact way.

In order to numerically simulate GPE, we first have to cast it into dimensionless form and to express all quantities through dimensionless variables. To achieve this, we start by choosing a reference frequency  $\tilde{\omega}$ , which may be equal to one of the trap frequencies, or have some other value. Now we can express the trap frequencies by the corresponding dimensionless quantities,  $\omega_x = \gamma \tilde{\omega}$ ,  $\omega_y = \nu \tilde{\omega}$  and  $\omega_z = \lambda \tilde{\omega}$ . We also use  $\tilde{\omega}$  to define a unit of time  $1/\tilde{\omega}$  and to introduce the dimensionless time as  $\tilde{t} = t \tilde{\omega}$ . The corresponding harmonic oscillator length  $l = \sqrt{\hbar/m\tilde{\omega}}$  is used as a unit of length, and the dimensionless coordinates are defined as  $\tilde{x} = x/l$ ,  $\tilde{y} = y/l$  and  $\tilde{z} = z/l$ . Finally, we express the wave function in a dimensionless form by  $\tilde{\Psi} = l^{3/2} \Psi$ . Using this rescaling, GPE (2.1) reduces to

$$i \frac{\partial}{\partial \tilde{t}} \tilde{\Psi}(\tilde{\mathbf{r}}; \tilde{t}) = \left[ -\frac{1}{2} \tilde{\nabla}^2 + \frac{1}{2} (\gamma^2 \tilde{x}^2 + \nu^2 \tilde{y}^2 + \lambda^2 \tilde{z}^2) + 4\pi \tilde{a}_s N_{\text{at}} |\tilde{\Psi}(\tilde{\mathbf{r}}; \tilde{t})|^2 + 3\tilde{a}_{\text{dd}} N_{\text{at}} \int d\tilde{\mathbf{r}}' V_{\text{dd}}(\tilde{\mathbf{r}} - \tilde{\mathbf{r}}') |\tilde{\Psi}(\tilde{\mathbf{r}}'; \tilde{t})|^2 \right] \tilde{\Psi}(\tilde{\mathbf{r}}; \tilde{t}), \quad (2.7)$$

where the dimensionless dipolar interaction potential reads

$$V_{\text{dd}}(\tilde{\mathbf{R}}) = \frac{1 - 3 \cos^2 \theta}{|\tilde{\mathbf{R}}|^3}, \quad (2.8)$$

and we have introduced the dimensionless scattering length  $\tilde{a}_s = a_s/l$  and the dipole-dipole interaction length  $\tilde{a}_{\text{dd}} = a_{\text{dd}}/l$ .

It is worth noting that the above conversion to dimensionless quantities is not the only way to do so. For example, we can additionally rescale the time so that instead of  $\tilde{t}$  we use  $2\tilde{t}$ . This would make the factor of  $1/2$  disappear from the kinetic energy and harmonic potential terms, which may be preferred in some cases. If we additionally drop the tilde symbol from all quantities, GPE takes the form

$$i \frac{\partial}{\partial t} \Psi(\mathbf{r}; t) = \left[ -\nabla^2 + \gamma^2 x^2 + \nu^2 y^2 + \lambda^2 z^2 + g N_{\text{at}} |\Psi(\mathbf{r}; t)|^2 + g_{\text{dd}} N_{\text{at}} \int d\mathbf{r}' V_{\text{dd}}(\mathbf{r} - \mathbf{r}') |\Psi(\mathbf{r}'; t)|^2 \right] \Psi(\mathbf{r}; t), \quad (2.9)$$

where the dimensionless nonlinearity is defined as  $g = 8\pi a_s/l$  and the dimensionless dipolar interaction strength is  $g_{\text{dd}} = 6a_{\text{dd}}/l$ .

In many cases, the geometry of the system is such that one can use effective equations with reduced dimensionality. Lower dimensional equations are analytically more complex and less tractable,

but their numerical complexity significantly reduces and thus this effective approach is widely used whenever possible.

In disc-shaped external potentials, when one of the trapping frequencies is much larger than the other two, we can effectively use 2D geometry and derive appropriate equations assuming that a BEC system remains confined to a ground state of harmonic oscillator in the third (confining) direction. We can consider two physically distinct cases, when the dipoles are oriented along the confining direction ( $\lambda \gg \gamma, \nu$ ) and when the dipoles are in the 2D plane (e.g.,  $\nu \gg \gamma, \lambda$ ). The corresponding dimensionless effective equations have been derived in Ref. [16]. For the first case, when the dipoles are orthogonal to the plane, the wave function is written as  $\Psi(\mathbf{r}; t) = \Psi_0(z; \lambda)\Psi_{2D}(\mathbf{r}_{2D}; t)$ , where the ground state in  $z$  direction is given as

$$\Psi_0(z; \lambda) = \left(\frac{\lambda}{\pi}\right)^{1/4} e^{-\lambda z^2/2}, \quad (2.10)$$

and  $\mathbf{r}_{2D} = (x, y)$  is 2D radius vector. Using this substitution, GPE (2.9) reduces to

$$i \frac{\partial}{\partial t} \Psi_{2D}(\mathbf{r}_{2D}; t) = \left[ -\nabla_{2D}^2 + \gamma^2 x^2 + \nu^2 y^2 + g_{2D} N_{\text{at}} |\Psi_{2D}(\mathbf{r}_{2D}; t)|^2 + g_{\text{dd},2D} N_{\text{at}} \int \frac{d\mathbf{k}_{2D}}{(2\pi)^2} e^{-i\mathbf{k}_{2D} \cdot \mathbf{r}_{2D}} \tilde{n}_{2D}(\mathbf{k}_{2D}; t) h_{2D}\left(\frac{k_{2D}}{\sqrt{2\lambda}}\right) \right] \Psi_{2D}(\mathbf{r}_{2D}; t), \quad (2.11)$$

where  $\nabla_{2D}$  represents 2D gradient operator, the effective dimensionless interaction strengths are  $g_{2D} = \sqrt{\lambda/(2\pi)} 8\pi a_s/l$  and  $g_{\text{dd},2D} = \sqrt{\lambda/(2\pi)} 8\pi a_{\text{dd}}/l$ ,  $\mathbf{k}_{2D} = (k_x, k_y)$  is the 2D  $k$ -vector, and

$$\tilde{n}_{2D}(\mathbf{k}_{2D}; t) = \int d\mathbf{r}_{2D} e^{i\mathbf{k}_{2D} \cdot \mathbf{r}_{2D}} |\Psi_{2D}(\mathbf{r}_{2D}; t)|^2, \quad (2.12)$$

$$h_{2D}(\xi) = 2 - 3\sqrt{\pi}\xi \exp(\xi^2) \{1 - \text{erf}(\xi)\}. \quad (2.13)$$

For the second case, when dipoles lie in the plane, the wave function can be expressed as  $\Psi(\mathbf{r}; t) = \Psi_0(y; \nu)\Psi_{2D}(\mathbf{r}_{2D}; t)$ , where the ground state is defined by Eq. (2.10) and 2D radius vector is now  $\mathbf{r}_{2D} = (x, z)$ . The corresponding effective 2D GPE has a form

$$i \frac{\partial}{\partial t} \Psi_{2D}(\mathbf{r}_{2D}; t) = \left[ -\nabla_{2D}^2 + \gamma^2 x^2 + \lambda^2 z^2 + g_{2D} N_{\text{at}} |\Psi_{2D}(\mathbf{r}_{2D}; t)|^2 + g_{\text{dd},2D} N_{\text{at}} \int \frac{d\mathbf{k}_{2D}}{(2\pi)^2} e^{-i\mathbf{k}_{2D} \cdot \mathbf{r}_{2D}} \tilde{n}_{2D}(\mathbf{k}_{2D}; t) j_{2D}\left(\frac{k_{2D}}{\sqrt{2\nu}}, k_z\right) \right] \Psi_{2D}(\mathbf{r}_{2D}; t), \quad (2.14)$$

where  $\nabla_{2D}$  now represents 2D gradient operator in  $x$ - $z$  plane,  $g_{2D} = \sqrt{\nu/(2\pi)} 8\pi a_s/l$  and  $g_{\text{dd},2D} = \sqrt{\nu/(2\pi)} 8\pi a_{\text{dd}}/l$  are the effective interaction strengths,  $\mathbf{k}_{2D} = (k_x, k_z)$  is the corresponding 2D  $k$ -vector, and

$$j_{2D}(\xi, k_z) = -1 + 3\sqrt{\pi} \frac{k_z^2}{2\nu\xi} \exp(\xi^2) \{1 - \text{erf}(\xi)\}. \quad (2.15)$$

Similar dimensional reduction can be performed for BEC systems in elongated, cigar-shaped traps [16], where again we can distinguish two cases: when the dipoles are orthogonal to the trap axis (without loss of generality, we consider just the geometry  $\nu, \lambda \gg \gamma$ ) and when they are parallel to the trap axis ( $\gamma, \nu \gg \lambda$ ). In a cigar-shaped trap, a BEC system is assumed to remain confined to a ground state of a harmonic oscillator in two transversal directions ( $y$  and  $z$  in the first case,

and  $x$  and  $y$  in the second), and we consider the wave function to depend only on one remaining coordinate. In the first case, the wave function is written as  $\Psi(\mathbf{r}; t) = \Psi_0(y; \nu)\Psi_0(z; \lambda)\Psi_{1D}(x; t)$ , and assuming axial symmetry ( $\nu = \lambda$ ) the effective 1D GPE can be expressed in a closed form

$$i \frac{\partial}{\partial t} \Psi_{1D}(x; t) = \left[ -\frac{\partial^2}{\partial x^2} + \gamma^2 x^2 + g_{1D} N_{\text{at}} |\Psi_{1D}(x; t)|^2 + g_{\text{dd},1D} N_{\text{at}} \int_{-\infty}^{\infty} \frac{dk_x}{2\pi} e^{-ik_x x} \tilde{n}_{1D}(k_x; t) j_{1D} \left( \frac{k_x}{\sqrt{2\nu}} \right) \right] \Psi_{1D}(x; t), \quad (2.16)$$

where interactions strengths are  $g_{1D} = 4\nu a_s/l$  and  $g_{\text{dd},1D} = 4\nu a_{\text{dd}}/l$ , and

$$\tilde{n}_{1D}(k_x; t) = \int_{-\infty}^{\infty} dx e^{ik_x x} |\Psi_{1D}(x; t)|^2, \quad (2.17)$$

$$j_{1D}(\xi) = \frac{1}{\sqrt{\pi}} \int_{-\infty}^{\infty} dk_y e^{-k_y^2} h_{2D} \left( \sqrt{\xi^2 + k_y^2} \right). \quad (2.18)$$

In the second case, when dipoles are parallel to the trap axis ( $\gamma, \nu \gg \lambda$ ), the wave function is written as  $\Psi(\mathbf{r}; t) = \Psi_0(x; \gamma)\Psi_0(y; \nu)\Psi_{1D}(z; t)$ . Again, assuming axial symmetry ( $\gamma = \nu$ ) the closed form of the effective 1D GPE reads

$$i \frac{\partial}{\partial t} \Psi_{1D}(z; t) = \left[ -\frac{\partial^2}{\partial z^2} + \lambda^2 z^2 + g_{1D} N_{\text{at}} |\Psi_{1D}(z; t)|^2 + g_{\text{dd},1D} N_{\text{at}} \int_{-\infty}^{\infty} \frac{dk_z}{2\pi} e^{-ik_z z} \tilde{n}_{1D}(k_z; t) h_{1D} \left( \frac{k_z}{\sqrt{2\gamma}} \right) \right] \Psi_{1D}(z; t), \quad (2.19)$$

where interaction strengths are  $g_{1D} = 4\gamma a_s/l$  and  $g_{\text{dd},1D} = 2\gamma a_{\text{dd}}/l$ , and

$$h_{1D}(\xi) = \int_0^{\infty} du \left[ \frac{3\xi^2}{u + \xi^2} - 1 \right] e^{-u}. \quad (2.20)$$

The above equations in 3D, as well as the effective equations in 2D and 1D are solved by the algorithms and programs presented in this thesis.

## 2.3 Split-step Crank-Nicolson method for the GPE

With the dimensionless form of GPE introduced in the previous section, we proceed with the description of *split-step semi-implicit Crank-Nicolson method* [30], the approach we use to numerically solve GPE in 3D as well as effective 2D and 1D equations. The basic idea behind this method is to divide the calculation of the solution into multiple steps at two levels. First, time propagation is discretized and the total propagation time  $T$  is divided into  $N$  time steps, each of length  $\Delta = T/N$ . At the second level, each time step is divided into several substeps, dealing with different parts of the Hamiltonian governing the time propagation. This second level of splitting uses semi-implicit Crank-Nicolson (CN) method, which we describe in detail here.

We first demonstrate how this method works on a simple example in 1D, before moving on to the solution of full dipolar GPE. Let us consider the equation of the general form

$$i \frac{\partial}{\partial t} \Psi(x; t) = H \Psi(x; t), \quad (2.21)$$

where the Hamiltonian  $H$  contains the kinetic term, i.e., the Laplacian, as well as other non-derivative terms, including the potential and nonlinear terms. Time propagation starts from a



known initial condition at  $t = t_0$  and proceeds in  $N$  time steps of the duration  $\Delta$ . To propagate the wave function in one time step, we split  $H$  into two parts,  $H = H_1 + H_2$ , where  $H_1$  contains all non-derivative terms, while  $H_2$  is the remaining Laplacian. This allows us to make an approximation and split Eq. (2.21) into two parts, which we consider and solve separately, one after the other,

$$i \frac{\partial}{\partial t} \Psi(x; t) = H_1 \Psi(x; t), \quad (2.22)$$

$$i \frac{\partial}{\partial t} \Psi(x; t) = H_2 \Psi(x; t). \quad (2.23)$$

Equation (2.22) is solved first, with an initial value  $\Psi(x; t_0)$  at  $t = t_0$ , to obtain an intermediate solution at  $t = t_0 + \Delta$ . This intermediate solution is then used as the initial value for Eq. (2.23), yielding the final solution  $\Psi(x; t_0 + \Delta)$  at  $t = t_0 + \Delta$ . This procedure then continues in the next time step until we finish time propagation. As we mentioned above, the time splitting is an approximation and introduces an error of the order of  $\Delta^2$ , which may be neglected when used with small time step  $\Delta$ . However, such a scheme allows us to treat the larger part of the Hamiltonian exactly, even for arbitrarily large nonlinear terms, as we show below.

Let us denote the wave function at time  $t_n = t_0 + n\Delta$  by  $\Psi^n(x)$  and the corresponding intermediate solution of Eq. (2.22) by  $\Psi^{n+1/2}(x)$ . Since  $H_1$  does not contain any derivatives, the intermediate solution can be explicitly calculated as

$$\Psi^{n+1/2}(x) = \mathcal{O}_{\text{nd}}(H_1) \Psi^n(x) \equiv e^{-i\Delta H_1} \Psi^n(x), \quad (2.24)$$

where we use  $\mathcal{O}_{\text{nd}}(H_1)$  to denote the time evolution operator with respect to  $H_1$  and the suffix “nd”, short for “non-derivative”, to differentiate it from other similar operators we use.

Next step is to perform the time propagation w.r.t.  $H_2$ , by using semi-implicit CN scheme,

$$i \frac{\Psi^{n+1}(x) - \Psi^{n+1/2}(x)}{\Delta} = \frac{1}{2} H_2 \left[ \Psi^{n+1}(x) + \Psi^{n+1/2}(x) \right], \quad (2.25)$$

where the partial derivative  $\partial/\partial t$  is approximated using a two-point formula connecting the values of the wave function in the present and the future time step. On the right-hand side hand, the Laplacian  $\partial^2/\partial x^2$  is written in a semi-implicit form, which averages the wave function over the present and the future time step. This averaging makes the scheme unconditionally stable [31, 32].

The solution of this time-discretized equation with the initial condition  $\Psi^{n+1/2}(x)$ , obtained in the previous substep, yields the wave function  $\Psi^{n+1}(x)$  at time  $t_{n+1}$ . Formally, it is solved as

$$\Psi^{n+1}(x) = \mathcal{O}_{\text{CN}}(H_2) \Psi^{n+1/2}(x) \equiv \frac{1 - i\Delta H_2/2}{1 + i\Delta H_2/2} \Psi^{n+1/2}(x). \quad (2.26)$$

Similarly to the  $\mathcal{O}_{\text{nd}}(H_1)$  operator, the  $\mathcal{O}_{\text{CN}}(H_2)$  operation denotes the time evolution operator w.r.t.  $H_2$ . To actually solve Eq. (2.25), we also have to discretize spatial coordinate and introduce a mesh with  $N_x$  equidistant points, separated by a spacing  $h$ . Mesh points are defined by  $x_i = -L + ih$ , where  $L = N_x h/2$  determines the size of the spatial interval  $[-L, L]$  considered to be relevant, i.e., such that the wave function is sufficiently small outside of it and we can safely set it to zero in our calculations. If we denote the wave function at time  $t_n$  and position  $x_i$  by  $\Psi_i^n$ , Eq. (2.25) can be written as

$$i \frac{\Psi_i^{n+1} - \Psi_i^{n+1/2}}{\Delta} = \frac{\Psi_{i+1}^{n+1} - 2\Psi_i^{n+1} + \Psi_{i-1}^{n+1} + \Psi_{i+1}^{n+1/2} - 2\Psi_i^{n+1/2} + \Psi_{i-1}^{n+1/2}}{2h^2}, \quad (2.27)$$

where each Laplacian is approximated using a three-point formula. The associated discretization error is of second order in  $h$ , the same as in time step  $\Delta$ , as discussed previously. This makes discretization errors negligible for small enough discretization steps. Equation (2.27) can be written as a series of tridiagonal equations in unknown wave function values at  $t_{n+1}$ ,

$$A_i^- \Psi_{i-1}^{n+1} + A_i^0 \Psi_i^{n+1} + A_i^+ \Psi_{i+1}^{n+1} = B_i, \quad (2.28)$$

where the coefficients

$$A_i^- = A_i^+ = -i \frac{\Delta}{2h^2}, \quad (2.29)$$

$$A_i^0 = 1 + i \frac{\Delta}{h^2}, \quad (2.30)$$

$$B_i = i \frac{\Delta}{2h^2} \left( \Psi_{i+1}^{n+1/2} - 2\Psi_i^{n+1/2} + \Psi_{i-1}^{n+1/2} \right) + \Psi_i^{n+1/2}, \quad (2.31)$$

are all known. To solve Eq. (2.28), we use forward recursion relation, and assume the solution in a form

$$\Psi_{i+1}^{n+1} = \alpha_i \Psi_i^{n+1} + \beta_i, \quad (2.32)$$

where coefficients  $\alpha_i$  and  $\beta_i$  need to be determined. If we substitute Eq. (2.32) into Eq. (2.28) we obtain

$$\Psi_i^{n+1} = \gamma_i \left( A_i^- \Psi_{i-1}^{n+1} + A_i^+ \beta_i - B_i \right), \quad (2.33)$$

where  $\gamma_i = -1/(A_i^0 + A_i^+ \alpha_i)$ . We can now use the backward recursion relation to obtain the coefficients  $\alpha_i$  and  $\beta_i$ ,

$$\alpha_i = \gamma_{i+1} A_{i+1}^-, \quad \beta_i = \gamma_{i+1} (A_{i+1}^+ \beta_{i+1} - B_{i+1}). \quad (2.34)$$

The strategy to apply the aforementioned recursion relations from Eqs. (2.32)–(2.34) is to perform a backward sweep of the mesh to determine  $\alpha_i$  and  $\beta_i$ , for  $i$  running from  $N_x - 2$  to 0. The initial values on endpoints are chosen to be  $\alpha_{N_x-1} = \beta_{N_x-1} = 0$ , which ensures that the value of the wave function at the border is always zero. With the coefficients  $\alpha_i$ ,  $\beta_i$  and  $\gamma_i$  determined, and the boundary condition  $\Psi_0^{n+1} = 0$  fixed to provide that the wave function is equal to zero at  $x_0 = -L$ , we can determine the solution for the entire space range at time  $t_{n+1}$  using a forward sweep from Eq. (2.32). This completes the single iteration and procedure continues on until end of propagation is reached.

Note that, apart from unconditional stability and small discretization errors, the above scheme has one additional advantage, namely it conserves the normalization of the wave function. This can be monitored throughout the calculation and used as a criterion to assess the quality and validity of the obtained solution.

The method illustrated above in 1D can be straightforwardly generalized to solving 3D problems. For a full 3D dipolar GPE (2.9), discretization of time is performed in the same way, while spatial coordinates are discretized analogously with  $N_x$ ,  $N_y$  and  $N_z$  mesh points in the corresponding direction, and with spacings  $h_x$ ,  $h_y$  and  $h_z$ , respectively. Time propagation in each iteration is now divided into four substeps, corresponding to the following parts the Hamiltonian:

$$H_1 = \gamma^2 x^2 + \nu^2 y^2 + \lambda^2 z^2 + g N_{\text{at}} |\Psi(\mathbf{r}; t)|^2 + g_{\text{dd}} N_{\text{at}} \int d\mathbf{r}' V_{\text{dd}}(\mathbf{r} - \mathbf{r}') |\Psi(\mathbf{r}'; t)|^2, \quad (2.35)$$

$$H_2 = -\frac{\partial^2}{\partial x^2}, \quad H_3 = -\frac{\partial^2}{\partial y^2}, \quad H_4 = -\frac{\partial^2}{\partial z^2}. \quad (2.36)$$

To propagate the wave function w.r.t.  $H_1$ , we need to calculate the expression (2.35) at each mesh point. This requires evaluation of the dipolar term given by the convolution integral, which can be written and calculated as

$$\int d\mathbf{r}' V_{\text{dd}}(\mathbf{r} - \mathbf{r}') |\Psi(\mathbf{r}'; t)|^2 = \mathcal{F}^{-1} \left[ \mathcal{F}[V_{\text{dd}}(\mathbf{r})] \mathcal{F}[|\Psi(\mathbf{r}; t)|^2] \right], \quad (2.37)$$

where  $\mathcal{F}$  represents Fourier transform and  $\mathcal{F}^{-1}$  its inverse. These can be computed using the standard Fast Fourier Transform (FFT) algorithm. Note that  $\mathcal{F}[V_{\text{dd}}(\mathbf{r})]$  does not change between time steps, and can be computed analytically, yielding

$$\mathcal{F}[V_{\text{dd}}(\mathbf{r})] = 3 \cos^2 \theta - 1, \quad (2.38)$$

where  $\theta$  is the angle between the orientation of dipoles and vector  $\mathbf{k}$ . As we assume that the dipoles are oriented in  $z$  direction,  $\cos \theta = k_z/k$ . This leaves us with two Fourier transforms in each time step and in particular motivates our interest for implementation of this approach on GPUs.

The above method is applicable to real-time dynamics of the system, given that its initial state is known. However, we can extend the method to also enable calculation of the ground state. This is achieved by switching to *imaginary-time* propagation, when real time  $t$  is replaced by an imaginary quantity  $\tau = it$ . Mathematically, this corresponds to a Wick rotation [33] and yields imaginary-time analogue of GPE,

$$-\frac{\partial}{\partial \tau} \Psi(\mathbf{r}; \tau) = H \Psi(\mathbf{r}; \tau), \quad (2.39)$$

where  $H = H_1 + H_2 + H_3 + H_4$  is the same as above. Although such a Wick rotation is just a formal mathematical transformation of GPE and does not have any physical motivation, time propagation of the above equation leads to a solution of the time-independent GPE (1.3),

$$\Psi_0(\mathbf{r}) = \lim_{\tau \rightarrow \infty} \Psi(\mathbf{r}; \tau), \quad (2.40)$$

provided that the normalization of the wave function is kept equal to unity during the time propagation, as it is not conserved any more. Equation (2.40) follows from the decomposition of the wave function in the eigenbasis of the Hamiltonian for the Schrödinger equation and also holds for GPE [12]. Therefore, to obtain a ground state we start with an arbitrary initial state (usually taken to be a Gaussian) and propagate it in imaginary time until the convergence is achieved. We also note that the ground state of a system described by GPE can always be represented by a real-valued wave function, which can be exploited to obtain the solution faster. Namely, working with real-valued data on a computer is faster than with the complex-valued ones, since all mathematical calculations require less operations. Practically, this works as follows: the time propagation w.r.t.  $H_1$  is performed using a formula analogous to Eq. (2.24),

$$\Psi^{n+1/2} = \mathcal{O}_{\text{nd}}(H_1) \Psi^n \equiv e^{-\Delta H_1} \Psi^n, \quad (2.41)$$

while the coefficients of the CN scheme (2.29)–(2.31) are now real-valued and (in 1D) defined as

$$A_i^- = A_i^+ = \frac{\Delta}{2h^2}, \quad A_i^0 = 1 - \frac{\Delta}{h^2}, \quad B_i = -\frac{\Delta}{2h^2} \left( \Psi_{i+1}^{n+1/2} - 2\Psi_i^{n+1/2} + \Psi_{i-1}^{n+1/2} \right) + \Psi_i^{n+1/2}. \quad (2.42)$$

Here  $\Delta$  represents time step in imaginary time and generalization to 3D is straightforward.

Similar procedure is applied to solve effective GPEs of reduced dimensionality. In 2D, depending on the orientation of dipoles, the Hamiltonian corresponding to GPE (2.11) or (2.14) is written in

the split form  $H = H_1 + H_2 + H_3$ . The non-derivative part  $H_1$  is already expressed in the form suitable for evaluation using FFT, while the Laplacians  $H_2$  and  $H_3$  are addressed using the CN scheme, as described above. Effective 1D GPEs (2.16) and (2.19) are solved in the same way, using the split form of the Hamiltonian  $H = H_1 + H_2$ . Similarly to 2D case, the non-derivative part  $H_1$  is written in  $k$ -space and allows for efficient evaluation using FFT.

With the main numerical method described above, we can define a computer algorithm for real- and imaginary-time propagation in 1D, 2D and 3D. While this is the main goal that will be addressed in subsequent chapters, we first give expressions for calculation of relevant physical quantities in the next section.

## 2.4 Calculation of relevant physical quantities

Numerical methods presented in this thesis enable us to solve different variants of GPE in a number of physically relevant problems. However, in order to study properties of physical systems described by those equations as well as their dynamical evolution, we have to extract values of relevant physical quantities from the current state of the system, i.e., from the wave function in a given time  $t$ . The most useful quantities are chemical potential, energy, and expectation value of the system size. In particular, convergence of these quantities during the imaginary-time propagation is used as a criterion for the convergence of the wave function to the ground state of the system.

Chemical potential  $\mu$  can be defined for stationary states of a BEC and follows from the time-independent GPE (1.3):

$$\begin{aligned} \mu = \frac{1}{2} \int d\mathbf{r} & \left[ |\nabla\Psi(\mathbf{r})|^2 + (\gamma^2 x^2 + \nu^2 y^2 + \lambda^2 z^2) |\Psi(\mathbf{r})|^2 + gN_{\text{at}} |\Psi(\mathbf{r})|^4 \right. \\ & \left. + g_{\text{dd}} N_{\text{at}} \int d\mathbf{r}' V_{\text{dd}}(\mathbf{r} - \mathbf{r}') |\Psi(\mathbf{r}')|^2 |\Psi(\mathbf{r})|^2 \right]. \end{aligned} \quad (2.43)$$

The corresponding energy of the system  $E$  is given by a similar expression in which the interaction energy terms contain an additional factor of 1/2:

$$\begin{aligned} E = \frac{1}{2} \int d\mathbf{r} & \left[ |\nabla\Psi(\mathbf{r})|^2 + (\gamma^2 x^2 + \nu^2 y^2 + \lambda^2 z^2) |\Psi(\mathbf{r})|^2 + \frac{1}{2} gN_{\text{at}} |\Psi(\mathbf{r})|^4 \right. \\ & \left. + \frac{1}{2} g_{\text{dd}} N_{\text{at}} \int d\mathbf{r}' V_{\text{dd}}(\mathbf{r} - \mathbf{r}') |\Psi(\mathbf{r}')|^2 |\Psi(\mathbf{r})|^2 \right]. \end{aligned} \quad (2.44)$$

The above equations are written for the rescaling of time  $\tilde{t} = 2\tilde{\omega}t$ , used throughout previous section. If other standard rescaling of time is used,  $\tilde{t} = \tilde{\omega}t$ , the factor of 1/2 in front of both Eqs. (2.43) and (2.44) should be removed. Note that the dimensionless values of  $\mu$  and  $E$  can be converted back to physical units by multiplying them with  $\hbar\tilde{\omega}$ .

Similarly, we can define chemical potential and energy for stationary states of effectively 2D and 1D systems. For 2D case when dipoles are oriented perpendicular to the plane, described by GPE (2.11), chemical potential is given by

$$\begin{aligned} \mu = \frac{1}{2} \int d\mathbf{r}_{2\text{D}} & \left[ |\nabla_{2\text{D}}\Psi_{2\text{D}}(\mathbf{r}_{2\text{D}})|^2 + (\gamma^2 x^2 + \nu^2 y^2) |\Psi_{2\text{D}}(\mathbf{r}_{2\text{D}})|^2 + g_{2\text{D}} N_{\text{at}} |\Psi_{2\text{D}}(\mathbf{r}_{2\text{D}})|^4 \right. \\ & \left. + g_{\text{dd},2\text{D}} N_{\text{at}} \int \frac{d\mathbf{k}_{2\text{D}}}{(2\pi)^2} e^{-i\mathbf{k}_{2\text{D}} \cdot \mathbf{r}_{2\text{D}}} \tilde{n}_{2\text{D}}(\mathbf{k}_{2\text{D}}) h_{2\text{D}} \left( \frac{k_{2\text{D}}}{\sqrt{2\lambda}} \right) |\Psi_{2\text{D}}(\mathbf{r}_{2\text{D}})|^2 \right], \end{aligned} \quad (2.45)$$

while the energy can be calculated as

$$E = \frac{1}{2} \int d\mathbf{r}_{2D} \left[ |\nabla_{2D} \Psi_{2D}(\mathbf{r}_{2D})|^2 + (\gamma^2 x^2 + \nu^2 y^2) |\Psi_{2D}(\mathbf{r}_{2D})|^2 + \frac{1}{2} g_{2D} N_{\text{at}} |\Psi_{2D}(\mathbf{r}_{2D})|^4 \right. \\ \left. + \frac{1}{2} g_{\text{dd},2D} N_{\text{at}} \int \frac{d\mathbf{k}_{2D}}{(2\pi)^2} e^{-i\mathbf{k}_{2D} \cdot \mathbf{r}_{2D}} \tilde{n}_{2D}(\mathbf{k}_{2D}) h_{2D} \left( \frac{k_{2D}}{\sqrt{2\lambda}} \right) |\Psi_{2D}(\mathbf{r}_{2D})|^2 \right]. \quad (2.46)$$

In the second 2D case, when dipoles lie in the plane, the chemical potential and energy are given by

$$\mu = \frac{1}{2} \int d\mathbf{r}_{2D} \left[ |\nabla_{2D} \Psi_{2D}(\mathbf{r}_{2D})|^2 + (\gamma^2 x^2 + \lambda^2 z^2) |\Psi_{2D}(\mathbf{r}_{2D})|^2 + g_{2D} N_{\text{at}} |\Psi_{2D}(\mathbf{r}_{2D})|^4 \right. \\ \left. + g_{\text{dd},2D} N_{\text{at}} \int \frac{d\mathbf{k}_{2D}}{(2\pi)^2} e^{-i\mathbf{k}_{2D} \cdot \mathbf{r}_{2D}} \tilde{n}_{2D}(\mathbf{k}_{2D}) j_{2D} \left( \frac{k_{2D}}{\sqrt{2\nu}}, k_z \right) |\Psi_{2D}(\mathbf{r}_{2D})|^2 \right], \quad (2.47)$$

$$E = \frac{1}{2} \int d\mathbf{r}_{2D} \left[ |\nabla_{2D} \Psi_{2D}(\mathbf{r}_{2D})|^2 + (\gamma^2 x^2 + \lambda^2 z^2) |\Psi_{2D}(\mathbf{r}_{2D})|^2 + \frac{1}{2} g_{2D} N_{\text{at}} |\Psi_{2D}(\mathbf{r}_{2D})|^4 \right. \\ \left. + \frac{1}{2} g_{\text{dd},2D} N_{\text{at}} \int \frac{d\mathbf{k}_{2D}}{(2\pi)^2} e^{-i\mathbf{k}_{2D} \cdot \mathbf{r}_{2D}} \tilde{n}_{2D}(\mathbf{k}_{2D}) j_{2D} \left( \frac{k_{2D}}{\sqrt{2\nu}}, k_z \right) |\Psi_{2D}(\mathbf{r}_{2D})|^2 \right]. \quad (2.48)$$

The same applies to 1D systems described by GPE (2.16), where in the case when dipoles are orthogonal to the trap axis, the chemical potential and energy can be expressed by

$$\mu = \frac{1}{2} \int_{-\infty}^{\infty} dx \left[ \left| \frac{d\Psi_{1D}(x)}{dx} \right|^2 + \gamma^2 x^2 |\Psi_{1D}(x)|^2 + g_{1D} N_{\text{at}} |\Psi_{1D}(x)|^4 \right. \\ \left. + g_{\text{dd},1D} N_{\text{at}} \int_{-\infty}^{\infty} \frac{dk_x}{2\pi} e^{-ik_x x} \tilde{n}_{1D}(k_x) j_{1D} \left( \frac{k_x}{\sqrt{2\nu}} \right) |\Psi_{1D}(x)|^2 \right], \quad (2.49)$$

$$E = \frac{1}{2} \int_{-\infty}^{\infty} dx \left[ \left| \frac{d\Psi_{1D}(x)}{dx} \right|^2 + \gamma^2 x^2 |\Psi_{1D}(x)|^2 + \frac{1}{2} g_{1D} N_{\text{at}} |\Psi_{1D}(x)|^4 \right. \\ \left. + \frac{1}{2} g_{\text{dd},1D} N_{\text{at}} \int_{-\infty}^{\infty} \frac{dk_x}{2\pi} e^{-ik_x x} \tilde{n}_{1D}(k_x) j_{1D} \left( \frac{k_x}{\sqrt{2\nu}} \right) |\Psi_{1D}(x)|^2 \right], \quad (2.50)$$

while in the case when dipoles are parallel to the trap axis these quantities are defined by

$$\mu = \frac{1}{2} \int_{-\infty}^{\infty} dz \left[ \left| \frac{d\Psi_{1D}(z)}{dz} \right|^2 + \lambda^2 z^2 |\Psi_{1D}(z)|^2 + g_{1D} N_{\text{at}} |\Psi_{1D}(z)|^4 \right. \\ \left. + g_{\text{dd},1D} N_{\text{at}} \int_{-\infty}^{\infty} \frac{dk_z}{2\pi} e^{-ik_z z} \tilde{n}_{1D}(k_z) h_{1D} \left( \frac{k_z}{\sqrt{2\gamma}} \right) |\Psi_{1D}(z)|^2 \right], \quad (2.51)$$

$$E = \frac{1}{2} \int_{-\infty}^{\infty} dz \left[ \left| \frac{d\Psi_{1D}(z)}{dz} \right|^2 + \lambda^2 z^2 |\Psi_{1D}(z)|^2 + \frac{1}{2} g_{1D} N_{\text{at}} |\Psi_{1D}(z)|^4 \right. \\ \left. + \frac{1}{2} g_{\text{dd},1D} N_{\text{at}} \int_{-\infty}^{\infty} \frac{dk_z}{2\pi} e^{-ik_z z} \tilde{n}_{1D}(k_z) h_{1D} \left( \frac{k_z}{\sqrt{2\gamma}} \right) |\Psi_{1D}(z)|^2 \right]. \quad (2.52)$$

Note that both of the above quantities can be generalized and calculated also during real-time propagation, however only time-dependent energy is physically relevant.

Another set of useful quantities is related to the size of the system, which can be calculated for the ground state as well as during dynamical evolution of the system. Typical measure of system's size in a given direction is expectation value of the coordinate square, such as

$$\langle x^2 \rangle = \int d\mathbf{r} x^2 |\Psi(\mathbf{r}; t)|^2, \quad \langle y^2 \rangle = \int d\mathbf{r} y^2 |\Psi(\mathbf{r}; t)|^2, \quad \langle z^2 \rangle = \int d\mathbf{r} z^2 |\Psi(\mathbf{r}; t)|^2, \quad (2.53)$$

and the system size is then estimated as a *root-mean-square* (RMS) of the corresponding expectation value

$$x_{\text{rms}} = \sqrt{\langle x^2 \rangle}, \quad y_{\text{rms}} = \sqrt{\langle y^2 \rangle}, \quad z_{\text{rms}} = \sqrt{\langle z^2 \rangle}. \quad (2.54)$$

The size of the whole system is usually estimated by the quadratic mean of coordinate RMS sizes,

$$r_{\text{rms}} = \sqrt{\langle x^2 \rangle + \langle y^2 \rangle + \langle z^2 \rangle}. \quad (2.55)$$

Note that all RMS sizes calculated in this way are dimensionless and can be converted to physical units by multiplying them with the harmonic oscillator length  $l$ .

We can define analogous quantities for systems of reduced dimensionality where, for example, coordinate RMS sizes of 2D systems can be expressed by

$$\langle x^2 \rangle = \int d\mathbf{r}_{2\text{D}} x^2 |\Psi_{2\text{D}}(\mathbf{r}_{2\text{D}}; t)|^2, \quad \langle y^2 \rangle = \int d\mathbf{r}_{2\text{D}} y^2 |\Psi_{2\text{D}}(\mathbf{r}_{2\text{D}}; t)|^2, \quad (2.56)$$

while for 1D systems the RMS size is given by

$$\langle x^2 \rangle = \int_{-\infty}^{\infty} dx x^2 |\Psi_{1\text{D}}(x; t)|^2. \quad (2.57)$$

Of course, the above equations depend on the 2D or 1D system geometry, i.e., on the coordinates used to describe the system.

Finally, in order to visualize the system, one can consider the appropriate density distribution. For 3D systems, the density is given by

$$n(\mathbf{r}; t) = |\Psi(\mathbf{r}; t)|^2, \quad (2.58)$$

and analogously for 2D and 1D systems. We can also consider effective, lower-dimensional density projections, that can be more easily visualized. For example, for 3D systems we can define the effective 2D and 1D density by integrating out the full 3D density over some of the coordinates,

$$n_{2\text{D}}(\mathbf{r}_{2\text{D}}; t) = \int_{-\infty}^{\infty} dz |\Psi(\mathbf{r}; t)|^2, \quad n_{1\text{D}}(x; t) = \int d\mathbf{r}_{2\text{D}} |\Psi(\mathbf{r}; t)|^2. \quad (2.59)$$

Similar density projections can be defined for 2D systems as well. All densities and density projections calculated in this way are dimensionless and can be converted to physical units by multiplying them with the appropriate unit ( $l^{-3}$  in 3D,  $l^{-2}$  in 2D and  $l^{-1}$  in 1D).

## Chapter 3

# Algorithm for shared memory systems

Numerical analysis from Chapter 2 outlines the algorithm we present in this chapter. Similarly to the numerical analysis, the algorithm is based on Refs. [12, 13, 16]. Here we present an improved serial and a new, parallelized version for shared memory computer systems.

Initial algorithm for solving the GPE with contact interaction term using Crank-Nicolson method was published in Ref. [12], which also includes a serial implementation in Fortran. This was followed by the serial and parallel implementations written in C [13]. An extension of the algorithm to include the dipolar interaction term was published in Ref. [16]. The authors provided two serial implementations, in C and Fortran. Their algorithm was used as a starting point for all algorithms and implementations presented in this thesis.

In Section 3.1 we will give an overview of the baseline algorithm for dipolar GPE [16] and discuss options for its parallelization on shared memory systems, followed by the description of the improved serial implementation in Section 3.2. Typical numerical simulations that make use of presented serial implementations of the algorithm are extremely computationally demanding, and therefore the parallelization approach was an obvious route forward, details of which are given in Section 3.3.

### 3.1 Description of the algorithm

In this section we give a brief description of the algorithm for the application of numerical methods presented in Chapter 2, which is instructive in understanding how to implement it on a shared memory computer.

We start with the description of the discretization scheme used by the algorithm. Let  $N_x$ ,  $N_y$  and  $N_z$  be the number of points in each direction, corresponding to the  $x$ ,  $y$  and  $z$  directions in one, two or three dimensions (1D, 2D and 3D, respectively). In other words, the number of points in 1D is defined by  $N_x$ , in 2D by  $N_x$  and  $N_y$ , and in 3D by  $N_x$ ,  $N_y$  and  $N_z$ . Discretization points in each direction are equidistant, with their spacing defined in dimensionless units by variables  $dx$ ,  $dy$  and  $dz$ , illustrated in Figure 3.1.




Figure 3.1: Example of 1D, 2D and 3D meshes with  $N_x$ ,  $N_y$  and  $N_z$  discretization points in  $x$ ,  $y$  and  $z$  direction, respectively, and the corresponding spacings  $dx$ ,  $dy$  and  $dz$ .

The mesh determines how the space is discretized, but the actual values of the wave function, which is also discretized by the mesh, need to be stored in a separate array. For convenience, we refer to this array as `psi` variable. This variable is defined as a vector of size  $N_x$  in 1D, as  $N_x \times N_y$  matrix in 2D and as  $N_x \times N_y \times N_z$  tensor in 3D. In real-time propagation, the wave function values are complex, while in imaginary-time propagation they are purely real, which determines the type of `psi` variable.

Time is discretized according to  $t_n = t_0 + n\Delta$ , with the time step  $\Delta$  stored in a variable `dt`. Taking into account the time step and the total amount of physical time we wish to simulate, we determine the number of iterations required ( $n$ ). The main part of the algorithm is a loop where each iteration corresponds to one time step of the propagation. Inside the loop, we update the wave function by propagating it in time with respect to different parts of the Hamiltonian, according to the split-step scheme described in Chapter 2. This is illustrated in Figure 3.2.




Figure 3.2: Main loop of the real-time propagation algorithm in 3D.

In each step of the main loop, wave function values are computed along the discretization mesh in several smaller steps. First such substep is the calculation of the dipolar interaction term. We have two further substeps in 1D, pertaining to Eq. (2.22) w.r.t.  $H_1$  and Eq. (2.23) w.r.t.  $H_2$ . In 2D and 3D we follow a similar procedure and add further substeps for  $H_3$  and  $H_4$ . Each substep produces intermediate values of the wave function that are used in the next substep. The order of the substeps dealing with parts of the Hamiltonian with spatial derivatives ( $H_2$ ,  $H_3$  and  $H_4$ ) can be arbitrary, a feature useful in hybrid and distributed memory algorithms. Propagation of the wave function w.r.t.  $H_1$  relies on the availability of the dipolar interaction term value, which is computed in the first substep, before we start with the update of the wave function.

The dipolar interaction term is computed according to Eq. (2.37), using the discrete Fourier transform (DFT). DFT of a sequence is commonly implemented by relying on the *Fast Fourier transform* (FFT) algorithm via an external library. From Eq. (2.37) it follows that we first need a



DFT of a sequence consisting of absolute squares of wave function values. The resulting complex-valued sequence is then multiplied by the Fourier transform of the dipolar potential. This transform is a known function (2.38), precomputed for a chosen spatial mesh and stored in an array. Next, by performing the inverse DFT we get the dipolar interaction term stored in the resulting array, which we then use to propagate the wave function w.r.t.  $H_1$ . In 2D and 3D, we need multidimensional DFT, which can be computed by the composition of a sequence of 1D DFTs along each direction.

The substep in which the wave function is propagated w.r.t. the Hamiltonian part without spatial derivatives ( $H_1$ ) proceeds according to Eq. (2.24), by employing further nested loops over mesh points. Part of the  $H_1$  is the trap potential  $V(\mathbf{r})$ , which, like Fourier transform of the dipolar potential, can be computed only once during the initialization, and then reused. We assume that the trap potential is static in all algorithms described in this thesis. In case of a time-varying trap potential, the algorithm can be modified to update the trap potential before this substep, i.e., at the beginning of each iteration.

The remaining substep, or substeps in 2D and 3D, propagates the wave function by relying on the Crank-Nicolson scheme. We need to perform a backward sweep of the mesh in each direction to determine corresponding coefficients  $\alpha$ ,  $\beta$  and  $\gamma$  for each direction, according to Eqs. (2.32)–(2.34), followed by a forward sweep to determine the solution for the entire space range. Since the coefficients  $\alpha$  and  $\gamma$  for each direction are independent of the current state of the wave function, they can be computed beforehand, during initialization.

The algorithm for imaginary-time propagation is similar to the one discussed above, with the addition of the normalization step at the end of each iteration of the main loop, as illustrated in Figure 3.3. To normalize the wave function, we first need to compute its norm, according to Eq. (2.3), and then to divide the wave function values with that norm.




Figure 3.3: Main loop of the imaginary-time propagation algorithm in 3D, which introduces an additional step of normalization.

In Section 2.4, we introduced other physical quantities of interest that are computed from the wave function, namely chemical potential, energy and the RMS of system’s physical extent (size), which give us more insight into the state of the simulation and allows us to observe and measure time evolution of a BEC. These quantities do not have to be computed in each iteration. Instead, they can be calculated periodically after a predefined number of iterations.

We see that the algorithm is computationally very demanding, as it requires multiple passes over the entire mesh in each iteration of the main loop. In 2D and 3D, the total number of discretized points sharply increases, making this problem even greater. As we have identified above, several variables can be moved out of the main loop and computed only once, however the remaining calculations are still very demanding, hence the need for parallelization.

The described substeps cannot be executed concurrently, as the intermediate values they produce are used in subsequent substeps. Instead, parallelization of the algorithm can only be achieved by focusing on the mesh loops inside substeps. This is an obvious choice for shared memory systems, as all data stored in the memory is local and accessible to all participating processes, removing the

need for data transfers between processes. The loops themselves can easily be parallelized if they do not contain recursive relations. This is true for the substeps dealing with the computation of the dipolar term and propagation w.r.t.  $H_1$ , as well as periodical computation of chemical potential, energy and RMS of BEC's size. Substep involving  $H_2$  cannot be easily parallelized in 1D, as it has recursive relations in both backward and forward sweeps. In general, recursive relations can be parallelized using the *scan* algorithm [34] (also known as the generalization of *prefix sum* algorithm [35]), however the implementation complexity of such an algorithm made us discard this approach. In 2D and 3D, we can exploit the fact that the recursive relations appear only in the innermost loops of the substeps involving  $H_2$ ,  $H_3$  and  $H_4$ . Thus, we can achieve parallelization by dividing the work among the processes at the level of the outermost loop.

Given the algorithm and parallelization strategy, in next two sections we present serial and parallel implementations.

## 3.2 Serial implementation

Basis for the programs presented in this section was originally developed in Fortran and published in 2009 in Ref. [12], however without considering the dipolar interaction term. Fortran and C implementations that take into account both contact and dipolar interaction term have been published in 2015 in Ref. [16] as package *DBEC-GP*, which can be downloaded from the journal repository [36] or alternatively from the local repository [37]. This package consists of 10 Fortran 90/95 programs, as well as 10 corresponding C programs. For each programming language, there is one program for solving full GPE in 3D, two programs for solving the effective 2D GPE and two for solving the effective 1D GPE, implemented separately for imaginary- and real-time propagation. Here we present improved versions of the latter C programs. Note that these programs may not be suitable for long-running simulations, due to their serial nature. They are superseded by all other programs presented in this thesis, which vastly out-perform programs presented in this section. Another issue with serial and other shared memory programs is their inherent limitation on mesh sizes due to memory constraints of a single computing node. This will be addressed in Chapter 6.

Implementation-wise, the main difference between real- and imaginary-time programs is that the wave function is complex in real-time propagation, and real in imaginary-time. Real values are stored in double precision, as single precision type is inadequate for wave function values since the lower precision in computation leads to faster accumulation of numerical errors. In C-based implementations, complex values are stored in `double complex` type, which is part of the C99 standard [38] of the C language. The implication of using complex numbers is that real-time propagation is more demanding both in terms of memory and computation. However, they have a different purpose: imaginary-time versions of programs are used for calculation of the stationary (ground) state of a physical system, which is then typically used as the initial value of the wave function for real-time propagation.

All programs share the similar structure and reuse common routines whenever possible. Here we present the general flow of implementation, with examples from 1D, 2D and 3D programs.

Before the main time-propagation loop can begin, we need to allocate space for the variables, and set them to their initial values. Matrices and tensors used in 2D and 3D programs are allocated as contiguous arrays of required type (real or complex), with a pointer scheme that allows access to

elements in 2D or 3D, see Listing 3.1. This allows us to access elements of matrices/tensors in a more intuitive way, e.g., `psi[i][j][k]`, as opposed to computing the 1D index every time (for example, `psi[i * Ny * Nz + j * Nz + k]`). The underlying flat allocation scheme is not required, but has turned out to be very useful during the implementation of GPGPU and heterogeneous programs, described in Chapters 4 and 5.

Listing 3.1: Allocation of a real-valued vector, matrix and tensor. Allocation of complex-valued arrays is analogous.

---

```

// Double vector allocation
double *alloc_double_vector(long Nx) {
    double *vector;

    if((vector = (double *) malloc((size_t) (Nx * sizeof(double)))) == NULL) {
        fprintf(stderr, "Failed to allocate memory for the vector.\n"); exit(EXIT_FAILURE);
    }
    return vector;
}

// Double matrix allocation
double **alloc_double_matrix(long Nx, long Ny) {
    long i;
    double **matrix;

    if((matrix = (double **) malloc((size_t) (Nx * sizeof(double *)))) == NULL) {
        fprintf(stderr, "Failed to allocate memory for the matrix.\n"); exit(EXIT_FAILURE);
    }
    if((matrix[0] = (double *) malloc((size_t) (Nx * Ny * sizeof(double)))) == NULL) {
        fprintf(stderr, "Failed to allocate memory for the matrix.\n"); exit(EXIT_FAILURE);
    }
    for(i = 1; i < Nx; i++)
        matrix[i] = matrix[i - 1] + Ny;

    return matrix;
}

// Double tensor allocation
double ***alloc_double_tensor(long Nx, long Ny, long Nz) {
    long i, j;
    double ***tensor;

    if((tensor = (double ***) malloc((size_t) (Nx * sizeof(double **)))) == NULL) {
        fprintf(stderr, "Failed to allocate memory for the tensor.\n"); exit(EXIT_FAILURE);
    }
    if((tensor[0] = (double **) malloc((size_t) (Nx * Ny * sizeof(double *)))) == NULL) {
        fprintf(stderr, "Failed to allocate memory for the tensor.\n"); exit(EXIT_FAILURE);
    }
    if((tensor[0][0] = (double *) malloc((size_t) (Nx * Ny * Nz * sizeof(double)))) == NULL){
        fprintf(stderr, "Failed to allocate memory for the tensor.\n"); exit(EXIT_FAILURE);
    }
    for(j = 1; j < Ny; j++)
        tensor[0][j] = tensor[0][j-1] + Nz;
    for(i = 1; i < Nx; i++) {
        tensor[i] = tensor[i - 1] + Ny;
        tensor[i][0] = tensor[i - 1][0] + Ny * Nz;
        for(j = 1; j < Ny; j++)
            tensor[i][j] = tensor[i][j - 1] + Nz;
    }

    return tensor;
}

```

---

Propagation of the wave function w.r.t. different parts of the Hamiltonian has been implemented in separate functions, where function `calcnu` takes into account  $H_1$  part of the Hamiltonian, `calclux` propagates the wave function w.r.t.  $H_2$ , `calcluy` corresponds to propagation determined by  $H_3$  (in 2D and 3D programs) and `calcluz` does the same for  $H_4$  (in 3D programs). Calculation of the dipolar interaction term, upon which function `calcnu` relies, has been implemented separately, in function `calcpsidd2`. Afterwards, the value of the wave function is propagated for each mesh point, as shown in Listing 3.2.

Listing 3.2: Loops are employed to propagate the wave function in each mesh point, as illustrated here on a 1D real-time `calcnu` function.

---

```
void calcnu(double complex *psi, double *psidd2) {
    long i;
    double psi2, psi2lin, psidd2lin, tmp;

    for(i = 0; i < Nx; i++) {
        psi2 = cabs(psi[i]);
        psi2 *= psi2;
        psi2lin = psi2 * g;
        psidd2lin = psidd2[i] * gd;
        tmp = dt * (pot[i] + psi2lin + psidd2lin);
        psi[i] *= cexp(- I * tmp);
    }
}
```

---

The function `calcpsidd2` computes the DFT of the input sequence by means of the FFT algorithm. The authors in Ref. [16] relied on the FFTW library [39] for this task, a well-known and widely used FFT library. In their implementation they have used the ordinary DFT of complex data, which takes a complex-valued input and produces a complex-valued output of identical size, a so-called *C2C* transform. This can be improved by observing the fact that in Eq. (2.37) the input of the transform is  $|\Psi(\mathbf{r})|^2$ , which is purely real, even when  $\Psi(\mathbf{r})$  is complex-valued. DFT of a real-valued input produces an output that has Hermitian symmetry, with one half of the output being the complex conjugate of the other half, making it redundant. Many FFT libraries, FFTW included, are able to exploit this symmetry to provide a special real-to-complex (*R2C*) transform. R2C transform is both faster, due to fewer operations involved, and requires less memory, because only half of the output is retained. To use R2C instead of C2C transforms, we have to ensure that the size of output array is sufficient (e.g.,  $N_x * N_y * (N_z / 2 + 1) * \text{sizeof}(\text{double complex})$  in 3D), and use the appropriate R2C functions from FFTW (see Listing 3.3). Performing any type of transform with FFTW involves creation and execution of a *plan*, which specifies all details of the transform.

Listing 3.3: Real-to-complex and complex-to-real FFTW calls used in `calcpsidd2` function. Shown here is the 1D variant of the function, used in real-time propagation.

---

```
// Here we use in-place FFT, so we need to ensure proper padding of input and output arrays
psidd2 = alloc_double_vector((Nx / 2 + 1) * 2);
psidd2fft = (fftw_complex *) psidd2;

// Create plans during initialization
plan_forward = fftw_plan_dft_r2c_1d(Nx, psidd2, psidd2fft, FFTW_MEASURE);
plan_backward = fftw_plan_dft_c2r_1d(Nx, psidd2fft, psidd2, FFTW_MEASURE);
```

---

```

void calcpsidd2(double complex *psi, double *psidd2, fftw_complex *psidd2fft) {
    long i;
    double tmp;
    // Prepare input for FFT
    for(i = 0; i < Nx; i++) {
        tmp = cabs(psi[i]);
        psidd2[i] = tmp * tmp;
    }
    // Execute R2C transform
    fftw_execute_dft_r2c(plan_forward, psidd2, psidd2fft);
    // Resulting output array has Nx / 2 + 1 elements
    for(i = 0; i < Nx / 2 + 1; i++) {
        psidd2fft[i][0] *= potdd[i];
        psidd2fft[i][1] *= potdd[i];
    }
    // Execute C2R transform
    fftw_execute_dft_c2r(plan_backward, psidd2fft, psidd2);
    // Rescale the result of FFT
    for(i = 0; i < Nx; i++) {
        psidd2[i] /= Nx;
    }
    // Handle boundary case
    psidd2[Nx - 1] = psidd2[0];
}

```

Functions `calclux`, `calcluy` and `calcluz` perform the time propagation with respect to parts of the Hamiltonian involving spatial partial derivatives by solving the tridiagonal system that arises from Crank-Nicolson method. Implementation of forward recursion and backward substitution is shown in Listing 3.4.

Listing 3.4: 3D real-time functions `calclux`, `calcluy` and `calcluz`.

```

// Time propagation with respect to H2 (x-part of the Laplacian).
void calclux(double complex ***psi, double complex *cbeta) {
    long i, j, k;
    double complex c;

    for(j = 0; j < Ny; j++) {
        for(k = 0; k < Nz; k++) {
            cbeta[Nx - 2] = psi[Nx - 1][j][k];
            for(i = Nx - 2; i > 0; i--) {
                c = -Ax * psi[i + 1][j][k] + AxOr * psi[i][j][k] - Ax * psi[i - 1][j][k];
                cbeta[i - 1] = c * (Ax * cbeta[i] - c);
            }
            psi[0][j][k] = 0.;
            for(i = 0; i < Nx - 2; i++) {
                psi[i + 1][j][k] = calphax[i] * psi[i][j][k] + cbeta[i];
            }
            psi[Nx - 1][j][k] = 0.;
        }
    }
}

// Time propagation with respect to H3 (y-part of the Laplacian).
void calcluy(double complex ***psi, double complex *cbeta) {
    long i, j, k;
    double complex c;

    for(i = 0; i < Nx; i++) {
        for(k = 0; k < Nz; k++) {
            cbeta[Ny - 2] = psi[i][Ny - 1][k];

```

```

for(j = Ny - 2; j > 0; j --) {
    c = - Ay * psi[i][j + 1][k] + AyOr * psi[i][j][k] - Ay * psi[i][j - 1][k];
    cbeta[j - 1] = cgamma[j] * (Ay * cbeta[j] - c);
}
psi[i][0][k] = 0.;
for(j = 0; j < Ny - 2; j ++){
    psi[i][j + 1][k] = calphay[j] * psi[i][j][k] + cbeta[j];
}
psi[i][Ny - 1][k] = 0.;
}
}
}
// Time propagation with respect to H4 (z-part of the Laplacian).
void calcluz(double complex ***psi, double complex *cbeta) {
    long i, j, k;
    double complex c;

    for(i = 0; i < Nx; i ++){
        for(j = 0; j < Ny; j ++){
            cbeta[Nz - 2] = psi[i][j][Nz - 1];
            for(k = Nz - 2; k > 0; k --){
                c = - Az * psi[i][j][k + 1] + AzOr * psi[i][j][k] - Az * psi[i][j][k - 1];
                cbeta[k - 1] = cgamma[k] * (Az * cbeta[k] - c);
            }
            psi[i][j][0] = 0.;
            for(k = 0; k < Nz - 2; k ++){
                psi[i][j][k + 1] = calphaz[k] * psi[i][j][k] + cbeta[k];
            }
            psi[i][j][Nz - 1] = 0.;
        }
    }
}

```

In Listing 3.4 we observe the pattern employed in higher-dimensional programs. Depending on the spatial direction of the computation we perform, i.e., the function used, nested loops are reordered appropriately so that the innermost one corresponds to the computation direction. This pattern is used in other functions whenever we need to access elements in a matrix/tensor in a specific spatial direction.

In imaginary-time propagation, the algorithm mandates that at the end of every time step an additional substep, normalization of the wave function, is required. To compute the wave function norm, we have to integrate over all space, according to Eq. (2.3). Numerical integration is implemented in function `simpint` using composite Simpson's rule,

$$\int_a^b dx f(x) \approx \frac{h}{3} \sum_{i=1}^{n/2} \left[ f(x_{2i-2}) + 4f(x_{2i-1}) + f(x_{2i}) \right], \quad (3.1)$$

where  $n$  is the number of subintervals, step size  $h$  is  $h = (b - a)/n$  and  $x_i = a + ih$ . Computing the norm and subsequent normalization of the wave function is illustrated in Listing 3.5. Note that in real-time propagation the norm of the wave function is preserved, and therefore we do not perform the normalization at the end of every time step. However, keeping track of the norm is still useful as a check of correctness, and we calculate it at the regular intervals. The implementation remains mostly the same, with the last step (performing the normalization of the wave function) omitted in real-time propagation.

Listing 3.5: Numerical integration via composite Simpson's rule and its usage in computation of the norm and normalization of the wave function, as used by the 1D imaginary-time programs. The equivalent real-time function is similar, with the last step removed.

---

```
// Numerical integration via composite Simpson's rule
double simpint(double h, double *f, long N) {
    int c;
    long i;
    double sum;

    sum = f[0];
    for (i = 1; i < N - 1; i++) {
        c = 2 + 2 * (i % 2);
        sum += c * f[i];
    }
    sum += f[N - 1];

    return sum * h / 3.;
}

void calcnorm(double *norm, double *psi, double *tmpx) {
    long i;
    double tmp;

    for(i = 0; i < Nx; i++) {
        tmpx[i] = cabs(psi[i]);
        tmpx[i] *= tmpx[i];
    }

    *norm = sqrt(simpint(dx, tmpx, Nx));
    tmp = 1. / *norm;

    for(i = 0; i < Nx; i++) {
        psi[i] *= tmp;
    }
}

```

---

The chemical potential and energy are computed in function `calcmuen`, while RMS size is calculated in function `calcrms`. Chemical potential and energy are computed in the same function because Eqs. (2.43) and (2.44) used to compute these quantities differ only in a coefficient in front of the nonlinear term. Computation of these quantities requires calculation of spatial derivatives of the wave function, which is implemented in function `diff` using Richardson extrapolation formula of the fourth order [40],

$$f'(x) \approx \frac{f(x-2h) - 8f(x-h) + 8f(x+h) - f(x+2h)}{12h}. \quad (3.2)$$

Finally, expressions for the chemical potential and energy are integrated using composite Simpson's rule. The implementation of `calcmuen` with the aforementioned calculations in 1d case is illustrated in Listing 3.6.

Listing 3.6: Calculation of chemical potential and energy in 1D imaginary-time `calcmuen` function. Also shown is the numerical differentiation function, based on Richardson extrapolation formula.

---

```
void calcmuen(double *mu, double *en, double *psi, double *dpsix, double *psidd2, double
    *tmpxi, double *tmpxj) {
    long i;
    double psi2, psi2lin, psidd2lin, dpsid2;

```

---

```

diff(dx, psi, dpsix, Nx);
calcpsidd2(psi, psidd2);

for(i = 0; i < Nx; i ++) {
    psi2 = psi[i] * psi[i];
    psi2lin = psi2 * g;
    psidd2lin = psidd2[i] * gd;
    dpsid2 = dpsix[i] * dpsix[i];
    tmpxi[i] = (pot[i] + psi2lin + psidd2lin) * psi2 + dpsid2;
    tmpxj[i] = (pot[i] + 0.5 * psi2lin + 0.5 * psidd2lin) * psi2 + dpsid2;
}
*mu = simpint(dx, tmpxi, Nx);
*en = simpint(dx, tmpxj, Nx);
}
// Calculation of numerical derivative using Richardson extrapolation formula
void diff(double h, double *f, double *df, long N) {
    long i;

    df[0] = 0.;
    df[1] = (f[2] - f[0]) / (2. * h);

    for(i = 2; i < N - 2; i ++) {
        df[i] = (f[i - 2] - 8. * f[i - 1] + 8. * f[i + 1] - f[i + 2]) / (12. * h);
    }

    df[N - 2] = (f[N - 1] - f[N - 3]) / (2. * h);
    df[N - 1] = 0.;
}

```

In 2D and 3D, partial derivatives need to be calculated for each direction. Previously published implementations [13, 16] allocate separate arrays for each partial derivative, that can consume significant amount of memory if the mesh size is large. By reorganizing the computation of derivatives slightly, we can reuse a single array for all calculations involving partial derivatives. This small change leads to significant reduction in memory usage, especially in 3D programs. For this reason we have relied on this type of memory usage optimizations for all programs presented in this thesis.

Calculation of RMS sizes in function `calcrms` is straightforward, given the simplicity of their definitions (2.53)–(2.57). Relevant portion of the implementation in 1D is shown in Listing 3.7. Implementation in 2D and 3D relies on the concepts we already presented, hence no further details will be given.

Listing 3.7: Computation of the RMS size in 1D in imaginary-time propagation.

```

void calcrms(double *rms, double complex *psi, double *tmpx) {
    long i;
    double psi2;

    for(i = 0; i < Nx; i ++) {
        psi2 = cabs(psi[i]);
        psi2 *= psi2;
        tmpx[i] = x2[i] * psi2;
    }
    *rms = sqrt(simpint(dx, tmpx, Nx));
}

```

In order to use the programs to simulate a given physical system, we need to provide the appro-



appropriate input parameters. The parameters of the programs, e.g., the discretization scheme, number of time steps, trap anisotropy coefficients, coefficients of nonlinear terms, interaction strengths, names of output files and others, are specified in the input file. Input files are passed to the programs as runtime arguments, where they are parsed, and the variables obtained in this way are used to initialize the space mesh, the wave function, trap and dipolar potentials, and Crank-Nicolson coefficients. Listing 3.8 shows how initialization of space mesh and associated wave function, trap potential and Crank-Nicolson coefficients is implemented in 1D. Initialization of the dipolar potential, which is not listed here, has a slightly longer implementation and is available in Ref. [16]. In 2D and 3D this step is performed in an analogous way, and is also not given here.

Listing 3.8: Initialization of space mesh and associated wave function, trap potential and Crank-Nicolson coefficients in 1D real-time programs.

---

```
// Initialization of the space mesh and the initial wave function.
void initpsi(double complex *psi) {
    long i;
    double cpsi = sqrt(sqrt(pi / vgamma));

    for(i = 0; i < Nx; i++) {
        x[i] = (i - Nx/2) * dx; x2[i] = x[i] * x[i];
    }
    for(i = 0; i < Nx; i++) {
        psi[i] = exp(-0.5 * vgamma * x2[i]) / cpsi;
    }
}
// Initialization of the potential.
void initpot() {
    long i;
    double vgamma2 = vgamma * vgamma;

    for(i = 0; i < Nx; i++) {
        pot[i] = 0.5 * par * (vgamma2 * x2[i]);
    }
}
// Crank-Nicolson scheme coefficients generation.
void gencoef() {
    long i;

    Ax0 = 1. + I * dt / dx2;
    AxOr = 1. - I * dt / dx2;
    Ax = -0.5 * I * dt / dx2;

    calphax[Nx - 2] = 0.;
    cgammax[Nx - 2] = -1. / Ax0;
    for (i = Nx - 2; i > 0; i--) {
        calphax[i - 1] = Ax * cgammax[i];
        cgammax[i - 1] = -1. / (Ax0 + Ax * calphax[i - 1]);
    }
}

```

---

With all the relevant variables allocated and initialized to the proper values, the time propagation can now proceed. This is achieved by calling the previously described computation functions in the main loop, as illustrated in Listing 3.9. Note that in the imaginary-time propagation, as mentioned earlier, the normalization has to be performed at the end of every time step. In the current implementation, the number of time iterations is predefined in the input file.

Listing 3.9: Main loop of the time propagation, for 3D imaginary-time programs.

---

```

//Read input, allocate and initialize variables
...
// Main loop of time propagation
for(i = 0; i < Nrun; i ++) {
    calcpsidd2(psi, psidd2, psidd2fft);
    calcnu(psi, psidd2);
    calclux(psi, cbeta);
    calcluy(psi, cbeta);
    calcluz(psi, cbeta);
    // In imaginary-time propagation, we need additional normalization step
    calcnorm(&norm, psi, tmpxi, tmpyi, tmpzi);
    // Optionally, physical quantities can be computed after certain number of steps
    if (i % 1000 == 0) {
        ... // Call 'calcmuen', 'calcrms' or output functions
    }
}
// Write summary file, clean up and exit
...

```

---

After the predefined number of time steps, programs can compute physical quantities like chemical potential or energy, or they can write to a file the current density profile (absolute square of the wave function), which can be used for further analysis and/or visualization. Final density profile of the imaginary-time propagation can be used as the initial value for wave function propagation in real time. 2D and 3D programs also provide integrated (effective) 1D density profiles, obtained by integrating the densities over all spatial variables but one. Furthermore, 3D programs can write integrated 2D density profiles, obtained in a similar way. Output mesh can coincide with the mesh used by the program, or can be coarser, controlled by the `outstpx`, `outstpy` and `outstpz` input parameters. Two examples of such output functions are given in Listing 3.10.

Listing 3.10: Examples of output functions from 2D real-time programs that compute and write density and integrated density to a file.

---

```

void outpsi2xy(double complex **psi, FILE *file) {
    long i, j;

    for(i = 0; i < Nx; i += outstpx) {
        for(j = 0; j < Ny; j += outstpy) {
            fprintf(file, "%81e %81e %81e\n", x[i], y[j], cabs(psi[i][j]) * cabs(psi[i][j]));
        }
        fprintf(file, "\n"); fflush(file);
    }
}

void outdenx(double complex **psi, double *tmp, FILE *file) {
    long i, j;

    for(i = 0; i < Nx; i += outstpx) {
        for(j = 0; j < Ny; j++) {
            tmp[j] = cabs(psi[i][j]) * cabs(psi[i][j]);
        }
        fprintf(file, "%81e %81e\n", x[i], simpint(dy, tmp, Ny)); fflush(file);
    }
}

```

---

Additionally, all programs write the summary output file, which contains the information on the

input parameters, and the computed norm, chemical potential and energy. A sample summary file is given in Listing 3.11.

Listing 3.11: Example of a summary output file of 3D real-time propagation.

---

```

Real time propagation 3D, OPTION = 2

Number of Atoms N = 5000, Unit of length AHO = 0.00000100 m
Scattering length a = 100.00*a0, Dipolar ADD = 132.70*a0
Nonlinearity G_3D = 332.4918, Strength of DDI GD_3D = 105.33272
Parameters of trap: GAMMA = 0.50, NU = 1.00, LAMBDA = 1.50

# Space Stp: NX = 128, NY = 96, NZ = 80
Space Step: DX = 0.200000, DY = 0.200000, DZ = 0.200000
# Time Stp : NSTP = 0, NPAS = 100, NRUN = 900
Time Step: DT = 0.005000
Dipolar Cut off: R = 10.000

* Change for dynamics: GPAR = 1.500, GDPAR = 1.000 *

-----
      Norm      Chem      Ener/N      <r>      |Psi(0,0,0)|^2
-----
Initial:      1.0000      5.81621      4.33909      2.54499      0.02162
After NPAS iter.: 1.0000      5.81920      4.33911      2.54358      0.02181
After NRUN iter.: 1.0000      6.93923      5.00916      2.63297      0.01448
-----

Clock Time: 290 seconds
CPU Time: 290 seconds

```

---

The 10 improved programs developed here use the same naming convention as the original C programs. They can be downloaded from the local repository [37].

### 3.3 Parallelization on shared memory systems

It can be observed that the implementation described in the previous section is computationally very demanding, especially in 2D and 3D. This results in long execution time when a large discretization scheme is used, which makes the programs less appealing for fine-grained simulations, thus making parallelization a necessity. Parallel implementation on shared memory systems using C language can be achieved in multiple ways, using a wide array of technologies. These range from automatic parallelizations built in compilers [41] to specialized language extensions like Unified Parallel C [42, 43] and Cilk Plus [44]. Given that the parallelization efforts of previous work in Ref. [13] was focused on using OpenMP, which was shown to achieve 80-90% of the ideal speedup, we adopted the same approach.

OpenMP is designed for shared memory systems containing multiple processors/cores. It provides an explicit programming model, based around the concept of using multiple threads to accomplish parallelism. OpenMP is published as an open specification [45], and there are numerous compilers which implement support for it. Usually, a special flag is passed to the compiler (`-fopenmp` in GCC or `-qopenmp` in newer versions of Intel's compiler) to instruct it to parallelize the code using OpenMP.

Parallelization in OpenMP is achieved using the *fork-join* model, where *master* thread spawns (or

*forks*) a number of *slave* threads to perform a task in parallel, at the end of which the slave threads *join* back to the master thread. It relies on compiler directives (called *pragmas*) and runtime library routines to achieve this task. Programmers declare parallel regions of the code by using compiler directives. When a master thread encounters such regions, it creates a team of parallel threads that execute the statements in the region. When threads complete all statements in the parallel region, they synchronize and terminate, leaving only the master thread to continue with the serial execution of the program. This flow is illustrated in Figure 3.4. The number of threads can be different for any parallel region, however in most cases the number of threads is equal to the number of processing cores available in the system. Such is the case with the programs presented here.




Figure 3.4: The fork-join model, used by OpenMP.

To create parallel regions, we must declare them with the compiler directives. Each OpenMP directive in C begins with the keyword `#pragma omp`, followed by a directive name and its optional clauses, and ends with the new line. When dealing with loops, the declaration of a parallel region can be shortened, as demonstrated in Listing 3.12. While the second approach results in a cleaner code that is easier to follow, the first one provides additional flexibility. We use both approaches in our programs, depending on whether we need to perform any additional work or just parallelize the loop at hand. Special care must be taken to ensure that each thread has a copy of intermediate variables it writes, defined in the surrounding parallel region. Otherwise, the compiler would assume that variable is shared among threads, and each thread would write its own value, overwriting the data of other threads and leading to concurrency problems.

Listing 3.12: OpenMP compiler directives.

---

```

#include<omp.h> // Include the appropriate OpenMP header
...
#pragma omp parallel // Spawns a team of parallel threads
{
    #pragma omp for private(i) // Divides loop iterations between the spawned threads
    for(i = 0; i < Nx; i ++){
        tmpx[i] = psi[i] * psi[i];
    }
}

#pragma omp parallel for private(i) // Shorter, combined parallel worksharing construct
for(i = 0; i < Nx; i ++){
    tmpx[i] = psi[i] * psi[i];
}

```

---

Our programs rely on OpenMP to parallelize the most computationally demanding functions, namely `calcnu`, `calclux`, `calcluy`, `calcluz`, `calcpsidd2`, `calcnorm`, `calcmuen` and `calcrms`. Unfortunately, OpenMP currently does not provide any support for parallel input and output (I/O),

leaving this task entirely to the programmer. Performing efficient read/write operations on a single file in multi-threaded environment is a complex task, and because of this no effort was made to improve the I/O functions in our programs, and they remain serial.

Since most of the computation in our programs is done in a series of loops, the parallelization is achieved by assigning the portion of each loop to a different thread. In 1D, that approach would take the form presented in Listing 3.13.

---

Listing 3.13: OpenMP parallelization of 1D loop, demonstrated on real-time `calcrms` function.

---

```
void calcrms(double *rms, double complex *psi, double *tmpx) {
    long i;
    double psi2;

    #pragma omp parallel for private(i)
    for(i = 0; i < Nx; i ++) {
        psi2 = cabs(psi[i]);
        psi2 *= psi2;
        tmpx[i] = x2[i] * psi2;
    }
    *rms = sqrt(simpint_th(dx, tmpx, Nx));
}
```

---

However, we can observe that this technique would not always work in 1D programs, because of the recursive relations in `calclux` function. The first loop requires that `cbeta[i]` is computed before `cbeta[i - 1]`, and similarly with the `psi` in the second loop. The *scan* algorithm could be used to parallelize the loops of this function, however improvements over the serial implementation would be visible only for simulations with very large number of discretization points (over a million), which are rarely needed. In multidimensional programs this approach would be even less effective, as the outer loops are independent and are better suited for parallelization than the innermost loop with recursive relations. Since the serial implementation in 1D can be considered to be sufficiently fast on modern hardware, and due to the complexity that parallelization of `calclux` would entail, the focus here is placed on improving 2D and 3D variants of the programs. 1D programs remain only partially parallelized, with all functions except for `calclux` being threaded.

In multidimensional programs, the outer loops are independent and their work can be divided between threads. Listing 3.14 illustrates how the loop of 2D real-time `calcnu` function is parallelized with OpenMP. The OpenMP runtime would divide the work in such a way that each thread would process approximately  $N_x / \text{nthreads}$  elements (with `nthreads` being the number of threads in the parallel region), where each element contains the whole inner loop. During computation, values of several expressions are stored in temporary variables. As mentioned before, each thread must have a private copy of these variables to avoid overwriting data, so we must specify them in the parallel construct. This concept is used in the parallelization of `calcnu`, `calclux`, `calcluy`, `calcluz` functions and portions of the `calcpsidd2` function.

When working with dynamically allocated arrays, for example when writing temporary results to an array, the variables cannot be made private just by declaring them in the parallel region. Instead, we have to extend the allocated array so that each thread has a separate portion. For instance, instead of allocating a vector of  $N_x$  elements, we allocate a  $\text{nthreads} \times N_x$  matrix, and use each row in a different thread. Even though this concept increases the amount of memory programs use, that increase is not significant as the number of threads is relatively low compared to typical

Listing 3.14: OpenMP parallelization of 2D loop, demonstrated on real-time calcnu function.

---

```

void calcnu(double complex **psi, double **psidd2) {
    long i, j;
    double psi2, psi2lin, psidd2lin, tmp;

    #pragma omp parallel for private(i, j, psi2, psi2lin, psidd2lin, tmp)
    for (i = 0; i < Nx; i++) {
        for (j = 0; j < Ny; j++) {
            psi2 = cabs(psi[i][j]);
            psi2 *= psi2;
            psi2lin = psi2 * g;
            psidd2lin = psidd2[i][j] * gd;
            tmp = dt * (pot[i][j] + psi2lin + psidd2lin);
            psi[i][j] *= cexp(- I * tmp);
        }
    }
}

```

---

vector sizes. An implementation of this concept is shown in Listing 3.15, which demonstrates how `calcnorm` function can be parallelized in 2D. The threads fill out their portion of the `tmpy` temporary array, and perform the integration on their portion of the array. The listing also shows how a single thread in a parallel region can be used to compute the final result. As the variable holding the final result is not declared to be private, its value is visible to all threads, and can be used in the parallel region that follows. Other than `calcnorm` function, parallelization is achieved in this way in `calcmuen` and `calcrms`, so we would not go into their specifics here.

Listing 3.15: Concept used in parallelization of functions requiring temporary storage, demonstrated on imaginary-time `calcnorm` function.

---

```

#pragma omp parallel
#pragma omp master
nthreads = omp_get_num_threads();
...
tmpx = alloc_double_matrix(nthreads, Nx); // Allocate separate array for each thread
tmpy = alloc_double_matrix(nthreads, Ny);
...
void calcnorm(double *norm, double **psi, double **tmpx, double **tmpy) {
    int threadid;
    long i, j;
    double tmp;

    #pragma omp parallel private(threadid, i, j)
    {
        threadid = omp_get_thread_num(); // Get a thread identification number

        #pragma omp for
        for (i = 0; i < Nx; i++) {
            for (j = 0; j < Ny; j++) {
                tmpy[threadid][j] = psi[i][j] * psi[i][j]; // Store data in separate slots
            }
            (*tmpx)[i] = simpint(dy, tmpy[threadid], Ny);
        }
        #pragma omp barrier // Wait for all threads to complete

        #pragma omp single

```

```

{
    *norm = sqrt(simpint(dx, *tmpx, Nx));
    tmp = 1. / *norm; // 'tmp' is a shared variable, visible to all other threads
}

#pragma omp for
for (i = 0; i < Nx; i++) {
    for (j = 0; j < Ny; j++) {
        psi[i][j] *= tmp;
    }
}
}
}

```

These functions rely on numerical integration, which essentially computes a sum of a sequence, an operation also known as *reduction*. Reduction can be handled with OpenMP via a special directive, as we can see in Listing 3.16, which shows how the `simpint` function was parallelized. Note that we need parallelized version of numerical integration only when invoked from outside the threaded loops, as the computation loops themselves are parallelized. Further attempts to perform nested parallelization would not yield any improvement in this case.

Listing 3.16: OpenMP parallelization of numerical integration, using the reduction clause.

```

double simpint_th(double h, double *f, long N) {
    int c;
    long i;
    double sum = f[0];

    #pragma omp parallel for private(i, c) reduction(+:sum)
    for (i = 1; i < N - 1; i++) {
        c = 2 + 2 * (i % 2);
        sum += c * f[i];
    }
    sum += f[N - 1];

    return sum * h / 3.;
}

```

The final portion of the programs that has to be parallelized is the DFT. Instructing FFTW to use OpenMP is a simple task, and requires minimal changes to the code. Before any plan is created or any FFTW function is used, we initialize the threaded FFTW by calling two functions shown in Listing 3.17. The plan creation and execution remain unchanged. Finally, the compiled programs should be linked with the threaded FFTW, by linking to `-lfftw3_omp` in addition to the `-lfftw3`.

Listing 3.17: Initialization of threaded FFTW.

```

#pragma omp parallel
#pragma omp master
nthreads = omp_get_num_threads();
...
fftw_init_threads();
fftw_plan_with_nthreads(nthreads);
// Create plans as usual, e.g., 3D FFT plans
plan_forward = fftw_plan_dft_r2c_3d(Nx, Ny, Nz, **psidd2, psidd2fft, FFT_MEASURE);
plan_backward = fftw_plan_dft_c2r_3d(Nx, Ny, Nz, psidd2fft, **psidd2, FFT_MEASURE);

```

This step completes the OpenMP parallelization of the programs from previous section. The

remaining portions of the code, mainly output functions, are not suitable for parallelization, and were thus not improved.

Using the concepts described, we have parallelized all programs in 1D, 2D and 3D, for both imaginary- and real-time propagation, resulting in 10 parallel programs. The programs form a package named *DBEC-GP-OMP*, a subset of a package *DBEC-GP-OMP-CUDA-MPI*, published in Ref. [18] and available for download from Refs. [46, 37]. Names of the programs are the same as in serial implementation with the suffix “-th” added. Detailed performance review of parallelized programs presented in this section will follow in Chapter 8.



## Chapter 4

# Solving NLSE using GPU accelerators

With the shared memory algorithm described and parallelized in the previous chapter, we can consider using an accelerator processor to achieve even better performance. Use of special-purpose hardware architectures to accelerate computationally demanding applications is becoming standard practice in the scientific community. In this chapter we will present the algorithm and implementation which uses one such accelerator, namely the graphics processing unit (GPU). As their name implies, GPUs have been originally designed for graphics-oriented tasks, however nowadays their enormous processing power can be used for general-purpose computation, giving rise to a practice that is now known as *general-purpose computing on graphics processing units* (GPGPU). They achieve this by relying on the large number of slower low-power processing cores, designed to process many parallel streams of data simultaneously, but are ill-suited for complex processing of few streams of data. This is in contrast to modern CPUs, which have few very fast latency-oriented cores, designed to handle single or few streams of data with complex processing (such as regular desktop applications). As a result, GPUs are better suited to handle tasks that significantly benefit from parallelization, which is why we considered using them for our algorithms.

Using GPUs for general-purpose computing by programming only with graphics primitives would be very time consuming and error-prone, severely limiting the usefulness of GPUs in high-performance computing. Fortunately, GPU vendors have recognized this and developed the interfaces for general-purpose programming on GPUs. There are two competing APIs in this space, CUDA [47] and OpenCL [48]. CUDA (formerly an acronym for Compute Unified Device Architecture) is a mature, proprietary, C, C++ and Fortran API and computing platform created by Nvidia, available since 2007. All Nvidia GPUs from that time onward support CUDA, with each generation of GPUs adding more features. GPUs from other vendors are not supported. It is available for free, but is not open-source.

On the other hand, OpenCL, initially developed by Apple, but since transferred to Khronos group, is the standardized approach to programming on heterogeneous platforms. As such, it is not limited only to GPUs from a single vendor (like CUDA is with Nvidia), rather it supports a plethora of different platforms, from CPUs and GPUs to digital signal processors (DSPs), field-programmable gate arrays (FPGAs) and more. Nvidia GPUs also support OpenCL, however it is implemented on

top of CUDA, and often has worse performance than using pure CUDA [49]. OpenCL specification is based on the subset of C language, and is similar in functionality to CUDA. Being a standard though, OpenCL often lags a bit with adding new functionality when compared to CUDA, which is controlled by a single company. Note that this is usually not a problem, even for advanced usage.

With the functionality being similar in both CUDA and OpenCL, the choice between the two is often based on the available hardware or developer's personal preference. Since the PARADOX supercomputing facility, where all our programs were developed and tested, has Nvidia GPUs, we have decided to use CUDA for implementation of our algorithm.

In the following sections of this chapter we will first describe the CUDA programming and execution model, taking note of the features we used in our programs, and then proceed with the explanation of the changes to the algorithm required by the implementation using CUDA, and finish with the important details of the implementation itself.

## 4.1 CUDA programming and execution model

Hardware architecture of Nvidia (and other) GPUs differs in some significant ways to the x86 CPU architecture. Understanding these differences is the key to exploiting the full potential of GPUs. CPU architectures have historically been focused on increasing performance for serial applications, however this trend has slowed down due to energy consumption and heat dissipation problems that arose with increasing clock frequency. Since then the CPU architectures switched to using multiple processing units, giving rise to the multi-core systems which are in use today. On the other hand, GPUs have been focused with increasing performance of floating-point calculations in response to the demand from video game industry. So instead of investing in advancing control logic and branch prediction that would rival the current CPU designs, GPU vendors have focused on increasing throughput through the use of massive number of threads. This approach has been dubbed *many-core*, signifying the larger number of threads than the multi-core systems have. The software written for this type of computing platform is expected to work with the large number of parallel threads, which the hardware takes advantage of to hide the latency of the memory accesses or arithmetic operations.

Hardware implementation of aforementioned concepts comes in a form of *Streaming Multiprocessor* (SM). The Nvidia GPU architecture is built around an array of these SMs, each executing in parallel with the others. Inside each SM there is a number of *CUDA cores* that execute a sequential thread, various caches, special function units (SFU) that are used for transcendental operations (e.g., sin, cos) and a small amount of shared memory available to all CUDA cores within the SM. A simplified architecture of a CUDA device is shown in Figure 4.1. Modern Nvidia GPUs have a large number of SMs and cores, which, when combined, allow for hundreds or even thousands of simultaneous threads to execute. For example, the previous-generation Tesla M2090 has 512 cores, while the newer Tesla K80 has 4992 CUDA cores. Cores execute in *Single Instruction Multiple Thread* (SIMT) model, similar to the Single Instruction Multiple Data (SIMD) model [50]. This means that all cores within the same group execute the same instruction at the same time, in lock-step fashion. These groups are called *warps*, and in current hardware implementations, each warp has 32 threads. One peculiar behavior of SIMT is in case of conditionals, which SIMT handles by disabling the threads that would not execute the same code, and then rerun the same code with the

previously disabled threads re-enabled. This has negative performance impact, therefore conditional computation that would cause thread divergence should be avoided in GPU code, if possible.

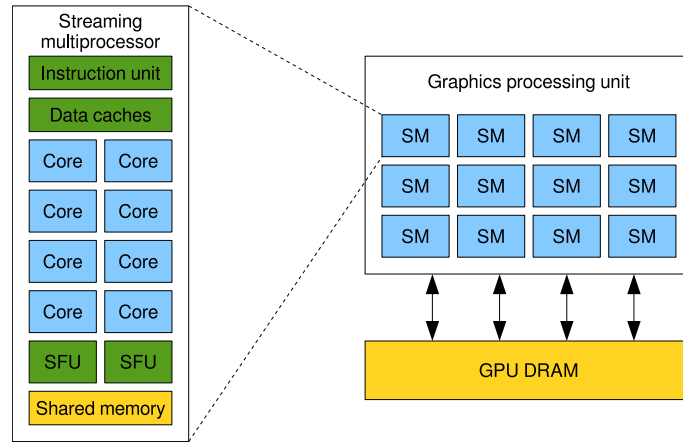


Figure 4.1: Simplified architecture of modern Nvidia GPU.

Within each SM, there is also a small amount of cache memory, named *shared memory*, which is available to all threads within the SM. Because it resides on-chip, shared memory is much faster than main GPU memory, but it is limited in size, with most GPUs having access to 48–96 KB. This type of memory is useful for many scenarios, such as user-managed data caches or high-performance cooperative parallel algorithms (like parallel reductions, which we use in our algorithm for numerical integration).

The main memory of the GPU (named *global memory* in CUDA terminology), its RAM, is separate from the main memory of the host computer on which the GPU is installed. GPU is connected to a host through PCI-Express, and all data transfers between the host and GPU device have to go through this I/O bus. Data is usually transferred between the GPU and host memory using direct memory access (DMA), either in synchronous or asynchronous fashion. Alternatively, GPU may access host’s memory through so called *mapped memory*, where data transfers happen on the fly, however, accessing memory in this way is usually far slower and should thus be used sparingly. Current-generation Nvidia GPUs offer between 2 GB and 24 GB of memory, which is enough for some use-cases, but still noticeably less than amount of RAM memory current high-end computing nodes have. This means that for memory-demanding programs, such as the ones we describe in this thesis, efficient use of GPU memory is paramount.

CUDA execution model states that instructions are issued and executed per warp. This is also true for global memory operations. The CUDA device combines, or *coalesces*, these global memory load/store operations issued by threads of a warp into as few transactions as possible, to maximize memory throughput. Additionally, for data to be read or written by memory transactions it must be naturally aligned, i.e., its first address must be a multiple of their size (32, 64, or 128 bytes). This means that the most favorable memory access pattern is achieved when all threads in a warp access consecutive global memory locations, and in the case of multidimensional data, its innermost dimension must be padded to the nearest natural boundary. The difference between coalesced and non-coalesced (non-sequential or unaligned) access pattern is illustrated in Figure 4.2. Failure to meet the coalesced access pattern will result in more transactions per instruction, reducing the memory throughput and increasing execution time. Memory coalescence is also very important

when dealing with multidimensional data. For accesses to this type of data to be fully coalesced, the number of elements in its innermost dimension must be a multiple of the warp size. In particular, this means that data whose innermost dimension is not a multiple of this size will be accessed much more efficiently if it is actually allocated with a number of elements in its innermost dimension rounded up to the closest multiple of this size.




Figure 4.2: Different memory access patterns on GPU device.

Programming for CUDA devices is done via a C-like language, which is essentially an extension of C language to support GPU computation. CUDA programming model reflects the heterogeneous nature of a system which has a GPU, by making a distinction between the *host* code, which executes on the CPU, and the *device* code, which executes on the GPU. Host code may contain any valid C code, whereas device code is declared by special language extensions. Any given CUDA source file may contain a mixture of host and device code. Because CUDA source file with device code is not valid C, it needs to be compiled with a special compiler, the Nvidia C compiler (`nvcc`). Internally, `nvcc` works by splitting the source files into host and device code, compiling them separately, and then linking them into a final executable. It can also be used to compile and link only the device code, which is useful when integrating CUDA code into existing programs.

A key component of CUDA programming model is the *kernel*, a specially annotated function that executes on the CUDA device. When the kernel is invoked, or *launched*, it is executed by a large number of threads on the device. Kernels do not create or manipulate threads directly, instead the number of threads is defined during kernel launch and device takes control of creating and scheduling them for execution. To differentiate between the threads and to assign different work for them, special variables exist which can only be used inside kernels. Kernel is launched using a special syntax, in the form of `kernel_name<< ... >>( ... )`, where the launch parameters are specified. Listing 4.1 demonstrates the basic definition and execution of a CUDA kernel.

Listing 4.1: Example of a simple CUDA program, with transferring data to/from GPU device, executing a kernel, and allocating/deallocating resources.

```
#include <stdio.h>
// Kernels are functions defined with '__global__' and returning 'void'
__global__ void square(int Nx, double *psi, double *psi2) {
    int i = blockIdx.x * blockDim.x + threadIdx.x; // Compute the thread index in a grid
    if (i < Nx) psi2[i] = psi[i] * psi[i]; // Ensure that we don't go out of bounds
}

void main() {
    int Nx = 1<<20; // 1048576
    double *h_data, *h_data2, *d_data, *d_data2;
    // Allocate data on host and device
    data = (double*) malloc(Nx * sizeof(double));
    data2 = (double*) malloc(Nx * sizeof(double));
```

```

cudaMalloc(&d_data, Nx * sizeof(double));
cudaMalloc(&d_data2, Nx * sizeof(double));
// Initialize 'h_data'
for (int i = 0; i < Nx; i++) { ... }
// Copy data from host to device
cudaMemcpy(d_data, h_data, Nx * sizeof(double), cudaMemcpyHostToDevice);
// Invoke kernel with 256 threads per block, 4096 blocks in total
square<<<(Nx + 255) / 256, 256>>>(Nx, d_data, d_data2);
// Copy the result from device back to host
cudaMemcpy(h_data2, d_data2, Nx * sizeof(double), cudaMemcpyDeviceToHost);
// Use 'h_data2' for anything
for (int i = 0; i < Nx; i++) { ... }
// Free resources
cudaFree(d_data); cudaFree(d_data2); free(h_data); free(h_data2);
}

```

Collection of threads that are generated by a kernel launch are referred to as a *grid*. When a host code launches a kernel, the CUDA runtime system generates a grid of threads that are organized in a two-level hierarchy. On the first level, each grid is organized into an array of *thread blocks*, which are in turn organized as an array of threads on the second level. All thread blocks of a grid are of the same size, and (on current CUDA devices) each thread block can contain up to 2048 threads. The number of threads in each thread block and the number of thread blocks are specified in the host code with kernel launch parameters. This hierarchy allows CUDA programming model to be scalable regardless of the underlying GPU architecture. CUDA runtime is in charge of scheduling thread blocks for execution on available SMs. Because there is no guarantee of which thread blocks are currently executing nor is their order of execution known in advance, there is no way to synchronize between thread blocks. Instead, it is expected that algorithms will be designed in a way so that they do not require synchronization between thread blocks. The array of thread blocks in a grid and the array of threads in a block are not necessarily 1D, and may be organized as 2D or 3D arrays. In fact, these arrays are always 3D, with unused dimensions initialized to 1. Figure 4.3 illustrates the thread hierarchy on an example with 2D array of thread blocks.

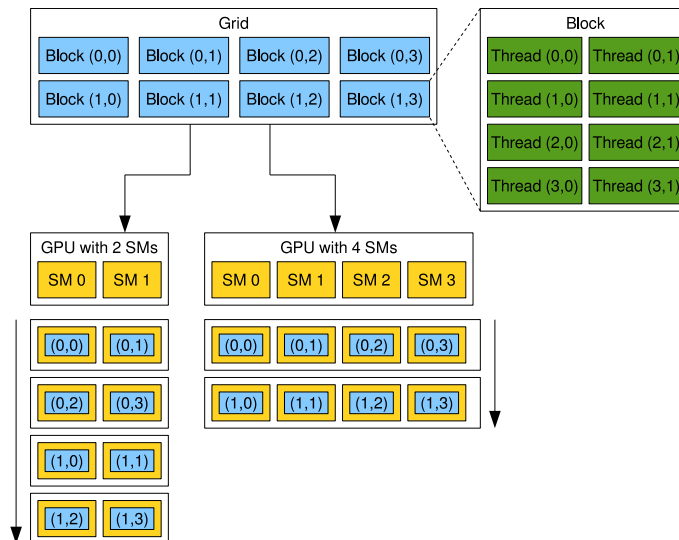


Figure 4.3: Execution of 2D array of thread blocks on a different number of SMs.

Kernels operate on device memory and in most cases have no access to host memory. Just

like the memory in a traditional C program, memory on the device must first be allocated, and then freed after use. Working with device memory is very similar to the way we work with host memory. There are four main memory operations, `cudaMalloc`, `cudaMemcpy`, `cudaMemset` and `cudaFree`, which are the counterparts of the standard C functions `malloc`, `memcpy`, `memset` and `free`, respectively. In most cases, CUDA programs follow the same processing flow, where the data required for computation is first copied from the host's memory to the device, then one or more kernels are invoked that perform some computation using that data, and finally the resulting data is copied back to the host. Since the memory transfers use the PCI-Express bus to copy the data between memories of host and device, they are in general much slower than the computation on either the host or the device. For this reason it is very important to minimize memory transfers, as well as to use asynchronous operations to overlap memory transfer with computation. We relied on these techniques in our heterogeneous implementation, with details presented in Chapter 5.

An important aspect of CUDA programming is error handling. All CUDA runtime functions return an error code and it is good practice to check for errors after any CUDA runtime function is invoked. Since the errors they produce usually do not terminate the whole program, failure to check for errors can lead to many unexpected situations and complicate debugging. In examples which we present throughout this thesis we will omit the error checking for brevity, however the proper error handling is in fact implemented in all of the GPU-related portions of the programs presented.

In this section we have described only the most basic concepts of GPU programming with CUDA. Many details have been omitted as we focus on the core concepts that we relied on in our implementations. More detailed description of CUDA can be found in [51, 52, 53].

## 4.2 CUDA implementation of shared memory algorithm

The significant differences in architecture of the GPU and its accompanying execution model on one side, and CPU architecture with its conventional execution model on the other side, did not result in significant changes to the main algorithm we employ. From the algorithm's point of view, both CPU and GPU are shared memory systems with multiple processing units, even though they are vastly different, and therefore can be used in a similar manner. Main loop of the algorithm remains unchanged, as well as each substep. Auxiliary algorithms, like the algorithm for numerical integration, which are more dependent on the execution model, need to be adapted to the new environment, as we will describe below.

While the algorithm and thus the flow of the programs remain conceptually the same, actual implementation on CUDA platform required significant changes to C/OpenMP programs. We have focused our efforts on using GPU for computationally demanding functions performing time evolution of the wave function (`calcpsidd2`, `calcnu`, `calclux`, `calcluy` and `calcluz`), normalization (`calcnorm`) and calculation of physical quantities (`calcmuen`, `calcrms`). As 1D variants of the programs are sufficiently fast even in their serial implementation, and the inherent difficulty in parallelizing recursive relations which are part of `calclux`, `calcluy` and `calcluz` functions, we have implemented the algorithm on GPU for 2D and 3D variants of the programs. In the resulting implementation, CPU host was used for initialization of variables, control of program flow, and I/O. Computation functions were reimplemented as CUDA kernels, while the initialization and output functions have remained host-only and serial. Since the initialization of variables uses host RAM

memory, programs require the same amount of host and device memory. Before any computation begins, variables are copied to GPU device, where they remain during computation, and only wave function variable (`psi`) is returned back to CPU memory when it is required for writing output. Implementation was done in several steps, each offloading more computation to the GPU.

The first step in migration to the CUDA platform was to use CUDA built-in types for complex numbers instead of the ones provided by the C99 standard. This was necessary only in real-time propagation programs because complex numbers, as implemented in C99 standard of the C language, are not directly supported in CUDA device code. However, converting them to CUDA built-in complex types is possible as they share the same binary representation. Complex numbers in CUDA do not support all the operations that C99 standard provides. Namely, the function to compute the complex exponential of a number, which we rely on in function `calcnu`, is missing. This function can be replaced by an approximation using `sincos` function available in Linux (though it is not part of a standard C library).

Next step was to replace the FFTW library used by C/OpenMP programs with the cuFFT library, which is available in CUDA. cuFFT provides two interfaces: the native one and the FFTW one [54]. FFTW interface is intended to be used as a drop-in replacement for FFTW, allowing programs written primarily with FFTW in mind to use CUDA GPUs with minimal modifications to the source code. While this interface could be used as a temporary solution during the development of CUDA programs, our goal of using CUDA for all computation relating to the propagation of the wave function could only be reached with the native cuFFT interface. This interface is modeled after FFTW but differs from it in the way plans are created and executed. In cuFFT, precision and type of data of the transform is determined by the functions initiating execution of plans, which results in a slightly different sequence of function calls. We used cuFFT in much the same way as FFTW, for doing R2C and C2R FFT transforms.

With the computation of FFT offloaded to GPU, we proceed and port the remaining wave function propagation functions to GPU. Most of the computation functions have nested loops, either two or three levels deep. This could naturally be mapped onto a 2D or 3D grid of threads, with each thread processing a single element of the corresponding arrays. This is the standard approach when writing CUDA kernels, called *monolithic kernel*, and is often found in CUDA textbooks. While this way of writing kernels is sufficient for our needs, we have opted to use *grid-stride* loops when writing kernels, as proposed in Ref. [55]. Basic grid-stride loop is illustrated in Listing 4.2. The stride of the loop is the total number of threads in the grid, which means that all memory accesses are coalesced for maximum performance, in the same way as in the monolithic approach. Addition of a loop in a kernel does not add to its instruction complexity, which would impact performance. This is because we need instructions of the same complexity in the form of conditional statements in the monolithic kernel to ensure that we do not access data out of bounds. We should also note that we have found grid-stride loops to be more flexible and easier to debug. Kernels with grid-stride loops can be launched with just a single thread and a single block, which would essentially make the kernel serial in nature and thus far easier to debug.

The remaining functions we ported, `calcnorm`, `calcmuen` and `calcrms`, rely on 1D spatial integration implemented with composite Simpson's rule. This function requires computation of a sum of elements of an array, which is trivial to implement in serial or OpenMP code, but can be much more difficult using CUDA due to lack of global synchronization between thread blocks.

Listing 4.2: Example of 1D grid-stride loop.

---

```

int i, tid_x, grid_stride_x;

tid_x = blockIdx.x * blockDim.x + threadIdx.x;
grid_stride_x = blockDim.x * gridDim.x;

for (i = tid_x; i < d_Nx; i += grid_stride_x) {
    ...
}

```

---

Nvidia GPUs have support for atomic operations, which could be used to significantly simplify the implementation, however, only the latest-generation Nvidia GPUs have the hardware support for atomic operations on double precision data. Alternative implementations of atomic operations on double precision data rely on compare-and-swap (CAS) instructions that may incur a performance penalty [56]. For this reason we have based our implementation of composite Simpson's rule on the reduction algorithm from Ref. [57]. We extended the reduction algorithm to support array sizes which are not powers of two and to compute three sums simultaneously. The algorithm operates in two steps, implemented as two kernel invocations. In the first invocation, each thread block creates a partial sum using a tree-based approach with shared memory and stores the partial result in main memory. In the second kernel invocation, this time with only one thread block, the partial sums from previous invocation are combined to produce the final result. An illustration of the algorithm is given in Figure 4.4, and its implementation is shown in Listing 4.3.




Figure 4.4: Two-phase reduction algorithm on CUDA device. In this example, four thread blocks are used in the first phase, while a single block computes the final result in the second phase.

Listing 4.3: Implementation of composite Simpson's rule with CUDA.

---

```

// Initiate temporary storage on device
void simpint_init(long N) {
    sum = alloc_double_vector_device(1);
    sumi = alloc_double_vector_device(1);
    sumj = alloc_double_vector_device(1);
    sumk = alloc_double_vector_device(1);
    tsumi = alloc_double_vector_device(ceil(1. * N / CUDA_BLOCK_SIZE / 2));
    tsumj = alloc_double_vector_device(ceil(1. * N / CUDA_BLOCK_SIZE / 2));
    tsumk = alloc_double_vector_device(ceil(1. * N / CUDA_BLOCK_SIZE / 2));
}

```



```

double simpint_gpu(double h, double *f, long N) {
    double f_sum;

    dim3 block; dim3 grid; int shmem;
    block.x = CUDA_BLOCK_SIZE; // For example, 128
    grid.x = ceil(1. * N / block.x / 2);
    shmem = 3 * block.x * sizeof(double);
    // Compute partial sums and reduce to a single block
    simpint1d_kernel1<<<grid, block, shmem>>>(tsumi, tsumj, tsumk, f+1, f, f+2, N-2, 2);
    // Compute final sums with just one block
    simpint1d_kernel1<<<1, block, shmem>>>(sumi, sumj, sumk, tsumi, tsumj, tsumk, grid.x,1);
    // Compute the final result of integration
    simpint1d_kernel2<<<1,1>>>(h, sum, sumi, sumj, sumk, f, N);
    cudaMemcpy(&f_sum, sum, 1 * sizeof(double), cudaMemcpyDeviceToHost);

    return f_sum;
}

__global__ void simpint1d_kernel1(double *sumi, double *sumj, double *sumk, double *ini,
    double *inj, double *ink, long N, int step) {
    extern __shared__ double psumi[]; // Shared memory for partial sums
    double *psumj = &psumi[blockDim.x]; double *psumk = &psumi[2 * blockDim.x];
    double tsumi = 0.; double tsumj = 0.; double tsumk = 0.;

    int tid = threadIdx.x;
    long idx = (blockIdx.x * blockDim.x + tid) * step;

    if (idx < N) {
        tsumi += ini[idx]; tsumj += inj[idx]; tsumk += ink[idx];
    }

    psumi[tid] = tsumi; psumj[tid] = tsumj; psumk[tid] = tsumk;

    __syncthreads();

    // Start the shared memory loop on the next power of 2 less than the block size.
    // If block size != power of 2, accumulate the intermediate sums in the remainder range.
    int pow2 = blockDim.x;

    if (pow2 & (pow2-1)) {
        while (pow2 & (pow2-1)) {
            pow2 &= pow2-1;
        }
        if (tid >= pow2) {
            psumi[tid-pow2] += psumi[tid];
            psumj[tid-pow2] += psumj[tid];
            psumk[tid-pow2] += psumk[tid];
        }
        __syncthreads();
    }
    for (int th = pow2 >> 1; activeThreads; activeThreads >>= 1) {
        if (tid < activeThreads) {
            psumi[tid] += psumi[tid+activeThreads];
            psumj[tid] += psumj[tid+activeThreads];
            psumk[tid] += psumk[tid+activeThreads];
        }
        __syncthreads();
    }
    if (tid == 0) {
        sumi[blockIdx.x] = psumi[0]; sumj[blockIdx.x] = psumj[0]; sumk[blockIdx.x] = psumk[0];
    }
}

```

```

}

__global__ void simpint1d_kernel2(double h, double *sum, double *sumi, double *sumj,
    double *sumk, double *f, long N) {
    *sum = *sumj + 4. * *sumi + *sumk;
    if(N % 2 == 0) *sum += (5. * f[N - 1] + 8. * f[N - 2] - f[N - 3]) / 4.;
    *sum = *sum * h / 3.;
}

```

---

The aforementioned algorithm works well in situations where we need just one call of the integration function. However, in our programs, this routine is often called inside nested loops. Invoking two kernels in each iteration creates an overhead that makes this approach slower than the OpenMP-threaded programs. In case of nested loops, we can use the similar approach used in `calclux`, `calcluy` and `calcluz` functions, by performing the whole integration in each GPU thread. Just like in the case of recursive relations, the integration via composite Simpson's rule is much slower on GPU threads than on CPU threads, however the large number of parallel GPU threads successfully hides their individual latency. Since both approaches have strengths and weaknesses, we used a combination of both in our implementation of `calcnorm`, `calcmuen` and `calcrms` functions. These functions share a common flow, shown in Listing 4.4.

Listing 4.4: Performing reduction along multiple directions, demonstrated on the 3D real-time `calcrms` function.

---

```

void calcrms(...) {
    // Prepare data in a tensor
    calcrms_kernel1x<<<dimGrid3d, dimBlock3d>>>(psi, tmpyzx, d_x2);
    // Reduce from 3D to 2D
    simpint3d_kernel<<<dimGrid2d, dimBlock2d>>>(dx, tmpyzx, tmpyz, Ny, Nz, Nx);
    // Reduce from 2D to 1D
    simpint2d_kernel<<<dimGrid1d, dimBlock1d>>>(dz, tmpyz, tmpy, Ny, Nz);
    // Reduce from 1D to scalar
    rms[1] = sqrt(simpint_gpu(dy, tmpy, Ny));
    ...
}

__global__ void simpint2d_kernel(double h, cudaPitchedPtr matrix, double *vector, long N1,
    long N2) {
    ...
    for (i = blockIdx.x * blockDim.x + threadIdx.x; i < N1; i += blockDim.x * gridDim.x) {
        matrixrow = get_double_matrix_row(matrix, i);
        vector[i] = simpint(h, matrixrow, N2); // Each thread calls sequential function
    }
}

__global__ void simpint3d_kernel(double h, cudaPitchedPtr tensor, cudaPitchedPtr matrix,
    long N1, long N2, long N3) {
    ...
    for (i = blockIdx.y * blockDim.y + threadIdx.y; i < N1; i += blockDim.y * gridDim.y) {
        matrixrow = get_double_matrix_row(matrix, i);
        for (j = blockIdx.x * blockDim.x + threadIdx.x; j < N2; j += blockDim.x * gridDim.x){
            tensorrow = get_double_tensor_row(tensor, i, j);
            matrixrow[j] = simpint(h, tensorrow, N3);
        }
    }
}

```

---

In Section 3.1 we mentioned that we perform flat allocation of CPU memory even when dealing with multidimensional data. This property was especially useful when transitioning to GPU. Contiguous (flat) memory allowed us to transfer whole multidimensional variables with one call to `cudaMemcpy`. While we also used flat allocations on GPU, we did not rely on double or triple pointers to index our variables. Using them is possible on GPU, however correct usage is much more complicated, and may reduce performance [58]. Since we are focused on 2D and 3D variants of our programs, we also had to ensure that we are optimally accessing memory of multidimensional variables. To this end, we have used *pitched* memory allocations, to make sure that the allocation is appropriately padded to allow coalesced access. Pitched memory can be allocated with `cudaMallocPitch` or `cudaMalloc3D` which perform the same allocation routine, just with a slightly different interface. Transferring pitched data to/from GPU is also done with special functions, `cudaMemcpy2D` and `cudaMemcpy3D`, which follow the same differences. We settled on using `cudaMalloc3D` and `cudaMemcpy3D` in our programs, which we found to be more convenient due to the fact that `cudaMalloc3D` neatly groups all allocation-related variables in one C structure. The pitch of allocation must be used to access elements of an allocated array, which makes the index calculation slightly different, as shown in Listing 4.5. Note that in these functions the innermost dimension (`Nz` in our case) is specified first, and the slowest changing (`Nx`) last, resulting in an unusual order when compared to the memory allocation functions we presented in the previous chapter. Nevertheless, this is just a cosmetic issue which does not impact performance in any way.

Listing 4.5: Accessing pitched memory on a device.

---

```
// Allocated memory resides in a special structure
struct cudaPitchedPtr tensor;
cudaMalloc3D(&tensor, make_cudaExtent(Nz * sizeof(double), Ny, Nx);

// Define code for accessing specified row
// Equivalent to double *row = tensor[slice][row] on host
__device__ double *get_double_tensor_row(cudaPitchedPtr tensor, long slice, long row) {
    return (double *)(((char *)tensor.ptr) + slice * tensor.pitch * tensor.ysize) + row *
        tensor.pitch);
}

// From device code, access the specified row, and k-th element inside
// Equivalent to double val = tensor[i][j][k] in host code
double *tensorrow = get_double_tensor_row(tensor, i, j);
double val = tensorrow[k];
```

---

Changes required to support pitched memory were in most cases straightforward, however special care has to be taken when working with `cuFFT`, as its basic interface does not assume that memory is padded. Similarly to the advanced interface of `FFTW`, the advanced interface of `cuFFT` supports working with a subset of a larger (in this case padded) multidimensional array. Using pitched memory does not cause a significant increase in memory usage on GPU in most cases. On the other hand, parallelization of outer loops requires additional per-thread memory, in the same way as in `OpenMP` parallelization. Since GPU usually executes more threads simultaneously than what is available on the CPU, this increase can be significant. Given that GPUs in most cases have less RAM than their host, high memory usage means that the GPU can be used for much smaller problem sizes.

To curb memory usage, we reused memory as much as possible, and reduced the memory foot-

print of some functions. We allocate up to four matrices/tensors (for 2D and 3D, respectively), which are used throughout the programs. These four matrices/tensors hold wave function values (variable `psi`), temporary data, trap potential (`pot`) and Fourier transform of the dipolar potential (`potdd`).

Allocating data for the two potentials that we use has been made optional, and can be exploited to further reduce memory usage. We exposed this functionality through `POTMEM` input parameter. The three possible values of `POTMEM` are 0, 1 and 2. Setting the value of `POTMEM` to 2 instructs the programs to allocate two separate matrices/tensors which will be used to store trap potential and dipolar potential in GPU memory. This provides the best performance at the expense of increased memory usage, and thus smaller maximum mesh size. Setting `POTMEM` to 1 will allocate only one matrix/tensor, which will store one of the two potentials, as needed. The required potential will be copied asynchronously in the background while other computations on GPU take place. Initially, the matrix/tensor will store trap potential. During execution of the `calcpsidd2` function, the dipolar potential will be copied over the trap potential, and after it is no longer needed, the trap potential will replace it for the rest of the computation. This cycle will repeat itself in each iteration of the wave function propagation. The last option, setting `POTMEM` to 0, will not allocate any memory for the potentials on GPU, and will instead expose the host memory of the potentials as mapped memory. Leaving the `POTMEM` parameter undefined will instruct the programs to try and select the optimal value based on the size of the mesh. Figure 4.5 illustrates the three different settings of the `POTMEM` parameter. We suggest using `POTMEM` value of 2 if possible, or leaving it undefined, and using values of 1 or 0 if programs do not correctly predict the optimal value, or the problem cannot fit into GPU memory.




Figure 4.5: Illustration of placement of relevant variables in CPU and GPU memory.

We allocate matrix/tensor for temporary data as a flat array of double or double complex values, in the same way as for the other matrices/tensors. This matrix/tensor is used in real-time propagation functions `calclux`, `calcluy` and `calcluz` to store the Crank-Nicolson coefficients. Alternatively, complex matrix/tensor can be further divided into two matrices/tensors of double values. Since the arrays we get in this way are also used in computation of R2C FFT, we had to ensure the matrix/tensor had the proper size, that is  $2 * N_x * (N_y / 2)$  and  $2 * N_x * N_y * (N_z / 2)$  in 2D and 3D variants of programs, respectively. To use this memory allocation scheme effectively, we had to reorganize some of the computations inside `calcmuen` and `calcpsidd2`. In fact, the memory optimizations of those two functions mentioned in Chapter 3 were inspired by the optimizations we had to introduce in GPU variants of programs. Changes required in these functions are conceptually very similar to the ones we presented in Chapter 3, therefore we will

not go into all of their details here. One notable difference is the additional restriction we have to place on FFT computation in order to keep its memory usage in check. We do this by ensuring that cuFFT reuses temporary storage we already allocated to store its intermediate results. By default, cuFFT will allocate memory for intermediate results separate from the already allocated input/output arrays. This can be changed, however we must ensure that we have allocated enough memory to meet cuFFT demands. The amount of memory required for FFT temporary data varies with the transform size due to different algorithms used. Some transform sizes require much more memory than others, in some cases up to eight times more [54]. This amount can be larger than what we have already allocated, in which case the programs will select a nearest larger transform size that can fit in the allocated memory and perform the computation as if the user has entered the corresponding number in the input file for the mesh size. The output of programs will report if it had to make these adjustments to the mesh size.

When all the memory optimizations are combined, we get an approximately 50% reduction in the memory usage compared to the initial GPU implementation. This not only allowed us to use this implementation for much larger mesh sizes, but also to use similar optimizations to improve all variants of programs described in this thesis.

CUDA programs have been grouped and published as *DBEC-GP-CUDA* package in Ref. [17], which can be downloaded from Refs. [59, 37]. The package consists of 2D (corresponding to  $x$ - $y$  and  $x$ - $z$  planes) and 3D programs, in both imaginary- and real-time propagation, resulting in 6 programs in total. We followed the similar naming convention as in the case of OpenMP-parallelized programs, by adding a suffix “-cuda” to the base program name. Performance evaluation of these programs is given in Chapter 8.

## Chapter 5

# Hybrid algorithm for heterogeneous computing systems

In this chapter we describe the hybrid algorithm targeting heterogeneous systems. While this algorithm can be implemented on every heterogeneous system where processing resources have separate memory (i.e., the memory is not shared), here we focus on platforms consisting of CPUs and GPUs on the same computer, as it is the most common heterogeneous platform. Thus, the algorithm described here is a combination of algorithms from Chapters 3 and 4, further extended to allow for simultaneous computation. The resulting hybrid algorithm and its implementation is made possible by the fact that pure GPU solution does not use CPU for any computation, rather the CPU is used only to control the execution of GPU kernels and perform I/O operations. Modern computers with powerful GPUs usually also have a powerful multi-core CPUs, so using only GPU presents a missed opportunity to achieve better performance through the use of all available resources.

Our goal was to extend the algorithm so that we can reuse as many portions of the CPU and GPU programs as we can. We did this by introducing GPU support over the base OpenMP-parallelized CPU programs from Chapter 3. Numerically most demanding functions were extended to asynchronously offload computation to GPU, and then proceed with the computation on CPU using multiple threads.

The most important part of creating a hybrid algorithm is the distribution of data between CPU and GPU in a way that maximizes the performance of any function involved in time propagation of the wave function. We first give a detailed description of how data is distributed, followed by a description of hybrid computation functions and hybrid FFT in Section 5.1. Details of basic implementation that uses all available CPU cores and GPUs are given in Section 5.2. The implementation of the data distribution scheme revealed that the time it takes to transfer the data to GPU and back dominates the execution time. To optimize the implementation and minimize the impact of data transfers, we have used CUDA streams to hide the latency of data transfers, as demonstrated in Section 5.3 along with other memory usage optimizations we applied.

## 5.1 Description of the hybrid algorithm

In order to use CPU and GPU at the same time, we have to divide the work efficiently between them. If we put aside the obvious hardware architecture difference between CPU and GPU which we described in Section 4.1, we can consider the computer with GPU as being similar to the distributed memory system consisting of two computing nodes with different characteristics. Both CPU and GPU have their own memory, and are connected through a fast interconnect, in this case the PCI-Express. This means that we need an approach similar to the one distributed memory algorithms use to achieve the desired parallelization. The most important difference here is that all data is available in main (CPU) memory, and only portions of it need to be transferred to GPU memory, as opposed to the true distributed memory systems, where no single part of the system has all the data. Therefore, the general approach to perform this type of hybrid computation is to:

- (i) designate a portion of data to be processed by GPU,
- (ii) copy that data to GPU while simultaneously using CPU for computation over the remaining data,
- (iii) transfer the data back from GPU, overwriting old data, and
- (iv) synchronize CPU and GPU.

This flow is illustrated in Figure 5.1.

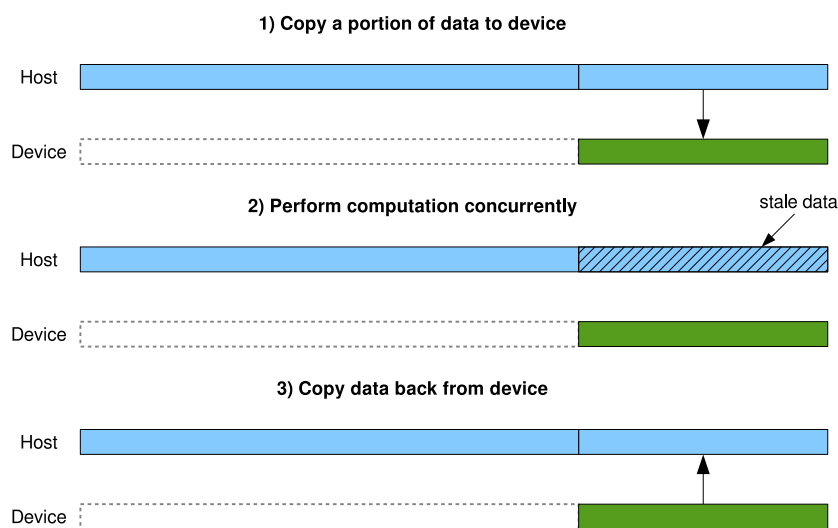


Figure 5.1: Flow of data between host and device.

Not every function is feasible for simultaneous computation on CPU host and GPU device. Functions calculating physical properties (chemical potential, energy, norm and RMS size), or functions not executed in each iteration are not suitable for computation on GPU for various reasons. Specifically, functions which require multiple invocations of numerical integration do not achieve noticeable speedup on GPU, and when coupled with data transfer, actually perform worse than on CPU alone. We have therefore focused our efforts on speeding up only the functions relating to time propagation of the wave function, which are invoked in each step of the main loop, i.e.

`calcnu`, `calclux`, `calcluy`, `calcluz` and `calcpsidd2`. The function computing the norm of the wave function was not offloaded to GPU for real-time propagation, as it is not invoked in each iteration. On the other hand, in imaginary-time propagation, in which normalization is performed as the last step of each iteration, we perform a concurrent computation on both CPU and GPU only for the normalization step (i.e., dividing the wave function by the norm), not the computation of the norm itself.

Data can be offloaded in various ways [60, 61] and, in general, any data distribution scheme which targets distributed memory systems may be used. However, we settled on a simple approach using 1D decomposition, also known as *slab decomposition*.

When using slab decomposition, data is distributed along one dimension, usually the slowest-changing one, and the remaining dimension (in 2D programs) or dimensions (in 3D programs) of data remain local. This means that we can perform computation on the local data efficiently, and no data exchanges are necessary. Depending on a data access pattern, in our programs we rely on decompositions either along the  $x$  direction (discretized with the  $N_x$  spatial points), or along the  $y$  direction (discretized with the  $N_y$  spatial points). In 3D programs, there is no need to decompose the data along the  $z$  direction. In case we need data from the distributed dimension to become local in order to perform some computation (e.g., to update the wave function values in `calclux` function when data are decomposed along the  $x$  direction), we have to reassemble data in host memory and decompose again along the appropriate dimension. Figure 5.2 illustrates this concept.




Figure 5.2: Two offload patterns used. If offloaded along  $x$  direction, the  $y$  direction remains local, however if we need to access whole  $x$  direction data, we need to decompose along  $y$  direction.

With the decomposition scheme in place, we can consider how to divide the computation of the wave function propagation in time between CPU and GPU. Once the data have been distributed between CPU and GPU, on the CPU side only the exit condition of the outer loop needs to be adjusted, so that the CPU processes a smaller number of elements (e.g., by replacing  $N_x$  with `cpuNx`). Similar changes are required on the GPU side as well, however we also need to ensure that GPU has valid data to work with, and that data are reassembled in host memory after the computation on GPU is done. Introduction of offloading to GPU did not require significant changes to the functions calculating individual parts of the wave function propagation, and all changes are similar.

Initially, we distribute the data along the slowest-changing dimension. This allows CPU and GPU to simultaneously perform computation of wave function propagation w.r.t.  $H_1$  and  $H_3$  parts



of the Hamiltonian (and  $H_4$  when working in 3D). The propagation of the wave function w.r.t.  $H_2$ , corresponding to the  $x$  direction, can be done locally only if we decompose the data along the  $y$  direction before invoking the function `calcflux` on CPU and GPU. Note that we do not have to transfer data from GPU back to host memory at the end of each function, due to the fact that time propagation w.r.t.  $H_2$ ,  $H_3$  and  $H_4$  parts of the Hamiltonian can be done in arbitrary order in each step. Without this, the time propagation workflow of 3D programs would have to involve the following steps, according to Figure 3.2:

1. calculate the dipolar term,
2. transfer the data decomposed along  $x$  direction to GPU,
3. propagate the wave function w.r.t.  $H_1$
4. reassemble the data in host memory,
5. transfer the data decomposed along  $y$  direction to GPU,
6. propagate the wave function w.r.t.  $H_2$
7. reassemble the data in host memory,
8. transfer the data decomposed along  $x$  direction to GPU,
9. propagate the wave function w.r.t.  $H_3$
10. propagate the wave function w.r.t.  $H_4$
11. reassemble the data in host memory.

Rearranging the order of time-propagation substeps allows us to remove one transfer of data to GPU and its subsequent reassembly in host memory:

1. calculate the dipolar term,
2. transfer the data decomposed along  $x$  direction to GPU,
3. propagate the wave function w.r.t.  $H_1$
4. propagate the wave function w.r.t.  $H_3$
5. propagate the wave function w.r.t.  $H_4$
6. reassemble the data in host memory.
7. transfer the data decomposed along  $y$  direction to GPU,
8. propagate the wave function w.r.t.  $H_2$
9. reassemble the data in host memory,

Therefore, we use the above, optimized sequence in our 3D programs, and similarly in 2D.

To complete our hybrid algorithm, we also need to distribute computation of the dipolar term between CPU and GPU, where complexity arises in performing DFT on distributed data. Currently available FFT libraries target either CPU or GPU for their computation, but unfortunately not both at the same time. There are numerous attempts to develop specialized FFT libraries which would enable this [62, 63, 64], however a full-featured library with support for R2C transforms with advanced data layout is still not available. The approach we used is to rely on existing libraries for actual transforms, but perform data distributions manually. The libraries we use, FFTW on CPU and cuFFT on GPU, support advanced data layouts as well as working on a subset of the whole data, allowing for an efficient implementation.

To perform the Fourier transform simultaneously on CPU host and GPU device, we split the single multidimensional transform into a series of 1D transforms along each dimension of the input data, which can be computed on CPU and GPU independently. This approach is known as the *row-column* algorithm [65], and is often used in FFT libraries. The essence of this algorithm can best be summarized in an example of a 2D FFT. Given a matrix  $N_x \times N_y$ , we compute the DFT in the following way:

1. Transfer portion of the input array, decomposed along the  $x$  direction, to device memory. We transfer last  $\text{gpu}N_x \times N_y$  consecutive array elements to GPU. This can be done as a single memory copy operation due to the flat allocation that was used. The choice whether to offload data to GPU memory from the beginning or the end of the input array is arbitrary.
2. Perform DFT along the  $y$  direction on both CPU and GPU concurrently. CPU will perform  $\text{cpu}N_x$  such transformations, while the GPU will perform  $\text{gpu}N_x$  1D DFTs. Each of these transforms will take a subset of the input array, which are  $N_y / 2 + 1$  elements apart. Note that CPU will not do anything with the last  $\text{gpu}N_x \times N_y$  elements.
3. Copy the array with the transform back from GPU to CPU, writing over the stale data with the relevant portion transformed on GPU. After this step, we have the complete array transformed along the  $y$  direction residing in host memory.
4. Transfer portion of the input array, decomposed along the  $y$  direction, to GPU memory. Similarly to step 1, we transfer last  $N_x \times \text{gpu}N_y$  elements to GPU.
5. Perform DFT along the  $x$  direction on both CPU and GPU. This time, CPU will perform  $\text{cpu}N_y / 2 + 1$  such transformations, while GPU will perform remaining  $\text{gpu}N_y / 2$  transformations. The halving of the number of transformations is due to use of R2C transformations. In each transform, elements are  $N_y / 2 + 1$  places apart, while the first element of each transform is adjacent to the previous one.
6. Copy the array with the transform back from GPU to CPU. With this step completed, we have the full FFT of the input array residing in host memory. This step may be omitted if the computation that follows does not require the FFT of input data to be fully assembled in host memory. This is the case in our algorithm, as we use the resulting transformed array only for subsequent computation on GPU. Since the memory transfer is expensive, omitting this step leads to significant performance improvement.

The inverse Fourier transform can be done in an analogous way. In 3D, the general principle remains the same, however, we do not separate the DFT into three 1D transforms, but to one 2D and one 1D transform. We do this due to better performance of 2D transforms, which in both libraries used is found to be better than two 1D transforms. Figure 5.3 illustrates the DFT algorithm described above.

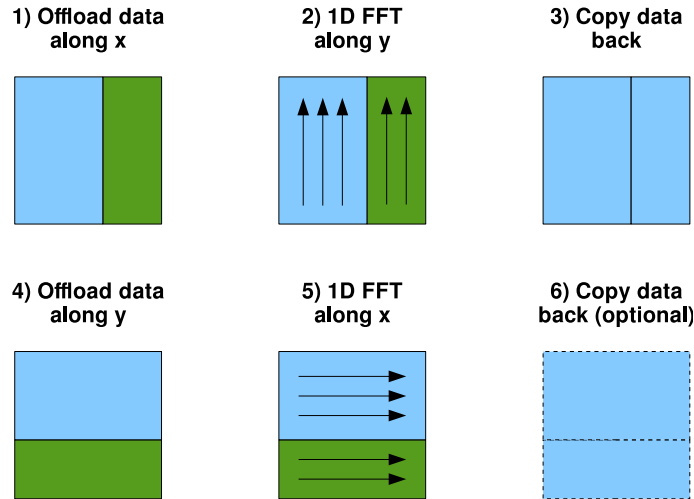


Figure 5.3: Hybrid algorithm for concurrent FFT on host and device.

## 5.2 Implementation of the hybrid algorithm

The hybrid algorithm was implemented by extending the C/OpenMP implementation with the data distribution scheme from the previous section, while the computation functions on the GPU device were taken from the CUDA implementation described in Section 4.2. This allowed us to reuse much of the existing code, with straightforward modifications.

Decomposition of data was implemented efficiently using the CUDA built-in functionality for copying memory to and from the GPU. In case of decomposition along the  $x$  direction, we divide the data into two (possibly uneven) parts with sizes `cpuNx` and `gpuNx`, where `cpuNx + gpuNx = Nx`. Now the CPU works with the first `cpuNx × Ny` elements in 2D programs and `cpuNx × Ny × Nz` elements in 3D programs, while the GPU works with remaining `gpuNx × Ny` elements in 2D programs and `cpuNx × Ny × Nz` elements in 3D programs. Listing 5.1 shows how the memory is copied from host to device in this scenario. The opposite, copying from device to host, is done in a similar way, where only the source and destination pointers are exchanged in the copy parameters.

Listing 5.1: Initialization and copying of data from host to device, distributed along the  $x$  direction.

```
// Allocate host data
double *h_matrix = alloc_double_matrix(Nx, Ny);

// Allocate only the required amount of data on device
struct cudaPitchedPtr d_matrix_x = alloc_double_matrix_device(gpuNx, Ny);

// Prepare memcpy operation
struct cudaMemcpy3DParms cpy_x_h2d;
cpy_x_h2d.srcPtr = make_cudaPitchedPtr(h_matrix[cpuNx], Ny * sizeof(double), Ny, gpuNx);
```

```

cpy_x_h2d.dstPtr = d_matrix_x;
cpy_x_h2d.extent = make_cudaExtent(Ny * sizeof(double), gpuNx, 1);
cpy_x_h2d.kind = cudaMemcpyHostToDevice;

// Copy data from host to device
cudaMemcpy3D(&cpy_x_h2d);

```

Similar concept is applied when distributing data along the  $y$  direction. This time CPU works with  $N_x \times \text{cpuNy}$  in 2D ( $N_x \times \text{cpuNy} \times N_z$  in 3D), while the GPU portion is  $N_x \times \text{gpuNy}$  in 2D,  $N_x \times \text{gpuNy} \times N_z$  in 3D (Listing 5.2).

Listing 5.2: Initialization and copying of data from host to device, distributed along the  $y$  direction.

```

// Allocate only the required amount of data on device
struct cudaPitchedPtr d_matrix_y = alloc_double_matrix_device(Nx, gpuNy);

// Prepare memcpy operation
struct cudaMemcpy3DParms cpy_y_h2d;
cpy_y_h2d.srcPtr = make_cudaPitchedPtr(h_matrix[0], Ny * sizeof(double), Ny, Nx);
cpy_y_h2d.srcPos = make_cudaPos((cpuNy) * sizeof(double), 0, 0);
cpy_y_h2d.dstPtr = d_matrix_y;
cpy_y_h2d.extent = make_cudaExtent(gpuNy * sizeof(double), Nx, 1);
cpy_y_h2d.kind = cudaMemcpyHostToDevice;

// Copy data from host to device
cudaMemcpy3D(&cpy_y_h2d);

```

We can also notice how the flat allocation of memory on CPU is once again useful. Without it, the transfer of data between CPU and GPU would be much more complicated, involving copying individual rows of the innermost dimension. Flat allocation allows for a single copy operation to be invoked to copy the whole GPU portion of the data. It is also much easier to further divide the copy operation when using CUDA streams, as we will demonstrate in Section 5.3.

With the data transfer mechanism in place, we can focus on enabling concurrent computation. CUDA kernel launches are asynchronous with regard to the CPU, allowing us to combine them with OpenMP-based computation on CPU. There are numerous ways to implement GPU offloading. We have investigated three possible approaches, which we now describe.

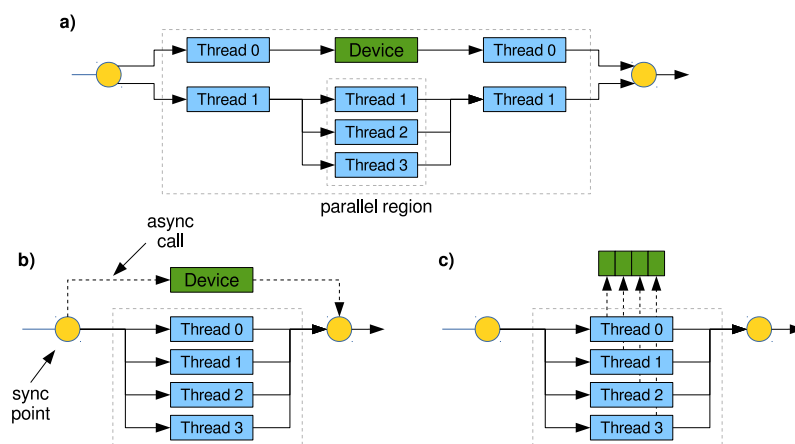


Figure 5.4: Three strategies considered for GPU offloading.

The approach we decided to use, illustrated on Figure 5.4(a), involves the use of OpenMP nested parallelism. We first create a parallel region that consists of two threads, one of which will invoke the GPU functions, while the other will do the same for CPU functions. The GPU-related thread invokes the required functions and waits for their completion. Meanwhile, the CPU-related thread spawns new (nested) threads, which perform the CPU computation with the changed loop exit condition (Listing 5.3). The optimal number of newly spawned threads is equal to `nthreads - 1`, where `nthreads` is the total number CPU cores in the system.

---

Listing 5.3: Use of nested OpenMP parallelism to control both host and device computation.

---

```
void hybrid_function(...) {
...
    #pragma omp parallel private(threadid) num_threads(2)
    {
        threadid = omp_get_thread_num();

        if (threadid == 0) {
            gpu_function(); // Async calls on device
            sync_with_gpu(); // Wait for completion
        }

        if (threadid == 1 || omp_get_num_threads() != 2) {
            cpu_function(...); // Invokes nested parallelism
        }
    }
...
}

void cpu_function(...) {
...
    nthreads = omp_get_num_threads();

    #pragma omp parallel for private(...) num_threads(nthreads - 1)
    for (...) { ... }
...
}
```

---

The second approach, illustrated in Figure 5.4(b), does not use nested parallelism. Instead of dedicating one thread to GPU-related tasks, here we asynchronously invoke GPU functions from the master thread. Before performing computation on CPU using OpenMP threads, we instruct the GPU to perform the work from the master CPU thread. Next we invoke the CPU computation via OpenMP threads (one for each CPU core), and finally, after CPU completes its computation, we synchronize with the GPU to avoid race conditions (Listing 5.4). This approach works because GPU supports asynchronous calls. Kernel invocations are naturally asynchronous with regard to the CPU, while the memory transfer calls can be asynchronous, a feature which will be explained in detail in Section 5.3. While simpler to implement, this solution performs worse than aforementioned implementation with a dedicated GPU thread. This is because the asynchronous GPU calls still require some CPU work, and the frequent context switches the operating system has to make to service the GPU calls negatively affects performance of OpenMP threads.

Final approach we considered involves using each OpenMP thread for both CPU and GPU functions. This way, each OpenMP thread was responsible for transferring optimal amount of data to GPU and back, and invoking kernels on GPU, as illustrated in Figure 5.4(c). This approach did

Listing 5.4: Use of asynchronous calls to device functions while using OpenMP for computation on host.

---

```

void hybrid_function(...) {
...
    gpu_function();

    #pragma omp parallel private(threadid, ...)
    {
        threadid = omp_get_thread_num();

        #pragma omp for
        for (...) { ... }
    }

    sync_with_gpu();
...
}

```

---

not bring observable performance improvement and led to a much more complex implementation, requiring more bookkeeping, especially with the introduction of multi-GPU support. Additionally, the flexibility of choosing an ideal amount of data to offload to GPU for each function was somewhat lost because the amount of data offloaded to GPU was now controlled by the number of active OpenMP threads. The introduction of per-thread default streams in CUDA 7.0 (not available at the time of development) could have been used to simplify the implementation, however it would have still been more complex than our other approach, and would not solve the flexibility issue. For these reasons we have discarded this approach.

At this point, we have an implementation that performs most of the computation simultaneously on CPU and GPU, but we also need to implement hybrid FFT routines. The hybrid row-column algorithm from the previous section for DFT can be implemented using existing libraries very efficiently, without the requirement to make the input data of every transform consecutive in memory. The only complexity of this approach is that we have to use the *advanced* interface of FFTW and cuFFT that allows setting appropriate offsets, strides and number of transforms manually. We can see how this works in Listing 5.5, which demonstrates how a 3D FFT can be divided between host and device. This approach does not require more memory than regular multidimensional CPU or GPU FFT. We have tested several combinations of in-place and out-of-place transforms, and have concluded that on the device, the transform along the  $x$  direction (when data is distributed along the  $y$  direction) performs much better when done out-of-place, while for all other transforms the difference between in-place and out-of-place operation is usually not significant. In our case, switching from one type of transform to another is easy, the only change needed is to allocate memory for the output array, and pass its pointer to the planner.

Listing 5.5: Creating and executing plans for concurrent FFT on host and device.

---

```

// Allocate proper amount of data on host
double ***tensor = alloc_double_tensor(Nx, Ny, Nz);
fftw_complex *fft_array = alloc_complex_vector(Nx * Ny * (Nz/2+1));

// Using FFTW's 'advanced' interface on host
int fft_rank = 2; int nfr[] = {Ny, Nz}; int howmany = cpuNx;
int idist = Ny * (Nz/2+1) * 2; int odist = Ny * (Nz/2+1);

```

```

int  istride = 1; int  ostride = 1;
int  *inembed = NULL, *onembed = NULL;
// Define an out-of-place plan for cpuNx transforms of size Ny * (Nz/2+1)
plan_fw_row = fftw_plan_many_dft_r2c(fft_rank, nfr, howmany, **tensor, inembed, istride,
    idist, fft_array, onembed, ostride, odist, FFTW_MEASURE);

fft_rank = 1; int  nfc[] = {Nx}; howmany = cpuNy * (Nz/2 + 1);
idist = 1; odist = 1;
istride = Ny * (Nz/2+1); ostride = Ny * (Nz/2 + 1);
// Define an in-place plan for cpuNy * (Nz/2 + 1) transforms of size Nx
plan_fw_col = fftw_plan_many_dft(fft_rank, nfc, howmany, fft_array, inembed, istride,
    idist, fft_array, onembed, ostride, odist, FFTW_FORWARD, FFTW_MEASURE);
...
// Plans can be executed with the following command
fftw_execute_dft_r2c(plan_fw_row, **tensor, fft_array);
fftw_execute_dft(plan_fw_col, fft_array, fft_array);

// Allocate proper amount of data on device
struct cudaPitchedPtr d_tensor_x = alloc_double_tensor_device(gpuNx, Ny, (Nz/2+1)*2);
struct cudaPitchedPtr d_tensor_y = alloc_double_tensor_device(Nx, gpuNy, (Nz/2+1)*2);
struct cudaPitchedPtr d_tensor_y_tran = alloc_double_tensor_device(Nx, gpuNy, (Nz/2+1)*2);

// Using equivalent cuFFT calls on device
fft_rank = 2; int  nfr[] = {Ny, Nz}; howmany = gpuNx;
idist = Ny * (d_tensor_x.pitch / sizeof(cufftDoubleReal));
odist = Ny * (d_tensor_x.pitch / sizeof(cufftDoubleComplex));
istride = 1; ostride = 1;
int  inembed_fwr[] = {Ny, d_tensor_x.pitch / sizeof(cufftDoubleReal)};
int  onembed_fwr[] = {Ny, d_tensor_x.pitch / sizeof(cufftDoubleComplex)};
// Define an in-place plan for gpuNx transforms of size Ny * (Nz/2+1), padded by 'pitch'
cufftMakePlanMany(plan_fw_row, fft_rank, nfr, inembed_fwr, istride, idist, onembed_fwr,
    ostride, odist, CUFFT_D2Z, howmany, &ws_fwr);

fft_rank = 1; int  nfc[] = {Nx}; howmany = gpuNy * (Nz/2+1);
idist = 1; odist = 1;
istride = gpuNy * (d_tensor_y.pitch / sizeof(cufftDoubleComplex));
ostride = gpuNy * (d_tensor_y_tran.pitch / sizeof(cufftDoubleComplex));
int  inembed_fc[] = {Nx}; int  onembed_fc[] = {Nx};
// Define an out-of-place plan for gpuNy * (Nz/2+1) transforms of size Nx, padded by 'pitch'
cufftMakePlanMany(plan_fw_col, fft_rank, nfc, inembed_fc, istride, idist, onembed_fc,
    ostride, odist, CUFFT_Z2Z, howmany, &ws_fwc);
...
// Plans can be executed with the following command
cufftExecD2Z(plan_fw_row, (cufftDoubleReal *) d_tensor_x.ptr, (cufftDoubleComplex *)
    d_tensor_x.ptr)
cufftExecZ2Z(plan_fw_col, d_tensor_y.ptr, d_tensor_y_tran.ptr, CUFFT_FORWARD);

```

Implementation described so far utilizes all available CPU cores, but only a single GPU. Since a computer may have multiple GPUs installed, we have implemented support for using all available GPUs. Currently, CUDA supports up to four GPUs in a single computer, meaning we can offload up to four times the amount of data to GPUs. We have tested the programs on a computer with up to three GPUs, maximum allowed by the hardware at PARADOX supercomputing facility.

CUDA provides a simple API for choosing which GPU is in use, via the `cudaSetDevice` function. Every command issued afterward targets the selected device. The developer only needs to keep track of pointers used for each device. This simple mechanism allows us to keep our data transfer routines intact, and add support for multiple GPUs with little effort. The general concept we used is to loop over available GPUs, and issue the desired operations on each device (Listing 5.6).

Listing 5.6: Changing the currently active GPU device.

---

```

cudaGetDeviceCount(&nDevices);

for (i = 0; i < nDevices; i++) {
    cudaSetDevice(i);
    ... // Issue commands as usual
}

```

---

This approach has only one drawback, and that is that it will issue the same amount of work to all GPUs, which may not be optimal if GPUs are different. We believe that the trade-off between the simplicity of implementation and feature-completeness was worthwhile, because in most HPC installations, all GPUs within a single computing node are of the same model. Alternatively, we could have also made our programs aware of the differences between multiple GPUs installed. CUDA has an API for discovering features of GPUs which would be useful in selecting the optimal parameters for the amount of work for each GPU, however it would still be a significant effort to implement this in our programs, and for the users to optimally use this functionality.

During implementation of the algorithm described in this section we noticed that the majority of the execution time on GPU is spent waiting for the memory transfers to complete. To perform a single multidimensional FFT, we need four memory transfers, which, if implemented in a sequential manner, becomes the main bottleneck. In the next section we describe a technique to mitigate this.

### 5.3 Optimization of data transfers

Transferring memory is usually a significantly slower process than computation, to such an extent that if memory copies between host and device are common, that defeats the whole purpose of using GPU. In our implementation, several copies are needed during one iteration of the wave function propagation, limiting the use of GPU to smaller mesh sizes.

There are several remedies to this problem. An obvious one is to reduce the amount of data transferred, by only transferring subsets of data that are required on the GPU side. The programs presented here minimize the amount of data involved by copying only the parts of arrays required by the computation that immediately follows. Although this technique brings observable improvement, it alone is not enough to consider the problem solved. In fact, reducing the amount of data that is copied to GPU and back also means that we use less GPU memory, allowing for use of larger meshes on GPU, which in turn take more time to transfer. Therefore, we rely on additional optimizations in order to improve memory transfer.

To be able to do this, we have to understand and exploit the nature of data transfers between CPU and GPU. Data allocations on the CPU side are *pageable* by default, and the operating system can swap-out memory pages at any point. The GPU cannot access data directly from pageable CPU memory, so when a data transfer from pageable host memory to GPU memory is invoked, the CUDA driver must first allocate a temporary page-locked, or *pinned*, host array, copy the host data to the pinned array, and then transfer the data from the pinned array to device memory.

As can be seen in Figure 5.5, pinned memory is used as a staging area for transfers. We can avoid the cost of the transfer between pageable and pinned host arrays either by directly allocating CPU arrays in pinned memory, or by declaring that already allocated array should be page-locked, or pinned. Pinned memory has its downsides: if the operating system cannot swap-out pages, it



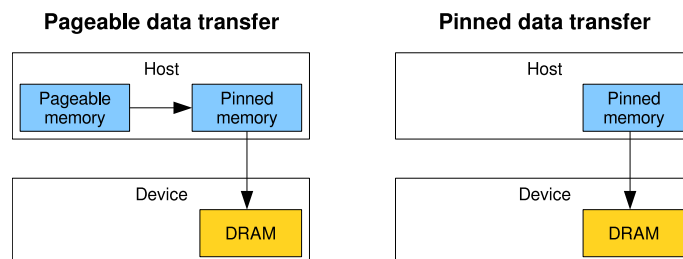


Figure 5.5: Memory copy between host and device using pageable (left) and pinned (right) memory.

may run out of resources, causing it to behave in an unexpected way. Consequently, pinned memory should not be overused. Since our programs transfer only subsets of whole arrays, pinning just those parts allows us to benefit from improved transfer speeds, as well as leave the rest of the memory allocated as pageable, in case the operating system wishes to swap it out.

Even with the use of pinned memory, the data transfer times are still large in comparison to the computation time, and leave the GPU idle during transfer. Most CUDA GPUs support simultaneous data transfers and computation via CUDA *streams*, which are sequences of operations that execute on the device in order by which they are issued by the host code. While operations within a stream are guaranteed to execute in the prescribed order, operations in different streams can be interleaved and, when possible, they can even run concurrently. Proper use of streams allows us to overlap data transfers and computation, thus mitigating the slow transfer speeds. Streams are actually always used internally within CUDA runtime, and unless a stream is specified explicitly, all operations are queued to the default stream where they appear to be synchronous due to their order of execution.

Streams are relatively simple to use when there are multiple operations on disjoint arrays queued for execution. However, when there is just one transfer to GPU, followed by computation on that data and transfer of data back from GPU, streams alone do not offer a way to improve performance. Therefore, to obtain speedup we divide the data into smaller chunks, which can then be transferred and processed independently (Figure 5.6). This idea allows us to reduce the idle time of GPU, from the time it takes to transfer the whole array to GPU and back, to the time it takes to transfer just one chunk to GPU and back, which can be significant.




Figure 5.6: Dividing data into smaller chunks and transferring them using streams reduces the GPU idle time. The idle time is reduced to the time it takes to complete one host to device (H2D) copy operation and one device to host (D2H) copy operation.

The overlap illustrated in Figure 5.6 is possible only on certain CUDA GPUs. In general, the overlap depends on the order the operations are issued to a stream, and the number of *copy engines* that the GPU has. Operations in each engine's queue are executed in the order they are issued. Therefore, to queue all operations, we have two choices, shown in Listing 5.7.

Listing 5.7: Two approaches to issuing commands to CUDA streams.

---

```

h_data = ... // Allocate data on host
d_data = ... // Allocate data on device
str = ... // Initialize 'nStreams' CUDA streams

// Asynchronous copy v1
for (i = 0; i < nStreams; i++) {
    int offset = i * streamSize;
    cudaMemcpyAsync(&d_data[offset], &h_data[offset], size, cudaMemcpyHostToDevice, str[i]);
    kernel<<<gridSize, blockSize, 0, str[i]>>>(d_data, offset);
    cudaMemcpyAsync(&h_data[offset], &d_data[offset], size, cudaMemcpyDeviceToHost, str[i]);
}

// Asynchronous copy v2
for (i = 0; i < nStreams; i++) {
    cudaMemcpyAsync(&d_data[offset], &h_data[offset], size, cudaMemcpyHostToDevice, str[i]);
}
for (i = 0; i < nStreams; i++) {
    int offset = i * streamSize;
    kernel<<<gridSize, blockSize, 0, str[i]>>>(d_data, offset);
}
for (i = 0; i < nStreams; i++) {
    int offset = i * streamSize;
    cudaMemcpyAsync(&h_data[offset], &d_data[offset], size, cudaMemcpyDeviceToHost, str[i]);
}

```

---

Depending on the number of copy engines, overlaps may be different, or not exist at all. If a GPU has two copy engines, one for transferring to and one for transferring back from the GPU, the overlap obtained by the first approach is ideal, while the second one is not. If there is one copy engine, the situation is different, with the second approach performing better. This is shown on Figure 5.7.

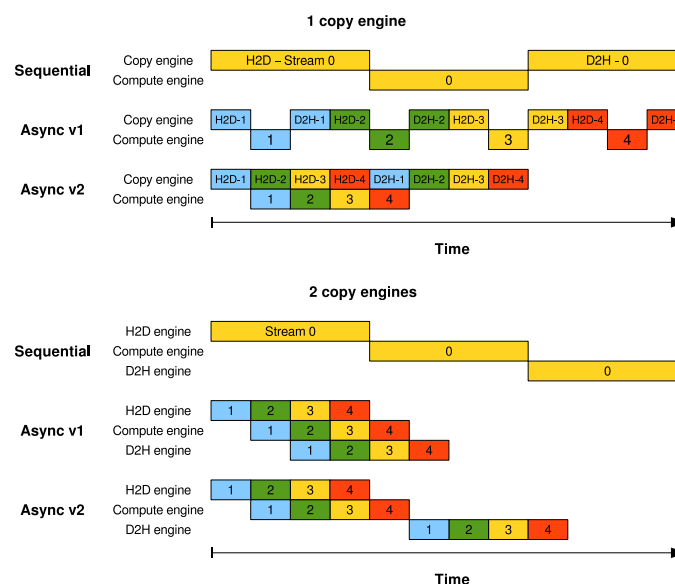


Figure 5.7: Difference between one and two copy engines for overlapping data transfer and computation.

Nvidia Tesla line of GPUs usually has two copy engines, while the consumer-oriented GeForce line has just one [66]. On newer Nvidia GPUs with Compute Capability 3.5 and higher this is not a problem, so both approaches have equal performance. However, we decided to aim for maximum compatibility and implemented both approaches, and our programs themselves select the correct approach at runtime. We created separate functions for different cases, and a function pointer to select between the two. Number of copy engines (among other things) for CUDA device in use can be obtained from CUDA runtime. With this piece of information programs choose the correct functions during initialization.

In Figures 5.6 and 5.7 we assume that the time required for transfer and computation is approximately the same. This, however, is not necessarily easy to achieve, as it depends on the size and number of chunks that the GPU will process. Users can control the size and number of chunks in our programs, and a detailed guide to choosing the best parameters is given in Chapter 8.

## Chapter 6

# Distributed memory algorithm

Algorithms presented in previous chapters target a single computer, where they use all available CPU and/or GPU resources. Next logical step is to extend them to multiple machines, i.e., clusters of computing nodes with distributed memory model, where no single process contains the whole dataset. Using multiple computing nodes can help us achieve two goals, namely faster computation on existing mesh sizes due to more computing resources being available, and availability of larger mesh sizes that cannot fit on a single computing node. In a distributed memory system, we often need to exchange data between processes running in parallel on different computing nodes. Explicit message passing is typically deployed to facilitate data exchanges and achieve parallelism.

The most popular standardized protocol for message passing paradigm is the Message Passing Interface (MPI) [67], which we use in our implementations targeting distributed memory clusters. For this development, focus was placed only on the 3D variants of our programs. We made this decision because modern computers (especially powerful cluster nodes) have enough memory for 1D or 2D calculations our programs require. An MPI algorithm with three different implementations came out as a result of this effort. The three implementations are:

1. pure CPU version built on top of shared memory algorithm from Chapter 3,
2. pure GPU version built on top of CUDA implementation of the shared memory algorithm described in Chapter 4 and
3. hybrid version based on the algorithm described in Chapter 5.

Henceforth, we refer to these implementations as OpenMP/MPI, CUDA/MPI and Hybrid/MPI, respectively. Two of the three implementations, the OpenMP/MPI and CUDA/MPI implementations, have been published in Ref. [18] and is available for download [46, 37]. The Hybrid/MPI implementation is also publicly available [37].

In the following sections, we demonstrate how the algorithms from previous chapters can be extended with a data distribution scheme to enable their execution on distributed memory systems, and how MPI can be successfully used in the implementation of these new algorithms. All three implementations we developed share a common data distribution scheme, and thus share the same general structure, so we describe them together in this chapter, emphasizing the differences in implementations where applicable. We start with the description of data decomposition among the nodes in the cluster (Section 6.1), and how the necessary transpose of data can be implemented.

In Section 6.2 we move on to the description of the changes in computation imposed by the distributed environment. Finally, in Section 6.3, we describe the changes needed to input and output components in the new environment.

## 6.1 Data distribution scheme

When we design algorithms for distributed memory systems there are two general approaches we can take. We can either focus our efforts on developing distributed algorithms, or we can dynamically redistribute the data by doing a transpose among processes in order to apply our existing algorithms on each computing node.

In the first approach, we would have to develop many novel algorithms from scratch, such as a parallel tridiagonal solver used in the functions `calclux`, `calcluy` and `calcluz`, or a parallel Fourier transform working on distributed data, and in the case of a heterogeneous architecture, further divided between CPU and GPU. This would require a significant effort, and while in the end it may provide significant performance improvement, the complexity of the implementation would be detrimental to future modifications.

The other approach, to redistribute the data, requires only development of transpose routines necessary to transform the data so that any desired dimension is local to a process. Thus, the second approach is preferred here due to its simplicity. The existing algorithms, optimized for a single machine, remain mostly unchanged and preserve all of the original logic. Conceptually, the most important changes are that we now initialize only the relevant portion of data on each process, and rely on data transposes to bring the nonlocal data to each process.

Working with 3D data allows us to have a 1D, 2D or 3D decomposition, as shown in Figure 6.1. In 1D or 2D decompositions we can arbitrarily choose the dimensions which would be split among the processes. The 3D decomposition would imply that no single dimension is local to any process, a scenario that would require us to develop distributed algorithms from scratch, which we do not pursue.




Figure 6.1: Decomposition of 3D data along one (left), two (center) and three (right) dimensions.

The 1D decomposition (slab decomposition, Figure 6.1, left), which we have already used as the data distribution method in our hybrid algorithm, is the simplest way to distribute the data among separate processes here as well. In this decomposition approach, each MPI process is assigned one or more 2D slices of data as shown in the figure. The downside of this type of decomposition is its scalability. In order to perform a transpose over slab-decomposed data, all processes communicating must exchange data, which in cases of large number of processes (over 1000), can lead to scalability

issues [68]. On the other hand, in 2D decomposition, also known as block or pencil decomposition, MPI processes are organized in a 2D array and each contains one or more 1D blocks of data as shown in Figure 6.1 (center). The block decomposition requires only between processes in a single logical row or column, which scales better.

Transposing data in both of these decomposition types can be implemented using all-to-all MPI functions, e.g., `MPI_Alltoall` or `MPI_Alltoallv`. However, 2D decomposition is more difficult to implement since the communication pattern is quite complex [69], e.g., the cost of communications is very sensitive to the orientations of blocks and their associated memory patterns. The packing and unpacking of memory buffers for the MPI library must be handled with great care for efficiency. All this is further complicated by the fact that we need two such decompositions, depending on the dimension which we need to be local to each process. In the implementation of 1D decomposition, there is an alternative way to transpose the data, via the FFTW library, whose implementation of transpose routines relies on internal mechanisms used in the computation of multidimensional FFT. This approach is simpler and has an additional benefit in that FFTW routines can operate in-place, whereas `MPI_Alltoall`-based approach can only operate out-of-place. Even for out-of-place transposes, FFTW's routine should be equal or faster, since one of the possible algorithms that FFTW uses for an out-of-place transpose is simply to call `MPI_Alltoall`.

Since PARADOX cluster used during the development has slightly more than a hundred nodes (with a total of 1696 cores), we have decided to base our algorithm and its implementations on the 1D decomposition. In order to avoid its scalability issues for large number of processes, we use MPI only for communication between nodes in a cluster, whereas all computation within a single node remains unchanged, with OpenMP, CUDA, or a combination of the two. Similarly to approach used in our hybrid algorithm, here we also decompose along the  $x$  and  $y$  direction, according to data access patterns, an example of which we show in the next section. This means that each process holds `localNx`  $\times$  `Ny`  $\times$  `Nz` amount of data (or `Nx`  $\times$  `localNy`  $\times$  `Nz`), where `localNx` = `Nx` / `nprocs` and `nprocs` is the number of MPI processes (`localNy` = `Ny` / `nprocs`). Transpose along the innermost, fastest changing dimension is not required by any computation.

With the data distributed among the MPI processes, we need an efficient transpose routine to redistribute the data along a different dimension. As mentioned before, one easy way to achieve this is to use the transpose routines from the FFTW library, which we use in the OpenMP/MPI and Hybrid/MPI implementations. Transpose interface in FFTW works in a similar way as the rest of the library, i.e., it involves creating and executing a plan (Listing 6.1). We use a special type of transpose routine from FFTW, where data is *locally* transposed. The full transpose of data would require performing an additional, local transpose operation, which we omit since we can perform all computation with locally transposed data. This means that the data after such partial transpose operation are stored structured as `Nx`  $\times$  `localNy`  $\times$  `Nz` in row-major order, instead of the expected `localNy`  $\times$  `Nx`  $\times$  `Nz` structure in row-major order. Similar procedure can be used to transpose the data back to its original layout. This way, we do not have to change the data access pattern, we only have to adjust the loop limits. To instruct FFTW to perform transpose in this way, we have to pass an additional flag when creating a transpose plan.

While it is very easy to use, FFTW transpose interface currently works only on data stored in host memory of each computing node, not the GPU memory. As a consequence, we had to implement our own transpose routines for CUDA/MPI version. We used the MPI data types and

Listing 6.1: The transpose operation implemented using FFTW, using locally transposed data layout.

---

```
// Allocate required memory
double ***tensor = alloc_double_tensor(localNx, Ny, Nz);
double ***tensor_t = alloc_double_tensor(Nx, localNy, Nz);

// Create a plan with a special flag specifying that output should be locally transposed
fftw_plan plan_tran_x = fftw_mpi_plan_many_transpose(Nx, Ny * Nz, 1, localNx, localNy *
    Nz, **tensor, **tensor_t, MPI_COMM_WORLD, FFT_MEASURE | FFTW_MPI_TRANSPOSED_OUT);
fftw_plan plan_tran_y = fftw_mpi_plan_many_transpose(Ny * Nz, Nx, 1, localNy * Nz,
    localNx, **tensor_t, **tensor, MPI_COMM_WORLD, FFT_MEASURE | FFTW_MPI_TRANSPOSED_IN);

//Execute like a regular FFTW plan
fftw_execute(plan_tran_x);
fftw_execute(plan_tran_y);
```

---

all-to-all communication (by using `MPI_Alltoall` function or a series of `MPI_Isend` and `MPI_Irecv` functions) to perform this task. We implemented the same locally transposed input/output data layout. Figure 6.2 shows the communication pattern required to perform such transpose operation. Since the result of the transpose is the same with FFTW and our own implementation, we made available the choice between the two transpose routines as a compile-time variable in Hybrid/MPI implementation, allowing users to select routine to be used.




Figure 6.2: Communication pattern of transpose operation between four processes. Note that the data after the operation are locally transposed, i.e., do not represent full transpose in general case.

Two custom MPI data types were required, representing a portion of the data which is to be exchanged in  $x$  and  $y$  dimension. Special care must be taken to account for the third dimension, which is padded in CUDA implementation and should thus be expressed in bytes and not in the actual number of elements. To achieve this, we use MPI vector type, as described in Listing 6.2.

Listing 6.2: The transpose operation implemented using MPI, using locally transposed data layout.

---

```
struct tran_params { // We use a structure to keep the data relevant to transpose
    void *orig_buf, *tran_buf; // Data buffers
    int *orig_cnts, *tran_cnts; // Counts
    int *orig_displ, *tran_displ; // Displacements
    MPI_Datatype *orig_types, *tran_types; // Custom MPI types
    int nprocs; // Number of processes
    MPI_Request *send_req, *recv_req; // Used to track async sends and recv's
};
```

---

```

struct tran_params init_transpose_double(int nprocs, long localNx, long localNy, long Ny,
    long Nz, void *send_buf, size_t send_pitch, void *recv_buf, size_t recv_pitch) {
    long i;
    MPI_Datatype tran_dtype;
    MPI_Datatype lxyz_dtype, xlyz_dtype;
    struct tran_params tran;

    // Create a custom data type representing a Ny * Nz matrix, pitched, in bytes
    MPI_Type_contiguous(localNy * (send_pitch / sizeof(double)), MPI_DOUBLE, &tran_dtype);
    MPI_Type_commit(&tran_dtype);

    // MPI vector representing the data to be sent to each process
    MPI_Type_create_hvector(localNx, 1, Ny * send_pitch, tran_dtype, &lxyz_dtype);
    MPI_Type_create_hvector(localNx, 1, localNy * recv_pitch, tran_dtype, &xlyz_dtype);
    MPI_Type_commit(&lxyz_dtype);
    MPI_Type_commit(&xlyz_dtype);

    ... // Allocate tran_params variables (*_cnts, *_displ, *_types, *_req)

    // Set the parameters for each process
    for (i = 0; i < nprocs; i++) {
        orig_cnts[i] = 1;
        tran_cnts[i] = 1;

        orig_displ[i] = i * localNy * send_pitch;
        tran_displ[i] = i * localNx * localNy * recv_pitch;

        orig_types[i] = lxyz_dtype;
        tran_types[i] = xlyz_dtype;
    }

    return tran;
}

void transpose(struct tran_params tran) {
    long i;

    // Exchange data using asynchronous send and receive
    for (i = 0; i < tran.nprocs; i++) {
        MPI_Irecv(((char *)tran.tran_buf) + tran.tran_displ[i], 1, tran.tran_types[i], i, 0,
            MPI_COMM_WORLD, &(tran.recv_req[i]));
        MPI_Isend(((char *)tran.orig_buf) + tran.orig_displ[i], 1, tran.orig_types[i], i, 0,
            MPI_COMM_WORLD, &(tran.send_req[i]));
    }
    for (i = 0; i < tran.nprocs; i++) {
        MPI_Wait(&(tran.recv_req[i]), MPI_STATUSES_IGNORE);
        MPI_Wait(&(tran.send_req[i]), MPI_STATUSES_IGNORE);
    }

    // Alternatively, we can use MPI_Alltoallw
    //MPI_Alltoallw(tran.orig_buf, tran.orig_cnts, tran.orig_displ, tran.orig_types,
        tran.tran_buf, tran.tran_cnts, tran.tran_displ, tran.tran_types, MPI_COMM_WORLD);
}

// To initialize the transpose operation
d_tensor = alloc_double_tensor_device(localNx, Ny, Nz);
d_tensor_t = alloc_double_tensor_device(Nx, localNy, Nz);
tran_tensor = init_transpose_double(nprocs, localNx, localNy, Ny, Nz, d_tensor.ptr,
    d_tensor.pitch, d_tensor_t.ptr, d_tensor_t.pitch);
transpose(tran_tensor);

```

---



An interesting property of the CUDA/MPI implementation is that the transpose routine is invoked on data residing in GPU memory, as opposed to the other two distributed memory implementations. This is possible due to the functionality available in some of the modern, CUDA-aware MPI implementations, where MPI runtime can directly access data in GPU memory without the need for copying to the host memory. Otherwise, we would need to copy the data to host memory before any MPI routine is called, a task known to be a bottleneck, which would make CUDA/MPI implementation far less usable. A list of CUDA-aware MPI implementations is available in Ref. [70]. During testing and development of the three distributed memory implementations, we have used Open MPI library [71], which can be compiled with support for CUDA.

## 6.2 Distributing the computation over multiple processes

With our data distribution scheme explained in the previous section, we can outline the changes required to the computation functions. We first have to ensure that each process works with a different portion of the data, or rather, initialize equal, non-overlapping portions of data on each process. This requires that the number of discretization points of each dimension that will be distributed must be divisible by the number of processes. By introducing additional processing logic it is possible to support the distribution of non-equal portions of data, however, we decided against such implementation since it leads to unbalanced computing load. Note that this is not a significant restriction because users can always adapt and select slightly larger or smaller, but evenly divisible mesh.

Data initialization begins by calculating the amount of data each process will get. Based on that we calculate each process' logical offset into the data, so we can initialize only the relevant portion. Initially, we distribute the data along the  $x$  dimension and perform a transpose operation whenever the data need to be local in that dimension. Listing 6.3 shows how each process can identify and initialize only the portion of data that logically belongs to it.

Listing 6.3: Example of initialization of the relevant portion of data on each process.

---

```

// Get the rank of the current process and total number of processes
MPI_Comm_rank(MPI_COMM_WORLD, &rank);
MPI_Comm_size(MPI_COMM_WORLD, &nprocs);
// Compute offsets
long localNx = Nx / nprocs; long offsetNx = rank * localNx;
// Allocate memory for data
psi = alloc_double_tensor(localNx, Ny, Nz);
// Example of initialization of relevant portion of wave function
for (i = 0; i < Nx; i++) {
    x[i] = (i - Nx / 2) * dx; x2[i] = x[i] * x[i];
}
...
for (i = 0; i < localNx; i++) {
    for (j = 0; j < Ny; j++) {
        for (k = 0; k < Nz; k++) {
            tmp = exp(- 0.5 * (vgamma * x2[offsetNx + i] + vnu * y2[j] + vlambada * z2[k]));
            psi[i][j][k] = tmp / cpsi;
        }
    }
}
}

```

---

Now that each process has initialized the relevant portion of the data, the only change to the computation functions is the adjustment of the exit clause (limit) of the outermost loop. Computation in the inner loops that require other two dimensions to be local remains unchanged. This means that functions `calcnu`, `calcluy` and `calcluz` do not require any changes other than modifications of the loop limits, as their relevant portion of data is already local. Function `calclux`, which requires  $x$  dimension to be local, is changed so that we first transpose the data, afterward perform the computation as before, and then transpose the data back to original layout (Listing 6.4).

Listing 6.4: Pseudocode of the changed `calclux` function.

---

```
// Using FFTW transpose for OpenMP/MPI implementation
void calclux(...) {
    ...
    fftw_execute(plan_transpose_x); // Transpose 'psi' and place the result in 'psi_t'
    ... // Compute calclux using 'psi_t', with the outer loop limit set at 'localNy'
    fftw_execute(plan_transpose_y); // Transpose 'psi_t' back to original 'psi'
}

// Using custom MPI transpose for Hybrid/MPI and CUDA/MPI implementations
void calclux(...) {
    ...
    transpose(tran_psi);
    ...
    transpose_back(tran_psi);
}

```

---

This is also a common pattern for other functions that require  $x$  dimension to be local, namely, parts of computation of chemical potential and energy, as well as RMS size. These functions also require communication between processes to gather results of computation at *root* process (usually the process with the lowest identification number, i.e., 0), which post-processes them (if needed) and writes the summary output. Listing 6.5 demonstrates this pattern.

Listing 6.5: Pseudocode of the data gathering pattern employed by `cacnorm`, `cacmuen` and `cacrms` functions.

---

```
void cacnorm(...) {
    ...// Place the temporary data in 'tmpx' as before
    void *sendbuf = (rank == 0) ? MPI_IN_PLACE : *tmpx;
    MPI_Gather(sendbuf, localNx, MPI_DOUBLE, *tmpx, localNx, MPI_DOUBLE, 0, MPI_COMM_WORLD);

    if (rank == 0) {
        *norm = sqrt(simpint(dx, *tmpx, Nx));
    }

    MPI_Bcast(norm, 1, MPI_DOUBLE, 0, MPI_COMM_WORLD);
    ...
}

```

---

The final piece is the computation of Fourier transforms in a distributed memory environment. Fourier transform for OpenMP/MPI implementation is a straightforward task, because FFTW library supports MPI, and uses the same slab decomposition. FFTW's MPI transform algorithms work by first computing transforms of the data local to each process, then by globally transposing the data to redistribute them among the processes and performing a transform of the new data

local to each process, and optionally transposing back. This sequence corresponds to the *row-column* algorithm, which we have already seen in our hybrid algorithm in Chapter 5. An example of the row-column algorithm in a distributed memory system for 2D Fourier transform is illustrated in Figure 6.3. If the final data transpose is omitted, the transform results are produced in transposed order, distributed over  $y$  dimension. This has a significant performance advantage, at the cost of having to change the data layout of the Fourier transform of the dipolar potential (variable `potdd`). This was a simple matter of allocating memory in a different layout to fit the transposed order and reordering the loops in the initialization of the dipolar term.




Figure 6.3: Illustration of four stages of the row-column FFT algorithm.

In the case of CUDA/MPI and Hybrid/MPI implementations though, it is not possible to use FFTW’s MPI transforms. The cuFFT library does not have support for MPI, and our Hybrid/MPI implementation relies on the heterogeneous Fourier transform we described in Section 5.2. However, we can still use transpose routines in conjunction with the row-column algorithm. To mimic the behavior of FFTW’s MPI transforms with transposed order, we rely on the same transpose routines we used earlier, either ones available in FFTW for the Hybrid/MPI implementation, or the ones we developed for the CUDA/MPI version. Listing 6.6 illustrates how we perform FFT using FFTW’s MPI transforms, using cuFFT and our own transpose routines CUDA/MPI, as well as using cuFFT and FFTW with FFTW transposes in Hybrid/MPI implementation.

Listing 6.6: Three different implementations of distributed FFT: using FFTW, cuFFT and both. We show only the *forward* FFT, the *inverse* transform that follows is done analogously.

```
// FFT plan creation and usage with FFTW MPI interface, used in OpenMP/MPI implementation
fftw_plan plan_fw = fftw_mpi_plan_many_dft_r2c(3, n, 1, localNx, localNy, psidd2[0][0],
        psidd2fft, MPI_COMM_WORLD, FFT_MEASURE | FFTW_MPI_TRANSPOSED_OUT);
...
fftw_execute(plan_fw); // Executed as any other plan

// FFT with cuFFT , used in CUDA/MPI implementation
int dim = 2;
long long nfr[] = {Ny, Nz};
long long howmany = localNx;
long long idist = Ny * (psidd2.pitch / sizeof(cufftDoubleReal))
long long odist = Ny * (psidd2.pitch / sizeof(cufftDoubleComplex));
long long istride = 1, ostride = 1;
long long inembedfr[] = {Ny, psidd2.pitch / sizeof(cufftDoubleReal)};
long long onembedfr[] = {Ny, psidd2.pitch / sizeof(cufftDoubleComplex)};

cufftMakePlanMany64(plan_fwr, dim, nfr, inembedfr, istride, idist, onembedfr, ostride,
        odist, CUFFT_D2Z, howmany, &ws_fwr);

dim = 1;
long long nfc[] = {Nx};
```

```

howmany = localNy * (Nz/2 + 1);
idist = 1; odist = 1;
istride = localNy * psidd2_t.pitch / sizeof(cufftDoubleComplex);
ostride = localNy * psidd2_t.pitch / sizeof(cufftDoubleComplex);
long long inembedfc[] = {Nx}, onembedfc[] = {Nx};
cufftMakePlanMany64(plan_fwc, dim, nfc, inembedfc, istride, idist, onembedfc, ostride,
    odist, CUFFT_Z2Z, howmany, &ws_fwr);
// Create a transpose plan
tran_dd2 = init_transpose_complex(nprocs, localNx, localNy, Ny, Nz/2 + 1, psidd2.ptr,
    psidd2.pitch, psidd2_t.ptr, psidd2_t.pitch);
...
cufftExecD2Z(plan_fwr, (cufftDoubleReal *) psidd2.ptr, (cufftDoubleComplex *) psidd2.ptr);
transpose(tran_dd2);
cufftExecZ2Z(plan_fwc, (cufftDoubleComplex *) psidd2_t.ptr, (cufftDoubleComplex *)
    psidd2_t.ptr, CUFFT_FORWARD);

// FFT using FFTW on host and cuFFT on device, as used in Hybrid/MPI implementation
... // Initialize FFT plans on host and device as in Listing 5.5
// Make transpose plan
plan_tr_psidd2_x = fftw_mpi_plan_many_transpose(Nx, Ny * (Nz/2 + 1), 2, localNx, localNy *
    (Nz/2 + 1), (double *) psidd2fft, (double *) psidd2fft, MPI_COMM_WORLD, FFTW_ESTIMATE
    | FFTW_MPI_TRANSPOSED_OUT);

#pragma omp parallel private(threadid) num_threads(2)
{
    threadid = omp_get_thread_num();
    if (threadid == 0) {
        ... // Copy data to GPU, perform FFT along Ny and Nz, and sync data back
    }
    if (threadid == 1 || omp_get_num_threads() != 2) {
        ... // Perform FFT on host, using FFTW without MPI
    }
}

fftw_execute(plan_tr_psidd2_x); // Transpose data using FFTW

#pragma omp parallel private(threadid) num_threads(2)
{
    threadid = omp_get_thread_num();
    if (threadid == 0) {
        ... // Copy data to GPU, perform FFT along Nx, without syncing back
    }
    if (threadid == 1 || omp_get_num_threads() != 2) {
        ... // Perform FFT on host, using FFTW without MPI
    }
}

```

---

The patterns we described in this section are used throughout the three implementations. Since their use is relatively straightforward, for brevity we will not go through all the individual changes.

### 6.3 Improvements of input and output operations

With the introduction of multiple computing nodes, each working on a separate set of data, the I/O components of our programs needed to be adjusted accordingly. Note that the *input* mentioned in this section is the output of imaginary-time propagation programs, used as input for real-time propagation. Input files containing parameters (passed via `-i` argument during program startup) are very small, and identical for every process, thus there is no need for handling them in any special

way.

A naive approach to reading input data in a cluster environment is to let one process read all data, and then distribute them to all other processes. This is inefficient, since it relies on I/O throughput of a single node to perform the read, and therefore scales badly. Additionally, the input may be very large, larger than the available memory of a single node, requiring additional processing logic to read and split the data. As far as output is concerned, a simple approach would be to let each process write its own portion of the data in a separate file, ensuring only that processes do not overwrite each other's output. While the simplicity of this approach to I/O is appealing, it may result in potentially many files being created, which complicates analysis of the results, forcing users to rely on some external process to collect, combine and otherwise process the output. Alternatively, we could send all data to one process, which would then write all output data aggregated. This approach again has scalability issues similar to the naive approach to reading input data, as it would limit the mesh size to the size of the single node's memory. Efforts to remedy this, for example by sequentially obtaining and writing data from other processes would further reduce the performance.

Another limiting factor is the sheer amount of data written. Previous implementations used a textual output, which is simple to analyze and process. With larger mesh sizes, however this output grows to become very large. Textual representation of floating-point values stored in memory that we used before takes up more space than the equivalent binary representation, and we can reduce the size of textual representation only by decreasing precision.

In light of these limitations, we have decided to use MPI for I/O as well, and to switch from textual to binary representation of data. While it is possible to write textual data with MPI, by calculating the offset based on the amount of bytes the text is going to occupy, this solution is not preferred due to the reasons mentioned above. Textual output could still be obtained from binary representation using `hexdump` command, available in most UNIX-like operating systems. We have included an example script that will convert binary output to a textual one as part of the documentation of our MPI programs.

MPI has a simple mechanism for writing to and reading from files based on the offset into data at which each process will perform its work (Figure 6.4). Behind the scenes, MPI runtime will choose an optimal way to perform I/O.




Figure 6.4: I/O using MPI, where each process accesses a portion of the file at a different offset.

Reading or writing data with MPI in binary form is straightforward, as illustrated by several output functions shown in Listing 6.7. We did not use the approach of single-node implementations, where small amount of data is written to a file in each iteration of the loop. Instead, we now use temporary variables to prepare data for output, and write them all at once. In case of very large outputs, like 3D density profiles, MPI buffers may be too small to write all data at once, so in that case we can simply write a 2D slice in a loop. Writing output this way has obvious performance benefits, and the amount of additional memory required for the temporary variables is negligible in a distributed memory system.

Listing 6.7: Examples of output functions implemented using MPI I/O.

---

```

// Write integrated XY density
void outdenxy(double ***psi, double ***outxy, double *tmpz, MPI_File file) {
    long i, j, k;
    // Compute the appropriate offset for each process
    MPI_Offset fileoffset= rank * 3 * sizeof(double) * (localNx / outstpx) * (Ny / outstpy);

    for (i = 0; i < localNx; i += outstpx) {
        for (j = 0; j < Ny; j += outstpy) {
            for (k = 0; k < Nz; k++) {
                tmpz[k] = psi[i][j][k] * psi[i][j][k];
            }
            // Prepare output data
            outxy[i / outstpx][j / outstpy][0] = x[offsetNx + i];
            outxy[i / outstpx][j / outstpy][1] = y[j];
            outxy[i / outstpx][j / outstpy][2] = simpint(dz, tmpz, Nz);
        }
    }
    // Write all prepared data at once, leaving MPI to choose the best strategy for I/O
    MPI_File_write_at_all(file, fileoffset, **outxy, (localNx / outstpx) * (Ny / outstpy) *
        3, MPI_DOUBLE, MPI_STATUS_IGNORE);
}

// Write integrated Y density
void outdeny(double ***psi_t, double **outy, double *tmpx, double *tmpz, MPI_File file) {
    long i, j, k;
    MPI_Offsetfileoffset = rank * 2 * sizeof(double) * (localNy / outstpy);
    fftw_execute(plan_transpose_x); // With a transpose we can make required data local

    for (j = 0; j < localNy; j += outstpy) {
        for (i = 0; i < Nx; i++) {
            for (k = 0; k < Nz; k++) {
                tmpz[k] = psi_t[i][j][k] * psi_t[i][j][k];
            }
            tmpx[i] = simpint(dz, tmpz, Nz);
        }
        outy[j / outstpy][0] = y[offsetNy + j];
        outy[j / outstpy][1] = simpint(dx, tmpx, Nx);
    }
    MPI_File_write_at_all(file, fileoffset, *outy, (localNy / outstpy) * 2, MPI_DOUBLE,
        MPI_STATUS_IGNORE);
}

// Write final density
void outdenxyz(double ***psi, double ***outxyz, MPI_File file) {
    long i, j, k;
    MPI_Offsetfileoffset = rank * sizeof(double) * localNx * Ny * Nz;

    for (i = 0; i < localNx; i += outstpx) {
        for (j = 0; j < Ny; j += outstpy) {
            for (k = 0; k < Nz; k += outstpz) {
                outxyz[0][j][k] = psi[i][j][k] * psi[i][j][k];
            }
        }
    }
    // MPI I/O returns error if the array is too large, so we write Ny * Nz at a time.
    MPI_File_write_at_all(file, fileoffset, **outxyz, (Ny / outstpy) * (Nz / outstpz),
        MPI_DOUBLE, MPI_STATUS_IGNORE);
    fileoffset += (Ny / outstpy) * (Nz / outstpz) * sizeof(double);
}
}

```

```
// Reading the output of imaginary-time programs in real-time programs
void readdenxyz(double complex ***psi, double *tmpz, MPI_File inputfile) {
    long i, j, k;
    MPI_Offset fileoffset = rank * sizeof(double) * localNx * Ny * Nz;

    for (i = 0; i < localNx; i++) {
        for (j = 0; j < Ny; j++) {
            // Read whole row at a time, faster then each element individually
            MPI_File_read_at(file, fileoffset, tmpz, Nz, MPI_DOUBLE, MPI_STATUS_IGNORE);
            for (k = 0; k < Nz; k++) {
                psi[i][j][k] = sqrt(tmpz[k]);
            }
            fileoffset += Nz * sizeof(double);
        }
    }
    MPI_File_close(&file);
}
```

---

Note that in some output functions we require data along  $x$  direction to be local to each process. We use the same techniques described in the previous two sections to achieve this, therefore we will not go into those details again.

## Chapter 7

# Interacting with simulation

Programs presented in previous chapters are not interactive. That is, programs are parametrized via an input file and produce a summary output, and optionally write the density profiles of the wave function to disk at predefined intervals. This output can then be analyzed further or visualized with some visualization tool. While this workflow is sufficient for many use-cases, sometimes it is useful to analyze and visualize data as it is being generated, as well as to interact with the simulation (for example, to start, pause, stop, or restart the simulation) and modify parameters it is running. This type of visualization is called *in-situ visualization*, meaning that it is applied on datasets in-place where they are computed (host memory of a computer), rather than writing files to disk, and then reading them back into a visualization tool. The ability to steer the visualized simulation, e.g., by stopping the simulation if we are not satisfied with the intermediate results, or modifying the parameters of a running simulation, can help reduce resource consumption and reduce the time it takes to obtain the desired results. In situ visualization can also be a valuable research tool, allowing users to investigate phenomena that appear during the simulation, even if not anticipated previously.

Creating a visualization component and embedding it in our programs would be an arduous task. The preferred approach therefore is to link a visualization component of an existing tool that supports in-situ visualization to our programs. There are plenty of such tools available, but two have garnered a lot of attention in the scientific community: ParaView [72] and VisIt [73]. Both are powerful and mature products with similar functionality, distributed under an open-source license. ParaView relies on a library called `Catalyst` to provide in-situ visualization support, which is a reduced version of the ParaView's server component, allowing a program to share data with ParaView for visualization. Similar capabilities are available within VisIt with the `libsim` library. In addition to data sharing between the simulation and VisIt, `libsim` also enables the opposite flow of information, sending data from the client to the simulation, enabling simulation steering. We chose to use `libsim` and VisIt for our visualizations and steering since it offers a simple C API, which was easier to include in existing programs.

In the following sections, we describe how we used `libsim` component of VisIt to provide visualization of data from our programs and enable simulation steering.



## 7.1 Visualization of data

VisIt is a free, open-source, platform-independent, distributed, parallel, visualization tool for visualizing data defined on 2D and 3D, structured and unstructured meshes, as well as plots of 1D data. VisIt uses a client-server model (viewer-compute engine), where the server is parallelized and can be distributed on a cluster. This architecture allows for remote visualization, allowing us to perform large-scale visualization on a powerful computer cluster while running viewer component on an ordinary desktop computer connected to a cluster. VisIt's compute engine handles coordination between distributed processes and, given the information on how the data are distributed, can reassemble and process them for visualization (e.g., generate a plot and send it to the viewer component). Figure 7.1 depicts VisIt's compute engine accessing data in parallel.

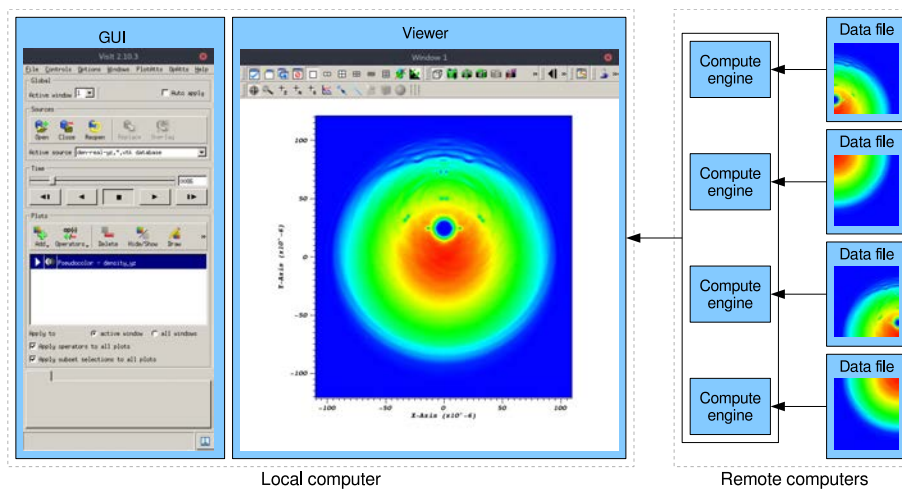


Figure 7.1: Illustration of VisIt's components. VisIt's compute engine reads data files in parallel and sends data to the viewer component, which user controls via GUI.

A scenario where processes of a program write separate data files is illustrated in Figure 7.1. Data created in such way are analyzed and visualized as a post-processing step. In order to make the simulation interactive, and to remove the need for writing data to disk, we have to allow VisIt's compute engine access to our data directly in memory. To do this, we rely on VisIt's `libsim` library. This library can be inserted into a program to make the simulation behave like a VisIt compute engine. The `libsim` library and the accompanying data-access code gives VisIt's data-processing routines access to the simulation's calculated data without the need for the simulation to write files to disk (Figure 7.2). Simulation instrumented in this way can begin its processing while periodically listening for connections from an instance of VisIt using `libsim`. When an incoming connection is detected, `libsim` loads its dynamic runtime library that contains the VisIt compute engine's data-processing functions. Once the runtime is loaded, the simulation connects to VisIt's viewer, which can make requests for plots as if the simulation was an ordinary VisIt compute engine.

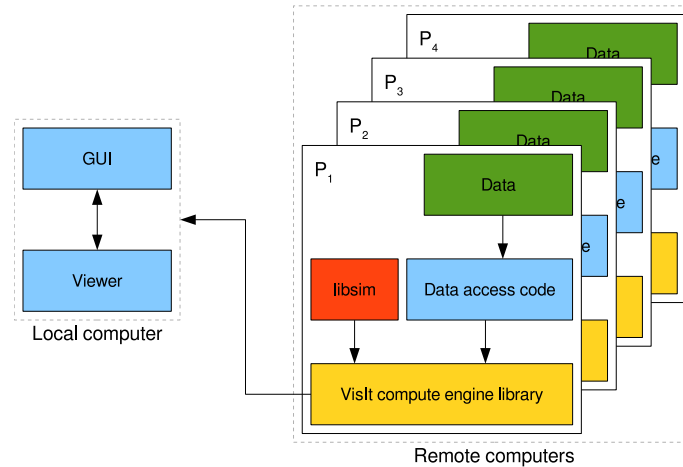


Figure 7.2: Schematic depiction of VisIt getting data from an instrumented parallel simulation.

User's requests for data from VisIt's viewer are passed to the simulation, which has to provide data through a set of callback functions. These callback functions are the main part of the integration with VisIt, as they provide information on what the simulation is doing, its current state, what data are available, how they can be represented, and how they are distributed among the processes. In addition to the callback functions, we also have to change the flow of our program to integrate `libsim`. General flow of a simulation with `libsim` integrated is illustrated in Figure 7.3.

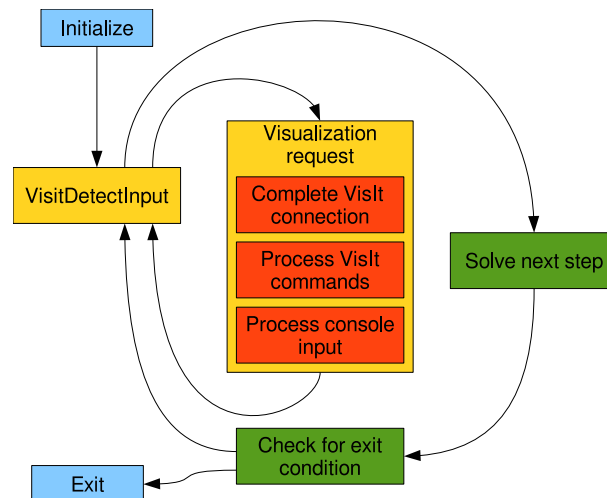


Figure 7.3: Main simulation loop with `libsim` integrated.

Restructuring programs to integrate `libsim` can be done in many ways, however we selected a pattern that VisIt uses in its examples and which is also recommended by VisIt's developers. Focus was placed on 3D variants of the programs, specifically, the MPI-parallelized versions described in Chapter 6, however other versions of programs would also be suitable candidates for integration with VisIt. We note that the necessary changes were not difficult to implement. While the changes are extensive, the only conceptually major change is the restructuring of the main loop to support connections from VisIt. Much of the restructured code can be shared between real- and imaginary-time programs, so we used a single implementation for most of the functions commu-

nicating with `libsim`. The differences were handled with conditional compilation, allowing us to enable the different functionality via a compiler switch. Listing 7.1 shows the flow of simulation after restructuring.

Listing 7.1: The flow of simulation in 3D programs after integration with `libsim`.

---

```

void main(...) {
    ... // Initialize variables and simulation data
    SimData *sim = ...; // Structure containing simulation state and variables
    SimMainLoop(sim);
    ... // Free resources and clean up
}

void SimMainLoop(SimData *sim) {
    int blocking, vstate, err = 0;

    // Register VisIt commands (start/stop, pause...), variables, on-demand output...
    SimRegisterCommands(sim);
    do {
        // Determines if the simulation should wait for user input or run anyway
        blocking = (sim->running == SIM_STOPPED) ? 1 : 0;

        if(sim->par_rank == 0) {
            vstate = VisItDetectInput(blocking, fileno(stdin)); // Detect user input, if any
        }
        MPI_Bcast(&vstate, 1, MPI_INT, 0, MPI_COMM_WORLD);

        switch(vstate) {
        case 0:
            SimNextStep(sim); // If there was no user input, proceed with the next iteration
            break;
        case 1:
            // Initial case, user attempts to connect to VisIt
            if(VisItAttemptToCompleteConnection() == VISIT_OKAY) {
                // Set command callbacks
                VisItSetCommandCallback(SimCommandCallback, (void*)sim);
                VisItSetSlaveProcessCallback2(SlaveProcessCallback, (void*)sim);
                // Set simulation metadata and variable callback
                VisItSetGetMetaData(SimGetMetaData, (void*)sim);
                VisItSetGetMesh(SimGetMesh, (void*)sim);
                VisItSetGetVariable(SimGetVariable, (void*)sim);
                VisItSetGetCurve(SimGetCurve, (void*) sim);
                VisItSetGetDomainList(SimGetDomainList, (void*)sim);
            } else {
                ... // Handle connection error
            }
            break;
        case 2:
            // User has sent an UI command, decode and process it
            if(!ProcessVisItCommand(sim)) {
                VisItDisconnect(); // Disconnect in case of failure
                sim->done = 1;
            }
            break;
        case 3:
            // User has sent an console command, decode and process it
            SimProcessConsoleCommand(sim);
            break;
        default:
            ... // Other states should not happen, handle errors if they do
            break;
        }
    }
}

```

```

    }
} while (!sim->done && err == 0); // Loop until simulation is stopped or reaches the end
}

```

---

As part of the restructuring, we introduced a structure called `SimData` that represents the simulation's global state, has pointers to all important variables, and is used to pass information inside callback functions. We define several callback functions that return information about data in our programs and register them with VisIt by calling the appropriate function from `VisItSetGet*` set.

First callback function to be invoked is `VisItSetGetMetaData`, which provides metadata about the simulation. This is the function invoked when VisIt connects to the simulation, and the obtained metadata is used to determine which meshes and variables are available. The associated callback function must return a handle to a `SimulationMetaData` object that contains lists of metadata objects. Listing 7.2 demonstrates creation of some of the metadata used. For 3D programs, we define one 3D mesh, and three 2D meshes (every combination of two axes). These meshes represent the base on top of which the corresponding 2D or 3D variables will be visualized. The 3D mesh is used as a base for the variable containing the absolute square of the wave function ( $|\psi|^2$ ), while the 2D meshes are used for the 2D integrated density profiles the programs generate. VisIt supports numerous types of meshes for many different use-cases. Due to the nature of the data in our programs, a *rectilinear* mesh is the best choice for our data, as it directly corresponds to the spatial grid we use. Rectilinear meshes are specified by lists of coordinate values for each axis, which map to spatial grid variables we already defined in our programs. For variables representing 1D arrays, there is no concept of 1D mesh, rather they are defined as *curves* in VisIt, and are visualized as ordinary plots. Using meshes and curves we have augmented the output system of our programs, which can now also send data to VisIt for visualization instead of just writing the computed output to a file.

Listing 7.2: The portion of the metadata exported from the 3D program to VisIt.

---

```

visit_handle SimGetMetaData(void *cbdata) {
    visit_handle md = VISIT_INVALID_HANDLE;
    SimData *sim = (SimData *)cbdata;
    // Allocate main SimulationMetaData object
    if(VisIt_SimulationMetaData_alloc(&md) == VISIT_OKAY) {
        // Define empty handles to meshes, variables and curves
        visit_handle xyzmd = VISIT_INVALID_HANDLE;
        visit_handle xynd = VISIT_INVALID_HANDLE;
        visit_handle xzmd = VISIT_INVALID_HANDLE;
#ifdef REAL
        visit_handle psimd = VISIT_INVALID_HANDLE;
#endif
        visit_handle psi2md = VISIT_INVALID_HANDLE;
        ...
        VisIt_SimulationMetaData_setMode(md, (sim->running == SIM_STOPPED) ?
            VISIT_SIMMODE_STOPPED : VISIT_SIMMODE_RUNNING);
        VisIt_SimulationMetaData_setCycleTime(md, sim->step, sim->time);
        // 3D mesh (XYZ) metadata
        if(VisIt_MeshMetaData_alloc(&xyzmd) == VISIT_OKAY) {
            VisIt_MeshMetaData_setName(xyzmd, "xyz_mesh");
            VisIt_MeshMetaData_setMeshType(xyzmd, VISIT_MESHTYPE_RECTILINEAR);
            VisIt_MeshMetaData_setTopologicalDimension(xyzmd, 3);
            VisIt_MeshMetaData_setSpatialDimension(xyzmd, 3);

```

```

    VisIt_MeshMetaData_setNumDomains(xyzmd, sim->par_size);

    VisIt_MeshMetaData_setXLabel(xyzmd, "z");
    VisIt_MeshMetaData_setYLabel(xyzmd, "y");
    VisIt_MeshMetaData_setZLabel(xyzmd, "x");

    VisIt_SimulationMetaData_addMesh(md, xyzmd);
}
// 2D mesh (XY) metadata
if(VisIt_MeshMetaData_alloc(&xymd) == VISIT_OKAY) {
    VisIt_MeshMetaData_setName(xymd, "xy_mesh");
    VisIt_MeshMetaData_setMeshType(xymd, VISIT_MESHTYPE_RECTILINEAR);
    VisIt_MeshMetaData_setTopologicalDimension(xymd, 2);
    VisIt_MeshMetaData_setSpatialDimension(xymd, 2);
    VisIt_MeshMetaData_setNumDomains(xymd, sim->par_size);

    VisIt_MeshMetaData_setXLabel(xymd, "y");
    VisIt_MeshMetaData_setYLabel(xymd, "x");

    VisIt_SimulationMetaData_addMesh(md, xymd);
}
... // Remaining 2D meshes
#endif REAL
// psi variable
if(VisIt_VariableMetaData_alloc(&psimd) == VISIT_OKAY) {
    VisIt_VariableMetaData_setName(psimd, "psi");
    VisIt_VariableMetaData_setMeshName(psimd, "xyz_mesh");
    VisIt_VariableMetaData_setType(psimd, VISIT_VARTYPE_SCALAR);
    VisIt_VariableMetaData_setCentering(psimd, VISIT_VARCENTERING_NODE);

    VisIt_SimulationMetaData_addVariable(md, psimd);
}
#endif
// psi2 variable
if(VisIt_VariableMetaData_alloc(&psi2md) == VISIT_OKAY) {
    VisIt_VariableMetaData_setName(psi2md, "psi^2");
    VisIt_VariableMetaData_setMeshName(psi2md, "xyz_mesh");
    VisIt_VariableMetaData_setType(psi2md, VISIT_VARTYPE_SCALAR);
    VisIt_VariableMetaData_setCentering(psi2md, VISIT_VARCENTERING_NODE);

    VisIt_SimulationMetaData_addVariable(md, psi2md);
}
// Density (XY) variable
if(VisIt_VariableMetaData_alloc(&denxymd) == VISIT_OKAY) {
    VisIt_VariableMetaData_setName(denxymd, "density_xy");
    VisIt_VariableMetaData_setMeshName(denxymd, "xy_mesh");
    VisIt_VariableMetaData_setType(denxymd, VISIT_VARTYPE_SCALAR);
    VisIt_VariableMetaData_setCentering(denxymd, VISIT_VARCENTERING_NODE);

    VisIt_SimulationMetaData_addVariable(md, denxymd);
}
... // Other 2D variables
/* Density (X) curve */
if(VisIt_CurveMetaData_alloc(&denxmd) == VISIT_OKAY) {
    VisIt_CurveMetaData_setName(denxmd, "density_x");
    VisIt_SimulationMetaData_addCurve(md, denxmd);
}
... // Remaining curves
// Custom commands
visit_handle cmd = VISIT_INVALID_HANDLE;
if(VisIt_CommandMetaData_alloc(&cmd) == VISIT_OKAY) {

```

```

    VisIt_CommandMetaData_setName(cmd, "START/PAUSE");
    VisIt_SimulationMetaData_addGenericCommand(md, cmd);
  }
  ... // Remaining commands
}

return md;
}

```

Once VisIt's GUI is populated by the metadata obtained, the user can request the visualization of variables from the simulation, which will call the associated callback. For example, if one wants to visualize a variable labeled  $\psi^2$ , one selects the appropriate command from the VisIt's GUI, and in the background, VisIt will use `libsim` to call the `SimGetMesh` (callback function registered in Listing 7.1) to get the info about the 3D mesh, then `SimGetVariable` to get the actual data, and other callback functions as needed. Listing 7.3 shows the portion of these callback functions. As part of the callback mechanism, functions receive the string with the name of the desired object (mesh, variable, curve, ...), which can be used to provide different data. In this way we can define a single function to return all possible meshes or variables.

Listing 7.3: Portions of the mesh, variable and curve callbacks.

```

visit_handle SimGetMesh(int domain, const char *name, void *cbdata) {
  visit_handle h = VISIT_INVALID_HANDLE;
  visit_handle hxc, hyc, hzc;
  SimData *sim = (SimData *)cbdata;

  if(strcmp(name, "xyz_mesh") == 0) {
    if(VisIt_RectilinearMesh_alloc(&h) != VISIT_ERROR) {
      VisIt_VariableData_alloc(&hxc);
      VisIt_VariableData_alloc(&hyc);
      VisIt_VariableData_alloc(&hzc);

      VisIt_VariableData_setDataD(hxc, VISIT_OWNER_SIM, 1, sim->var.localNx /
        sim->out.outstpx + sim->ghost, sim->var.xstp + (sim->var.offsetNx /
        sim->out.outstpx));
      VisIt_VariableData_setDataD(hyc, VISIT_OWNER_SIM, 1, sim->var.Ny /
        sim->out.outstpy, sim->var.ystp);
      VisIt_VariableData_setDataD(hzc, VISIT_OWNER_SIM, 1, sim->var.Nz /
        sim->out.outstpz, sim->var.zstp);
      VisIt_RectilinearMesh_setCoordsXYZ(h, hzc, hyc, hxc); // We use inverted order

      return h;
    }
  }
  ... // Remaining meshes
  return h;
}

visit_handle SimGetVariable(int domain, const char *name, void *cbdata) {
  visit_handle h = VISIT_INVALID_HANDLE;
  long ntuples;
  SimData *sim = (SimData *)cbdata;

  if(strcmp(name, "psi^2") == 0) {
    VisIt_VariableData_alloc(&h);
    ntuples = (sim->var.localNx / sim->out.outstpx + sim->ghost) * (sim->var.Ny /
      sim->out.outstpy) * (sim->var.Nz / sim->out.outstpz);
    VisIt_VariableData_setDataD(h, VISIT_OWNER_SIM, 1, ntuples, **(sim->var.psi2));
  }
}

```

```

    return h;
}
... // Remining variables
return h;
}

visit_handle SimGetCurve(const char *name, void *cbdata) {
    visit_handle h = VISIT_INVALID_HANDLE;
    SimData *sim = (SimData *)cbdata;

    // Note that only rank 0 calls this function
    if(strcmp(name, "density_x") == 0) {
        if(VisIt_CurveData_alloc(&h) != VISIT_ERROR) {
            visit_handle hxc, hyc;
            VisIt_VariableData_alloc(&hxc);
            VisIt_VariableData_alloc(&hyc);

            VisIt_VariableData_setDataD(hxc, VISIT_OWNER_SIM, 1, sim->var.Nx /
                sim->out.outstpx, sim->var.xstp);
            VisIt_VariableData_setDataD(hyc, VISIT_OWNER_SIM, 1, sim->var.Nx /
                sim->out.outstpx, sim->var.denx);
            VisIt_CurveData_setCoordsXY(h, hxc, hyc);

            return h;
        }
    }
    ... // Remaining curves
    return h;
}

```

In our simulations, we distribute data among multiple processes, and this information must also be passed via `libsim`. VisIt supports working with distributed data through the concept of *domains*. A domain is a unit of work assigned to a given process that corresponds to a portion of the mesh. By decomposing the meshes into multiple domains, VisIt is able to handle large-scale simulations. Therefore, in a distributed simulation, we have to provide VisIt with information on how domains abut, via another callback function. Additionally, each domain needs information on its adjacent domains, to ensure continuity between domains. This information is provided via *ghost nodes*. There are a few ways to provide information on ghost nodes, in our case the simplest way was to extend the data on each process to include the first spatial points of adjacent processes. The overlap of data obtained in this way is designated as a ghost node between adjacent domains. This has proved to be simple to implement, as it does not require changes to any computation. Instead, the processes exchange data just prior to sending them to VisIt, with a call to `MPI_Sendrecv`. Since the amount of data is small, the impact on memory usage and simulation performance is negligible. Figure 7.4 illustrates this concept.




Figure 7.4: Exchange of data between neighboring processes to construct ghost nodes.

This and other callback functions are implemented in a similar manner to the ones presented above, i.e., they all return a handle to an object that contains the necessary information. By providing the aforementioned metadata, VisIt can now connect to our programs, and visualize variables as they are updated. Figure 7.5 illustrates some of the visualizations made possible with this integration.

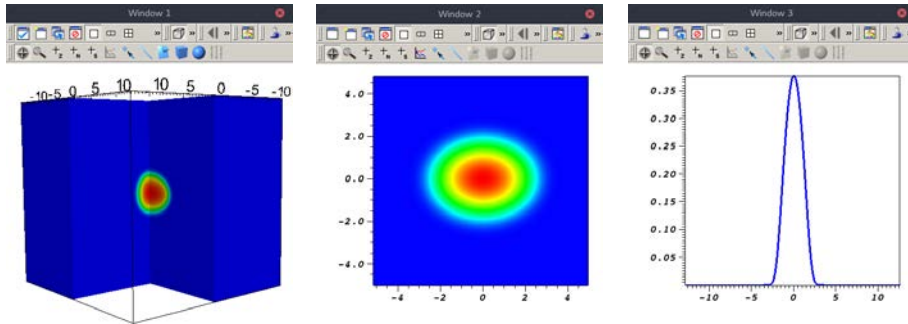


Figure 7.5: Examples of visualizations in VisIt: 3D pseudocolor plot (left), 2D pseudocolor plot (center) and curve plot (right).

We note that, while interacting with the simulation is very useful, sometimes the classic, approach where analysis of data happens after the simulation, is preferred. To enable this usage scenario, we added support of exporting data to VisIt as files independent of the running simulation. For visualization purposes, VisIt supports numerous formats via plugins, however two formats stand out as being the most commonly used: Silo and VTK. Silo [74] is data format and library for reading and writing a wide variety of scientific data, developed together with VisIt. It is mostly used for data intended to be visualized with VisIt, however it can be used with other tools as well. VTK [75] stands for Visualization Toolkit and is a software system for 3D computer graphics, image processing and visualization. It is developed by the same company that is also the primary developer of ParaView. VTK supports two different file formats, the legacy textual format that is easy to write either by hand or a library, and the newer XML-based format which has more advanced features and is often written with its own library. As all three file formats (Silo, textual VTK and XML-based VTK) meet our needs, the choice between them was based on personal preference and ease of use. We decided to use textual VTK, which we found to be more widespread in scientific communities, has good compatibility with the major visualization tools, and is simple to use. VisIt provides sample programs for writing textual VTK files, which we used as the base for our implementation. In this way the same set of meshes and variables can be exported as VTK files. Since the binary output produced by our programs can still be useful for data analysis without visualization, we decided to keep it together with VTK output.

## 7.2 Simulation steering

With changes shown in the previous section VisIt viewer component can connect to our programs and visualize the exported variables. However, the simulation is still independent of VisIt, and runs even if there is no connection from VisIt. In this section we describe how our simulation is extended to accept commands from VisIt.



VisIt supports issuing custom commands that can be used to send data to the simulation. The data sent can be interpreted as instructions to alter the simulation state. There are three ways to send data to the simulation: via simulation window in viewer GUI, via a custom user interface, or via console interface. In order to cover more use-cases, we implemented support for all three ways of commanding.

Simulation window, available after VisIt connects to a simulation, allows custom commands to be specified, which are accessible as push buttons the user can interact with. Pushing a button invokes a callback function in the simulation, which can perform some action based on the button selected. While useful for some basic commanding, this approach is very limited because there is no way to send any parameters to the simulation. That means that the callback function can only perform work based on the name of the command. For this reason, we use this functionality only for the most basic commands, i.e., start/pause the simulation, reset the simulation to the initial state, perform one step of time propagation, force the update of VisIt's plots, and stop the simulation and free the resources. To create these commands, they first must be registered with VisIt as part of the metadata callback (Listing 7.1). This essentially provides labels that will be written on the buttons. After that, along with other callbacks, we provide a callback to a function that will be invoked when a user clicks on a button in the simulation window. Listing 7.4 demonstrates this.

Listing 7.4: Registering basic VisIt commands.

---

```
// Callback for VisIt's Simulation window
void SimCommandCallback(const char* cmd, const char* args, void* cbdata) {
    if (strcmp(cmd, "START/PAUSE") == 0) {
        SimUIStartPauseClick(cbdata);
    } else if (strcmp(cmd, "RESET") == 0) {
        SimUIResetClick(cbdata);
    } else if (strcmp(cmd, "NEXT STEP") == 0) {
        SimUIStepClick(cbdata);
    } else if (strcmp(cmd, "UPDATE PLOTS") == 0) {
        SimUIUpdatePlotClick(cbdata);
    } else if (strcmp(cmd, "STOP") == 0) {
        SimUIStopClick(cbdata);
    }
}
void SimUIStartPauseClick(void* cbdata) {
    SimData* sim = (SimData*) cbdata;
    sim->running = !sim->running;
    // Temporary disable the Simulation window so that the user can not issue multiple
    // commands and cause concurrency problems
    SimSetupUI(sim, !sim->running);
    VisItTimeStepChanged(); // Notify VisIt so that the GUI is refreshed with the new data
}
void SimUIStepClick(void* cbdata) {
    SimData* sim = (SimData*) cbdata;
    SimSetupUI(sim, 0);
    SimNextStep(sim); // Perform next step in the time propagation and update GUI
    SimSetupUI(sim, 1);
}
... // Remaining three functions omitted for brevity
```

---

Greater flexibility can be obtained by creating a custom GUI. They are made using Qt toolkit [76], in which VisIt's GUI itself is made. Custom GUIs are designed using Qt Designer tool, which is a part of Qt software package. While simulation window only allows push buttons to be

programmed, custom UIs support three different widgets (at the time of writing): push button, check box, and spin box widget. Push button widgets work the same way as in simulation window, allowing some predefined action to be taken. On the other hand, check box widgets can be used to change the value of some boolean parameter in the simulation, while spin box widgets can be used to change a numeric value of some variable. Noticeably missing is the support for text input widget, which would be useful to send some textual (or decimal) value to the simulation. Support for additional widgets may come in a future version of VisIt, however we find the current widgets to be adequate for most cases. By using these three widget types we implemented support for: toggling computation of output variables (which may significantly speed up each step of the simulation), toggling automatic updates of VisIt's plots, writing current state of variables on-demand in multiple formats, and performing the selected number of steps in the simulation (Listing 7.5). The basic commands of the simulation window were also implemented in the custom GUI. In addition to sending data to the simulation, custom GUIs can also receive data from it. In this way, the custom GUI may present some information to the user. For example, we use this functionality to display the current iteration number on the custom UI.

Listing 7.5: Registering VisIt commands for the custom GUI. The callbacks for individual actions are similar to the functions in Listing 7.4

---

```
void SimRegisterCommands(SimData* sim) {
    // Simulation control
    VisItUI_clicked("btnStartPause", SimUIStartPauseClick, sim); // Button
    VisItUI_clicked("btnStep", SimUIStepClick, sim);
    VisItUI_clicked("btnReset", SimUIResetClick, sim);
    VisItUI_clicked("btnStop", SimUIStopClick, sim);
    VisItUI_clicked("btnUpdatePlot", SimUIUpdatePlotClick, sim);
    VisItUI_stateChanged("boxAutoUpdate", SimUIAutoUpdateChange, sim); // Check box
    VisItUI_valueChanged("spnStep", SimUIStepValue, sim); // Spinner
    // Variables to track
    VisItUI_stateChanged("boxMuen", SimUIMuenChange, sim);
    VisItUI_stateChanged("boxRms", SimUIRmsChange, sim);
    VisItUI_stateChanged("boxPsi2", SimUIPsi2Change, sim);
    VisItUI_stateChanged("boxDenX", SimUIDenXChange, sim);
    VisItUI_stateChanged("boxDenY", SimUIDenYChange, sim);
    VisItUI_stateChanged("boxDenZ", SimUIDenZChange, sim);
    ...
    // Output
    VisItUI_stateChanged("boxTextOut", SimUITextOutChange, sim);
    VisItUI_stateChanged("boxVtkOut", SimUIVtkOutChange, sim);
    VisItUI_stateChanged("boxAutoOut", SimUIAutoOutChange, sim);
    VisItUI_clicked("btnOutput", SimUIOutputClick, sim);
}

```

---

The custom GUI can be invoked from the simulation window, and interaction with it causes the registered functions to be called. Again, a simple mechanism based on a name of the widget user is interacting with is used to differentiate between the widgets in the GUI. The simulation window and the custom GUI are shown in Figure 7.6.




Figure 7.6: Simulation window with five programmable buttons (left) and more advanced custom GUI (right).

GUIs created in Qt Designer are essentially XML files, and the user does not have to compile them or link them in any special way with VisIt. The user only has to place the XML file of the GUI in a directory where VisIt can find it. They are completely independent of the simulation itself, and their (un)availability does not interfere with the simulation.

The third option for issuing commands to the simulation is via the console interface. This interface allows the user to type commands in the console window of the running simulation. This way of commanding predates other two approaches and is perhaps the most flexible. This flexibility comes from the fact that anything the user types in the console gets sent to the simulation as a string that can be further processed in any way. Therefore, with sufficiently capable parser, the console interface can alleviate the current limitations of GUI commanding. Additionally, this is the only available interface when VisIt is started without its GUI component. VisIt can be started without its GUI and scripted to perform many advanced actions, however this scenario is out of scope for this thesis. Since we did not focus on advanced scripting with VisIt, we implemented the console interface to have the same functionality as the graphical one.

We added support for the console interface as a special case in the main VisIt loop inside the simulation, as illustrated in Listing 7.1. `VisItDetectInput` returns a value of 3 to denote that the command issued was a console command, and we process it with the `SimProcessConsoleCommand` function. This function may then perform any work based on user's input. We do not present its implementation because it covers the same functionality already presented.

## Chapter 8

# Performance evaluation and modeling

In previous chapters we presented various parallel implementations of programs solving the dipolar GP equation without providing details of their measured performance. This chapter is dedicated to filling that gap. Before performance evaluation, all programs were tested for correctness to make sure they produce the same results as original serial programs, and extended to record execution time of the main loop.

We tested various scaling scenarios, strong and weak, from a single computing node to the cluster. All testing was done at the PARADOX supercomputing facility located at the Scientific Computing Laboratory, Center for the Study of Complex Systems of the Institute of Physics Belgrade.

Performance of hybrid implementations is highly dependent on the amount of work offloaded to GPU(s), so the key to maximizing performance of these versions is proper selection of parameters which control this feature. We used evolutionary computation techniques to find the optimal solutions, as described in Section 8.1. In the remaining two sections we give an overview of the tests and the methodology they relied on (Section 8.2), and present the performance results and their modeling in Sections 8.3 and 8.4. Section 8.5 discusses how to select the optimal algorithm for a given hardware platform.

### 8.1 Optimization of input parameters for hybrid implementations

Hybrid implementations, running on a single computer or a cluster with multiple computing nodes, can potentially offer the best performance of all algorithms presented in this thesis, by utilizing all allocated computing resources. For this to happen, work must be divided between CPU and GPU in such a way as to maximize their throughput and minimize their idle time. A question that naturally arises is how to achieve this for any possible mesh size for a given CPU/GPU combination. Unfortunately, there is no single best way to divide the work, since processing powers of CPUs and GPUs vary significantly. For example, in a computer where a powerful new GPU is paired with an older type of CPU, more work should be offloaded to the GPU, and vice versa for computers with a powerful CPU and low-performing GPU. Even if the installed CPU and GPU offer similar

performance in terms of floating point operations per second (as is the case for the PARADOX cluster), the differing architectures mean that certain portions of the algorithm are better suited to CPU or GPU. If we were to divide the work naively, e.g., equally, this would lead to unbalanced computation, as illustrated in Figure 8.1(a), which in turn increases the computation time.

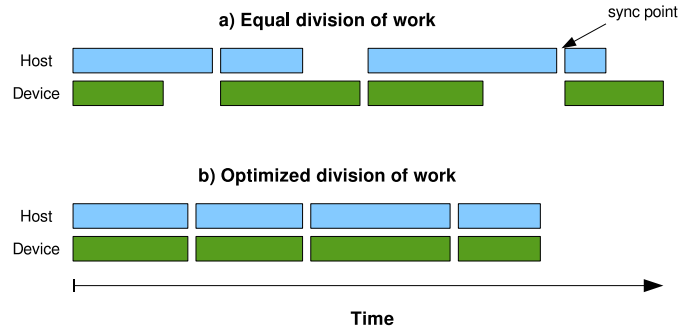


Figure 8.1: Execution timeline in a hybrid system: (a) with equal data distribution, leading to unbalanced computation time; (b) with optimized data distribution and ideal computation load.

To get over this problem, our hybrid implementations have an option to manually specify each parameter that controls the work being done on a GPU. Ideally, this flexibility would allow us to overlap computation on CPU and GPU as much as possible, thus minimizing the execution and idle time of each resource. The execution timeline for the ideal case looks like the illustration in Figure 8.1(b).

The aforementioned parameters include the total amount of data transferred to GPU, the number of chunks, kernel grid size parameters and a special parameter controlling whether some functions will even be offloaded to GPU or not, which is useful in situations where the CPU is much more powerful than the GPU. In 3D programs, this amounts to 33 integer-valued parameters for which we would like to find the optimal values for the desired mesh size on a given computing system (with given hardware characteristics). Furthermore, there are constraints that the parameters must satisfy. For example, the total amount of data transferred to GPU must be divisible by the number of chunks, and the size of a block that kernel launches must fit within the limits imposed by the underlying GPU.

To evaluate any combination of parameters, the programs have to be executed, and their execution time measured. The search space of 33 parameters is clearly too large to be exhaustively traversed in a reasonable amount of time. Fortunately, we can reduce the number of parameters we have to search through at any single time by grouping them based on the synchronization point. After the CPU and GPU synchronize, data have to be divided again, so the only relevant parameters in the region between two synchronization points are the ones controlling the division of work and the ones controlling the kernels which are executed in that region. During a single iteration, CPU and GPU have to synchronize four times:

1. after the FFT on local dimensions,
2. after the inverse FFT on the non-local dimension,
3. before `calc_lux` function, and

4. after `calclux` (which represents the end of a single iteration).

This gives us four distinct parameter sets to optimize:  $PS_1$  with 6 parameters,  $PS_2$  also with 6 parameters,  $PS_3$  with 15 parameters and  $PS_4$  with 5 parameters. We can consider each parameter set independently.

We investigated several approaches to optimization of parameter sets: naive brute-force search, iterative optimization via gradient descent, and metaheuristic via genetic algorithm. Our goal was not to find the best possible optimization algorithm for our problem, but rather than to implement a reasonably good one. By this we mean the algorithm which is not difficult to implement, which will give us a set of parameters that are close to the optimal ones, and will not take a long time to finish. Our focus was on 3D variants of both single node and MPI versions of hybrid programs. The concepts presented here apply equally to 2D variants and there is no important difference between the single node and the MPI version, so we will not make this distinction in the remainder of this section. We will now describe how the three approaches were implemented.

A brute-force search (BFS), or exhaustive search, considers every possible combination of parameter values in order to find the optimal one. Even with parameters divided into smaller sets, the exhaustive search takes too long, and is only feasible if we narrow down the ranges of all parameters involved. Unfortunately, this is only possible when we have a deep understanding of the performance of the CPU and GPU models used, when we can provide a reasonable guess for the optimal solution. An alternative is to consider only certain values along the range of a single parameter, rather than every value in the range. This approach is feasible and we were able to find reasonably good parameter values using it. The downside is that this way we may miss the optimal values of parameters, as they often lie between the selected test points. For instance, on one PARADOX computing node and a mesh size  $256 \times 256 \times 256$ , with the work divided between CPU and GPU along the outermost,  $x$  dimension, if we test only for the values of offloaded data to GPU that are multiples of 32 (i.e.,  $32 \times 256 \times 256$ ,  $64 \times 256 \times 256$ , etc.), we will miss the optimal value which lies around  $140 \times 256 \times 256$  (assuming all other parameters are fixed). Even though BFS is often not feasible, it is useful as the baseline for the evaluation of other methods we implemented.

With the BFS implemented as a baseline, we switched our focus to the implementation of an iterative optimization algorithm. A well-known and widely used algorithm is *gradient descent* (GD) [77]. It is an iterative procedure in which every iteration aims to get closer to the minimum of the given function  $F(\mathbf{x})$ , where  $\mathbf{x}$  represents a vector of function parameters. If the function  $F(\mathbf{x})$  is defined and differentiable around some point  $\mathbf{a}$ , then  $F(\mathbf{x})$  decreases fastest in the direction of the negative gradient of  $F$  at  $\mathbf{a}$ . GD exploits this fact and defines  $\mathbf{b}$  as

$$\mathbf{b} = \mathbf{a} - \gamma \nabla F(\mathbf{a}), \quad (8.1)$$

so that  $F(\mathbf{b}) \leq F(\mathbf{a})$ , for  $\gamma$  small enough. This observation is used as a building block for the algorithm. We start from an initial guess  $\mathbf{x}_0$  for a location of the minimum of  $F$  and construct a sequence  $\mathbf{x}_n (n \geq 1)$  such that:

$$\mathbf{x}_n = \mathbf{x}_{n-1} - \gamma \nabla F(\mathbf{x}_{n-1}). \quad (8.2)$$

This sequence will eventually converge to the local minimum, if one exists. GD is most applicable to functions for which the gradient can be computed analytically, however it can also be used when the derivative can only be obtained numerically. To apply it here, the execution time of our programs will be the function  $F$  that we minimize, and its arguments are the  $n = 33$  parameters that control

the GPU. The execution time is measured as the average time of a single iteration of the main time-propagation loop, sampled over a given number of iterations. As explained earlier, we can divide the problem into four separate minimization problems, with four sets of parameters of dimensionality  $m_i$  ( $i = 1, 2, 3, 4$ ), in the same way as in BFS. To numerically compute the  $k$ -th partial derivative of one of the minimization functions at a set of candidate parameters  $a_1, a_2, \dots, a_m$ , we use first order approximation:

$$\frac{\partial F}{\partial x_k}(a_1, \dots, a_m) \approx \frac{F(a_1, \dots, a_k + h_k, \dots, a_m) - F(a_1, \dots, a_k, \dots, a_m)}{h_k}, \quad (8.3)$$

where  $h_k$  is the increment of a given parameter. Since the parameters are integer-valued, we have to carefully choose the value of  $h_k$ , to be as small as possible, while still satisfying all the constraints that the corresponding parameter may have w.r.t. other parameters. We note that in order to evaluate the full gradient of  $F$  at  $(a_1, \dots, a_m)$  we need to execute our programs  $m + 1$  times.

We stress that execution times of our programs are not always the same due to the hardware, software and OS scheduling issues, making the minimization functions noisy. This noise affects the calculation of the gradient and could point the algorithm in the wrong direction. We can detect this by checking if the value of minimization function has increased in subsequent GD iteration, and discard this move if necessary. While useful, resorting to this tactic means that in case the GD gets stuck in a local minimum, it will not be able to get out and thus will never reach the global minimum. A naive way to check if this is the case would be to start over from a different initial point, and see if the GD algorithm converges to the same minimum. There are more sophisticated approaches to addressing this issue, which we discuss later.

In general, it is not possible to completely eliminate noise, due to inherent problems of accurately measuring time on a computer. However, we can try to minimize it by increasing the precision of execution time measurement of our programs. This can be achieved by averaging execution times over larger number of iterations of the main time-propagation loop. Unfortunately, this also means that the total execution time of the GD algorithm will significantly increase.

Another issue in our implementation of GD arises when we have to select the next values of our parameters based on the output values of previous GD iteration. These output values are real-valued, meaning that we need to convert them to integers somehow. Simply rounding them up or down to the nearest integer will not be enough, especially for small value of  $\gamma$ . For example, if we set the initial value of some parameter to 10, and after one GD iteration the proposed value of that parameter is 10.2, rounding it down will reset the value back to 10, effectively discarding the whole GD iteration. To avoid getting stuck in this way, our implementation selects the next possible value in the direction of change. In the example above it would mean selecting 11 as the new value. This change makes it impossible for the algorithm to converge, but it will usually stay close the minimum where all proposed solutions are of similar quality.

With the GD algorithm implemented as described above, we are able to get a set of optimized parameters from a random set of initial values. However, such optimized parameters are often suboptimal, and do not have the best performance. This is due to the fact that in most cases the GD will converge to the nearest local minimum, and there is no guarantee that that local minimum is also the global one. Using our previous example from BFS, a mesh of  $256 \times 256 \times 256$ , with GD we obtain that local minima exist for all power of two values, i.e., parameters suggest offloading  $2^n \times 256 \times 256$  to the GPU. If we randomly select a small initial value, GD will get stuck in the nearest local minimum, which in this example will be far from the global one, thus producing a

very bad solution. We can attempt to avoid getting stuck in a local minimum by adapting the parameter  $\gamma$ , also known as *learning rate* in machine learning literature. Keeping the learning rate high for a first few iterations would allow the algorithm to find the general location of the minimum, after which we could gradually lower the learning rate until convergence is achieved, in a process called *annealing* [78]. In practice this did not completely solve the problem, as we found that the optimal learning rate to start with varies with the initial set of parameters and the mesh size, and thus has to be manually selected. This complicates attempts to automate the process of finding the optimal parameters for a range of mesh sizes, which we needed as part of the tests of both hybrid implementations. A method to remove the manual tuning of the learning rate exists [79], but we find it to be too complex to implement for our programs, because it would require significant changes to the way the parameters are selected.

From the discussion above, we conclude that the GD is not the best-suited method for the optimization in our case, mostly because our minimization function is noisy and the numerical computation of the gradient is costly. Derivative-free optimization methods would be better suited to the problem, e.g., stochastic approximation algorithms like simultaneous perturbation (SPSA) [80, 81], or metaheuristics like genetic algorithms. We decided to use a genetic algorithm approach, as it is simple to understand and implement, as well as easy to adjust to get the desired behavior.

Genetic algorithm (GA) is an optimization method based on natural selection that mimics the process found in biological evolution [82]. GA works by creating a population of individual solutions, which it then evaluates and modifies, creating a new population, and iterating this process. Unlike the classical algorithms like GD, which iterate a single candidate solution towards the optimal one, the GA iterates a population of solutions in which the best individuals approach the optimal solution. The initial population is usually created at random, giving the GA different points in the search space to start from. The individual solutions are evaluated using a user-supplied fitness function, giving each individual a score based on how well they perform the given task. Individuals with the highest score are then selected to “reproduce” and create new offspring, after which they may be mutated randomly. The offspring form a new population, and the process can be repeated again. The GA continues until a suitable solution is found, or after a certain number of generations has passed.

To implement a GA, we first have to decide how to represent an individual. Individuals are created based on their blueprint, called *chromosome* in GA terminology. The most often used representation is bit-string [83], where all properties of an individual (its *genes*) are serialized to an array of bits, and then concatenated. More advanced representations exist, e.g., for encoding real values, permutations and general data structures [84, 85], however since our individuals are sets of integer parameters, we did not develop any special representation and instead we used ordinary arrays of integers as a chromosome. Therefore, each gene in a chromosome is a single GPU parameter of our programs. Individuals need to be evaluated using a fitness function. In our case, this means executing the programs with the parameters extracted from the individual’s chromosome, and reporting the execution time as fitness score, with lower execution time being better.

Next step is the implementation of the three GA operators: selection, reproduction (*crossover* in GA terminology) and mutation. Each operator can be implemented in different ways. Most common type of the selection operator is *roulette wheel* selection. This is a type of fitness-proportionate



selection, with the idea to give each individual a slice of the circular roulette wheel based on their fitness. The wheel is then spun, and when the roulette ball stops, the individual in whose region the ball stopped is selected. In this way, the fittest individuals have the greatest chance of being selected for reproduction, while the ones with very low score quickly die out. We also tested an alternative type of selection, the *tournament* selection. In this type of selection, we randomly select several individuals, and host a tournament for them. The individual with the best fitness score wins, and is selected for reproduction. By changing the size of the tournament we can control the selection pressure, e.g., by using small tournament size in the first few iterations we can prevent premature convergence and increase it in later generations when we have explored the search space enough. In our tests, the tournament selection gave slightly better results for smaller populations, by keeping the population diverse in early generations. For larger population size, both selection rules performed equally. Alongside the main selection algorithm, we also used *elitist* selection, where first few individuals are copied to the next generation without changing their chromosome. This prevented the loss of the best individuals in the next generation, but has to be used carefully as it may lead to premature convergence to a suboptimal solution.

To produce the next generation, selected individuals should combine their chromosomes and produce offspring. This is the task of the crossover operator. Crossover exchanges parts of the chromosomes, mimicking the biological recombination from nature. In its most basic form, crossover works by randomly selecting a point and exchanging segments before and after the point to create two new offspring from two parents. This type of crossover is called *single-point crossover*, and is illustrated in Figure 8.2(a). Other popular crossover techniques are *two-point crossover* and *uniform crossover*, illustrated in Figures 8.2(b) and 8.2(c). Two-point crossover is similar to the single-point crossover, just with two points instead of one. On the other hand, in the uniform crossover each gene of the offspring is selected randomly, either from the first parent or from the second one, with some fixed probability, typically 0.5. Using the uniform crossover leads to a wider exploration of the search space [86], but this may not always result in better performance of the operator [87]. We tested single- and two-point crossover and found that there is little difference in terms of performance between them. Since some of our parameters have constraints, it is important to select only the crossover points which lead to an allowed recombination of genes. Unfortunately, this made the implementation of uniform crossover impossible, and we did not pursue it further.




Figure 8.2: Three different crossover techniques in GA: a) single-point crossover, b) two-point crossover and c) uniform crossover.

The final step in producing the next generation is to apply the mutation operator on the new population. The mutation operator randomly changes genes to new values, which is equivalent of performing a random walk through the search space. Changing every gene would not be desired, so mutation operates with a very small probability, meaning most of the genes will be left unchanged.

When chromosomes are implemented as bits, the mutation would be equivalent to flipping a random bit. However, we could not implement mutation this way, as random changes to the chromosome would often result in non-functioning individuals, e.g., parameter constraints would not be satisfied. Special care must be taken to ensure that the mutation produces a healthy individual, similarly to the crossover operator. The importance of mutation operator and its relation to the crossover is often debated [88, 89], with its role being defined as “to maintain diversity within the population and inhibit premature convergence” [90], as the crossover operator does not introduce new information to the population. In our tests, the mutation plays a crucial role if the population size is small, e.g., less than 20 individuals, in accordance with Refs. [91, 92]. If the population size is large, e.g., over 200 individuals, the positive effects of mutation on the population fitness are not that evident.

GA is typically iterated for a fixed number of generations, as is the case in our implementation. Alternatively, we could have implemented some exit clause in the main loop of the GA which would stop the evolution after the best individuals have not been improved for some number of generations. After each run of the GA, there is usually several highly fit individuals in the population. Since randomness is an integral part of any GA, different runs of the algorithm produce slightly different results. Therefore, our GA does not converge to a single solution, but produces candidate solutions that have very similar parameter values and overall fitness. Among the candidate solutions, there may be small differences in the amount of data offloaded to GPU or the kernel parameters, with the execution times negligibly different. The number of generations required to produce a good solution also varies due to the randomness in the initial population and the population size. In our tests, it has a value between 10 and 50.

The implementation described above provides only the basis for a successful application of GA to the problem at hand. As can be seen, there are several important parameters of GA to tune in order to get optimal results. These include the population size, the number of generations, mutation rate, and the selection parameters (e.g., elite selection rate, the tournament size). We did not perform thorough testing of the performance of the GA that would allow us to obtain the best values for these parameters, as we were more focused on the quality of the candidate solutions GA creates. However, we observed that our GA finds good solutions faster if the initial population size is between 100 and 200, with the number of generations between 10 and 20, depending on the set of parameters we wish to optimize.

To get a perspective of how the three optimization methods perform, we tested them by comparing their final solutions, as well as by recording the number of program executions needed to get to the optimal set of parameters. For this test, we used 3D real-time propagation hybrid program on a single computing node. All three methods were allowed to execute the program up to 1000 times. The mesh sizes used range from  $80 \times 80 \times 80$  to  $600 \times 600 \times 600$ , all of which could potentially be offloaded to GPU. We tested parameter set  $PS_3$  (with 15 parameters), which control the execution of FFT and the kernels in the subsequent `calcnu`, `calcluy` and `calcluz` functions. The BFS algorithm was used as the baseline. Note that the range of parameters and a small number of allowed executions implies that the BFS algorithm takes the values of parameters with large stride, potentially missing the optimal solution. GD method is used as described, with the learning rate  $\gamma$  initially set to a higher value, which was gradually decreased. We performed the GD for 60 iterations, which amounts to 960 program executions. The GA was run on a population of 100 individuals, for 10 generations, amounting to the same the number of program executions as the

BFS. We have used mutation rate of 5%, and the tournament selection, with the tournament size equal to 10% of the population size. Elitism was also included, with the top 2% of the population copied over. All algorithms were tested five times, to minimize the effects of random initialization of both GD and GA methods. The results are shown in Table 8.1.

Table 8.1: A comparison of three optimization methods. Reported times are given in milliseconds, for a single iteration of the main time-propagation loop, averaged over 50 iterations. The last column contains the minimal execution times obtained by manual tuning. The reported time for each algorithm represents the minimal achieved value in five test runs.

Mesh size	BFS	GD	GA	Best
$80 \times 80 \times 80$	7	6	6	6
$128 \times 128 \times 128$	18	27	16	16
$240 \times 240 \times 240$	154	272	135	126
$256 \times 256 \times 256$	181	199	169	156
$360 \times 360 \times 360$	343	387	312	298
$480 \times 480 \times 480$	981	1049	868	829
$512 \times 512 \times 512$	1452	1628	1312	1242
$600 \times 600 \times 600$	2591	2984	2227	2159

From the results we conclude that the GA was the most effective optimization method of the three approaches. As expected, GA found better solutions than BFS due to large strides BFS had to use. The randomness in the initial population has a big effect on the convergence of the GA method, sometimes enabling it to find the optimal solution after just three generations. Even if the fitness of the initial population is very bad, the GA still converges to very good solutions after 10 generations. On the other hand, GD performed very poorly, mainly because the noise in program execution times has often thrown it in the wrong direction. Also, GD would get stuck in the nearest local minimum, which often was not the global one. When the initial position of GD is near the global minimum it converged to toward the optimal solution, which was rarely the case. For mesh sizes larger than the tested ones (relevant for MPI-based implementations), corresponding to a larger range of parameters, GD would be even less effective. Both BFS and GD can be made more usable if the range of parameters can be narrowed, i.e., if we know the relative performance of the GPU in comparison to the CPU. However, the GD would still be somewhat inefficient, due to its higher susceptibility to noise.

Since the GA method is shown to be superior to the other two approaches, we use GA as our optimization method of choice in the next two sections.

## 8.2 Testing methodology

All programs were tested on the PARADOX-IV cluster, which is a part of the PARADOX super-computing facility. This cluster is comprised of computing nodes with two Intel Xeon E5-2670 Sandy Bridge CPUs (with a total of  $2 \times 8 = 16$  cores), with 32 GB of RAM and one Nvidia Tesla M2090 GPU with 6 GB of GPU RAM, each connected by InfiniBand QDR interconnect. We used Intel's compiler (version 2016) to compile the serial and OpenMP programs, and CUDA 7.5 for

the GPU portions of the CUDA and hybrid programs. MPI-based implementations were compiled with Open MPI (version 1.10), which itself relied on underlying Intel and CUDA compilers. In the case of Hybrid/MPI programs, we performed tests using both FFTW and our own transpose routines. We found the minimal execution times to be about the same for both approaches, but the FFTW transpose would sometimes exhibit very bad performance due to the creation of suboptimal communication plan. For this reason, the execution times reported for Hybrid/MPI implementation are obtained using only our own transpose routines.

The base of all performance evaluations was the measured execution time of critical regions of the programs, i.e., the portions performing wave function propagation in imaginary or in real time. This measurement excluded the time spent in other parts of the programs, e.g., initialization of OpenMP/CUDA/MPI environment, memory allocation and deallocation, creation and destruction of FFTW plans, initialization of variables and I/O operations. Measuring average execution time of a single iteration of the main time-propagation loop allows us to predict the performance and total execution time of a given simulation, as the number of iterations is specific to the problem at hand and may vary significantly between different simulations. All measurements were collected using high precision timers based on `clock_gettime` POSIX function on the CPU side, and CUDA event API on the GPU side. The execution time of a single iteration of serial programs depends on the mesh size, controlled by variables `Nx`, `Ny` and `Nz`. For parallel programs, we can measure speedup and scaling efficiency as a function of the varying number of processing elements (OpenMP threads or MPI processes).

We tested performance of 1D, 2D and 3D programs on a range of mesh sizes, for a varying number of OpenMP threads and MPI processes, as shown in Table 8.2. Mesh sizes were chosen from the corresponding range, and we did not focus solely on mesh sizes which maximize performance of the programs, to obtain a more realistic assessment. We varied the number of OpenMP threads from 1 to 16, in increments of one thread. Similarly, we tested the MPI-based implementations by varying number of MPI processes, each bound to a different cluster node. In this way we tested MPI-based implementations on 2, 4, 8, 12, 16, 20, 24, 28 and 32 computing nodes. Note that varying the number of processing elements is not applicable to the single-GPU implementations of the programs, as they always use all available processing resources on a GPU. Only one of the two 1D and 2D programs were tested (corresponding to  $x$  direction in 1D and  $x$ - $y$  plane in 2D), because, performance-wise, there is no difference between them (e.g., `imag1dX-th` vs. `imag1dZ-th`).

The main performance indicator is the execution time, the wall-clock time of one iteration of the main loop, averaged over 5 executions of 1000 iterations, reported in milliseconds. Using the results obtained, we calculated the speedup of all programs compared to the published serial C implementation. We were also interested in examining the scaling efficiency, or *scalability*, of the OpenMP and MPI programs. We tested both *strong* scaling, when the mesh size stays the same but the number of processing elements varies, and *weak* scaling, when the amount of work each processing element performs stays the same while the number of processing elements increases. More formally, given the execution time of a single iteration of serial programs ( $T(1)$ ), and the corresponding execution time for parallel programs performed with  $N$  processing elements ( $T(N)$ ), we calculated speedup as  $S(N) = T(1)/T(N)$  and strong scaling efficiency as  $E(N) = S(N)/N$ . Weak scaling is computed as  $E_W(N) = T_W(1)/T_W(N)$ , where  $T_W(1)$  is the execution time of a program using single processing element performing the work assigned to it, while  $T_W(N)$  is the

Table 8.2: Performance testing matrix, showing the mesh sizes and numbers of processing elements (threads or processes) used to test the programs, as well as the baseline program used for comparison.

Program	Mesh size		Processing elements		Baseline
	Min	Max	Min	Max	
OpenMP programs					
imag1dX-th	1000	1000000	1	16	imag1d
real1dX-th	1000	1000000			real1d
imag2dXY-th	$1000 \times 1000$	$15000 \times 15000$	1	16	imag2dXY
real2dXY-th	$1000 \times 1000$	$13000 \times 13000$			real2dXY
imag3d-th	$50 \times 50 \times 50$	$800 \times 800 \times 800$	1	16	imag3d
real3d-th	$50 \times 50 \times 50$	$800 \times 800 \times 800$			real3d
CUDA programs					
imag2dXY-cuda	$1000 \times 1000$	$15000 \times 15000$	1	1	imag2dXY
real2dXY-cuda	$1000 \times 1000$	$13000 \times 13000$			real2dXY
imag3d-cuda	$50 \times 50 \times 50$	$600 \times 600 \times 600$	1	1	imag3d
real3d-cuda	$50 \times 50 \times 50$	$540 \times 540 \times 540$			real3d
Hybrid programs					
imag2dXY-hetero	$1000 \times 1000$	$15000 \times 15000$	16+1	16+1	imag2dXY-th
real2dXY-hetero	$1000 \times 1000$	$13000 \times 13000$			real2dXY-th
imag3d-hetero	$50 \times 50 \times 50$	$800 \times 800 \times 800$	16+1	16+1	imag3d-th
real3d-hetero	$50 \times 50 \times 50$	$800 \times 800 \times 800$			real3d-th
OpenMP/MPI programs					
imag3d-mpi	$480 \times 480 \times 250$	$1920 \times 1920 \times 960$	$1 \times 16$	$32 \times 16$	imag3d-th
real3d-mpi	$480 \times 480 \times 250$	$1920 \times 1920 \times 960$			real3d-th
CUDA/MPI programs					
imag3d-mpicuda	$480 \times 480 \times 250$	$1920 \times 1920 \times 960$	1	32	imag3d-cuda
real3d-mpicuda	$480 \times 480 \times 250$	$1920 \times 1920 \times 960$			real3d-cuda
Hybrid/MPI programs					
imag3d-mpihetero	$480 \times 480 \times 250$	$1920 \times 1920 \times 960$	$1 \times (16+1)$	$32 \times (16+1)$	imag3d-hetero
real3d-mpihetero	$480 \times 480 \times 250$	$1920 \times 1920 \times 960$			real3d-hetero

execution time of a program using  $N$  processing elements performing  $N$  times more work. We achieve this by increasing the mesh size.

### 8.3 Performance test results and modeling of single node programs

In this section we present the results obtained for single computing node OpenMP, CUDA and hybrid programs, and compare them to the previously published [13] serial implementation.

Strong scaling performance test results for the OpenMP-based implementation using methodology described in the previous section are given in Table 8.3 and Figure 8.3. They show the obtained execution times, speedups and strong scaling efficiencies for different number of OpenMP threads. Columns  $N_{th} = 1$ ,  $N_{th} = 2$ ,  $N_{th} = 4$ ,  $N_{th} = 8$  and  $N_{th} = 16$  in Table 8.3 correspond to the number of threads used, while the last column shows the obtained speedup  $S(16)$  with 16 OpenMP threads

compared to one OpenMP thread. Strong scaling efficiency  $E(N_{\text{th}}) = S(N_{\text{th}})/N_{\text{th}}$  in Figure 8.3 is calculated as a fraction of the obtained speedup compared to a theoretical maximum. The mesh size used in 1D is  $10^5$ , in 2D  $10^4 \times 10^4$ , while in 3D the mesh size is  $480 \times 480 \times 480$ . Execution times and speedups of imag1dZ-th, real1dZ-th, imag2dXZ-th, and real2dXZ-th (not reported here) are similar to those of imag1dX-th, real1dX-th, imag2dXY-th, and real2dXY-th, respectively.

Table 8.3: Wall-clock execution times of a single iteration of the main time-propagation loop of single-node OpenMP programs (in milliseconds) for different number of OpenMP threads  $N_{\text{th}}$  and speedup  $S(16)$  in strong scaling tests. The speedup is calculated w.r.t. the execution times of previously published serial versions of programs [13], given in the second column.

Program	Serial	$N_{\text{th}} = 1$	$N_{\text{th}} = 2$	$N_{\text{th}} = 4$	$N_{\text{th}} = 8$	$N_{\text{th}} = 16$	$S(16)$
imag1dX-th	9.1	7.1	4.7	3.4	2.9	2.8	2.5
real1dX-th	15.2	14.2	10.5	8.2	7.3	7.2	2.0
imag2dXY-th	13657	7314	4215	2159	1193	798	9.2
real2dXY-th	17281	11700	6417	3271	1730	1052	11.1
imag3d-th	16064	9353	5201	2734	1473	888	10.5
real3d-th	22611	17496	9434	4935	2602	1466	11.9

The change from C2C to R2C FFT routine has a big impact on the execution time of single-threaded ( $N_{\text{th}} = 1$ ) programs compared to the previous serial programs. As we can see from the table, these improvements alone yield a speedup of 1.3 to 1.9 in 2D and 3D programs, and somewhat smaller speedup for 1D programs, 1.1 to 1.3. The use of additional threads brings about further speedup (reported in the last column) of 2 to 2.5 for 1D programs, and 9 to 12 for 2D and 3D programs. In Figure 8.3 we see that the efficiency rapidly decreases for 1D programs, even though speedup increases with the number of threads used. This is expected, as parts of the algorithm dealing with the recursive relations for calculation of the CN coefficients are inherently serial. In 1D, already with  $N_{\text{th}} = 4$  threads we almost achieve the maximal speedup, while still keeping the efficiency around 50%. We also see that, as expected, speedup and efficiency of multidimensional programs behave quite well as we increase the numbers of threads. In particular, we note that the efficiency always remains above 60%, making the use of all available CPU cores worthwhile.

From Figure 8.3 we observe that the speedup of 1D programs saturate quickly due to inherent serial nature of the portion of the algorithm, while in 2D and 3D the speedup behaves almost linearly. Despite their obvious differences, all curves in Figure 8.3 can be successfully modeled based on Amdahl's law [93]. Namely, the measured execution time  $T(N_{\text{th}})$  of one iteration of the main loop can be expressed as

$$T(N_{\text{th}}) = T(1) \left( s + \frac{p}{N_{\text{th}}} \right), \quad (8.4)$$

where  $N_{\text{th}}$  is the number of threads used,  $T(1)$  is the execution time of a single-threaded run,  $s$  is the serial fraction of the loop code, and  $p$  is the corresponding parallel fraction. By definition,  $s + p = 1$ , and therefore the speedup can be modeled by

$$S(N_{\text{th}}) = \frac{T(1)}{T(N_{\text{th}})} = \frac{1}{1 - p + p/N_{\text{th}}}, \quad (8.5)$$

where the parallel fraction of the main loop code  $p$  is the only fit parameter. Table 8.4 gives the obtained model fit parameter values for the data from Figure 8.3. As we can see, the fits match

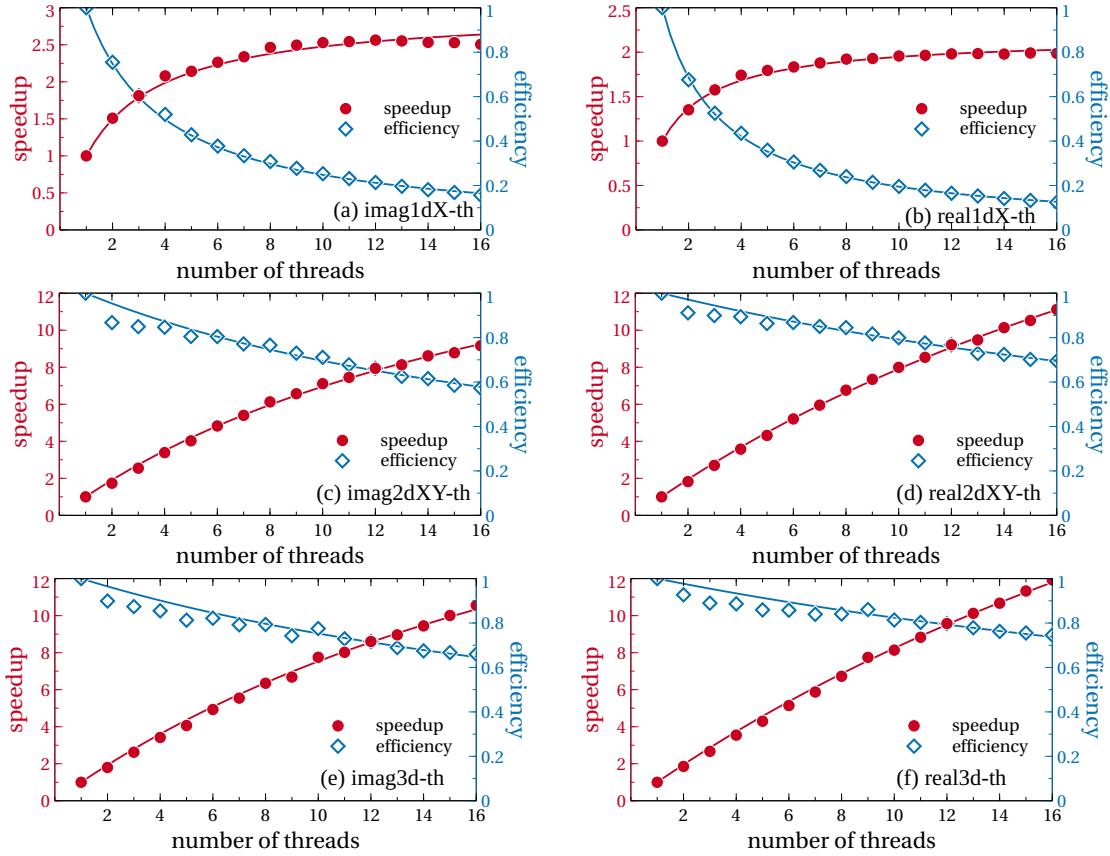


Figure 8.3: Speedup in the execution time and strong scaling efficiency of OpenMP programs compared to single-threaded runs for: (a) `imag1dX-th`, (b) `real1dX-th`, (c) `imag2dXY-th`, (d) `real2dXY-th`, (e) `imag3d-th`, (f) `real3d-th`. Solid lines represent fits to measured data, where fit model functions are given in the text and obtained fit parameters are listed in Table 8.4. Note that the model parameter  $p$  is fitted only to the speedup data, and then used to plot the efficiency model curve.

the obtained measurement data very well. Note that the efficiency can be expressed in terms of the speedup model

$$E(N_{\text{th}}) = \frac{S(N_{\text{th}})}{N_{\text{th}}} = \frac{1}{N_{\text{th}}} \frac{1}{1 - p + p/N_{\text{th}}}, \quad (8.6)$$

and is not fitted independently. Figure 8.3 shows that this model with the parameter  $p$  fitted on the speedup data matches very well with the efficiency data points.

Table 8.4: Values of obtained strong scaling model fit parameter  $p$  and the estimated fit errors for OpenMP-parallelized programs.

Program	$p$	$\Delta p$
<code>imag1dX-th</code>	0.662	0.004
<code>real1dX-th</code>	0.541	0.003
<code>imag2dXY-th</code>	0.9514	0.0007
<code>real2dXY-th</code>	0.9706	0.0004
<code>imag3d-th</code>	0.9635	0.0008
<code>real3d-th</code>	0.9762	0.0006

From the obtained values of the parallel fraction of the algorithm, we see that 2D and 3D programs are almost ideally parallelizable, with  $p$  over 95%. In 1D case, however, the parallel fraction is around 66% for imaginary-time propagation and around 54% for real-time propagation, due to the fact that calculation of CN coefficients in the function `calclux` cannot be parallelized. Note that the parallel fraction of imaginary-time 1D program is higher than that of the real-time 1D program due to the larger amount of arithmetic operations required to process complex-valued data in the real-time version of the serial function `calclux`.

OpenMP programs were also tested for weak scalability. We were mostly interested in 3D variants of the programs, which we tested by fixing the amount of work to 6,912,000 spatial points of the mesh, which corresponds to a mesh size of  $240 \times 240 \times 120$ . By increasing the number of OpenMP threads, we also increase the mesh size to be the multiple of 6,912,000. This means that the mesh has to be increased in such a way that, when divided among the threads, each thread gets 6,912,000 spatial points to process. The mesh sizes that satisfy this requirement cannot always be obtained by multiplying all three array dimensions with the same number, so we have to work with mesh sizes in which the size of each dimension may be different. For such mesh sizes, we tested all possible combinations, however no significant difference has been observed. Table 8.5 contains mesh sizes we used for testing for  $N_{\text{th}} = 1, 2, 4, 8, 12, 16$  threads, along with the execution times (in milliseconds) and the weak scaling efficiency. The testing was also done for  $N_{\text{th}} = 6$  and 10 threads, but the corresponding mesh sizes are omitted from the table for brevity. However, Figure 8.4 presents complete data collected in this test, for all  $N_{\text{th}}$  values, averaged over mesh sizes used. The figure shows that the real-time 3D programs have better weak scaling efficiency, which is about 75% at 16 threads, while imaginary-time programs demonstrate smaller, but still significant efficiency of about 60%.




Figure 8.4: Weak scaling efficiency of OpenMP-parallelized 3D programs, averaged over all mesh sizes tested for each value of  $N_{\text{th}}$ . Solid lines represent fits to measured data, where fit model functions are given in the text.

To model weak scaling efficiency, we compare execution times of a single iteration of the main loop  $T_W(N_{\text{th}})$ , performed with  $N_{\text{th}}$  threads, where for each value of  $N_{\text{th}}$  the total amount of work is  $N_{\text{th}}$  times the work being performed in a single-threaded run, i.e., the amount of work per thread is constant. The weak scaling efficiency is defined as

$$E_W(N_{\text{th}}) = \frac{T_W(1)}{T_W(N_{\text{th}})}, \quad (8.7)$$

where  $T_W(1) = T(1)$  is the execution time of a single-threaded run for a given workload (in our case, 6,912,000 spatial points). The expected execution time of a workload assigned to a simulation with



Table 8.5: Weak scaling efficiency of OpenMP-parallelized 3D programs. Wall-clock execution times  $T_W$  are given in milliseconds, efficiencies  $E_W$  in percents.

Mesh size	imag3d-th		real3d-th	
	$T_W$	$E_W$	$T_W$	$E_W$
$N_{\text{th}} = 1$				
$120 \times 240 \times 240$	699	98.2	1259	99.5
$240 \times 120 \times 240$	711	96.5	1271	100
$240 \times 240 \times 120$	686	100	1275	98.2
$N_{\text{th}} = 2$				
$240 \times 240 \times 240$	718	95.5	1254	99.9
$N_{\text{th}} = 4$				
$240 \times 240 \times 480$	755	90.9	1312	95.5
$240 \times 480 \times 240$	778	88.2	1301	96.3
$480 \times 240 \times 240$	754	91	1310	95.6
$N_{\text{th}} = 8$				
$240 \times 480 \times 480$	875	78.4	1405	89.9
$480 \times 240 \times 480$	847	81	1393	89.1
$480 \times 480 \times 240$	886	77.5	1431	87.7
$N_{\text{th}} = 12$				
$360 \times 480 \times 480$	1010	67.9	1532	81.7
$480 \times 360 \times 480$	981	69.9	1514	82.7
$480 \times 480 \times 360$	1004	68.4	1581	79.2
$N_{\text{th}} = 16$				
$480 \times 480 \times 480$	1155	59.4	1680	74.5

$N_{\text{th}}$  threads, executed in a single-threaded run, is  $N_{\text{th}}T(1)$ . In weak scaling tests, this workload is executed with  $N_{\text{th}}$  threads and therefore the expected execution time can be modeled by

$$T_W(N_{\text{th}}) = N_{\text{th}}T(1) \left( 1 - p + \frac{p}{N_{\text{th}}} \right). \quad (8.8)$$

According to this argument, we model the weak scaling efficiency by a single-parameter function

$$E_W(N_{\text{th}}) = \frac{1}{p + (1-p)N_{\text{th}}}. \quad (8.9)$$

We fitted this model to data presented in Figure 8.4 and the obtained fit parameters are given in Table 8.6. As we can see from the figure, the above model is an excellent fit to experimental data, and both 3D OpenMP-parallelized programs have high (above 95%) parallel fraction of the code, in agreement with the results of strong scaling tests (Table 8.4).

Table 8.6: Values of obtained weak scaling model fit parameter  $p$  and the estimated fit errors for OpenMP-parallelized programs.

Program	$p$	$\Delta p$
imag3d-th	0.958	0.002
real3d-th	0.9792	0.0009

Overall, we conclude that imaginary-time OpenMP-parallelized programs show smaller speedup and efficiency across all tests, except in 1D. This can be attributed to the additional step of wave function normalization in each iteration of the imaginary-time propagation, as well as the fact that real-time programs work with complex-valued data that require more arithmetic operations for the same mesh size. Since much of the computation inside loops requires simple arithmetic, the throughput of the CPU is often not fully exploited in imaginary-time programs, thus the pressure of memory bandwidth makes these programs less efficient. The real-time programs are also affected by this, however to a somewhat lesser extent.

Next, we consider the CUDA implementation of the shared memory algorithm. GPU functions as a single processing element, therefore we cannot test CUDA implementation by varying both the mesh size and the number of processing elements. However, just varying the mesh size gives us valuable insight into the behavior of this implementation, due to the difference in programming models and the libraries used. Tables 8.7 and 8.8 show the execution times (in milliseconds) for a number of mesh sizes tested, as well as the average speedup compared to the serial programs [13]. Figure 8.5 shows the speedup obtained for all mesh sizes tested, and red horizontal lines represent average speedups obtained for each program. Note that the dispersion of data is due to the use of `FFTW_ESTIMATE` flag in library calls to FFTW in the serial programs. Use of this flag results in a choice of suboptimal FFT algorithm for some mesh sizes. The vertical lines in Figure 8.5 denote the change in `POTMEM` parameter. The speedups left of the first vertical line are obtained with `POTMEM=2` and thus demonstrate the best speedup. Second group of results, between the two vertical lines, is obtained with `POTMEM=1`, and we note that the speedup decreases slightly, while the results right of the second vertical line are obtained with `POTMEM=2`, and show the smallest speedup due to the use of mapped memory. The 2D programs in  $x$ - $z$  plane exhibited very similar performance to those of  $x$ - $y$  plane, and therefore we did not include them in the figures.

Table 8.7: Wall-clock execution times of a single iteration of the main time-propagation loop of single-node 2D CUDA programs (in milliseconds) for different mesh sizes, and average speedup w.r.t. to the execution times of serial 2D programs [13].

Program	$2000^2$	$4000^2$	$6000^2$	$8000^2$	$10000^2$	$12000^2$	Avg. speedup
imag2dXY-cuda	24.1	104.3	235.2	386.1	657.1	1150.4	10
real2dXY-cuda	29.9	112.4	266.4	444.0	749.0	1528.3	14

Table 8.8: Wall-clock execution times of a single iteration of the main time-propagation loop of single-node 3D CUDA programs (in milliseconds) for different mesh sizes, and average speedup w.r.t. to the execution times of serial 3D programs [13].

Program	$100^3$	$200^3$	$300^3$	$400^3$	$500^3$	Avg. speedup
imag3d-cuda	10.6	79.3	298.8	674.5	1260.2	7.1
real3d-cuda	10.9	84.1	302.5	682.4	1467.4	13.5

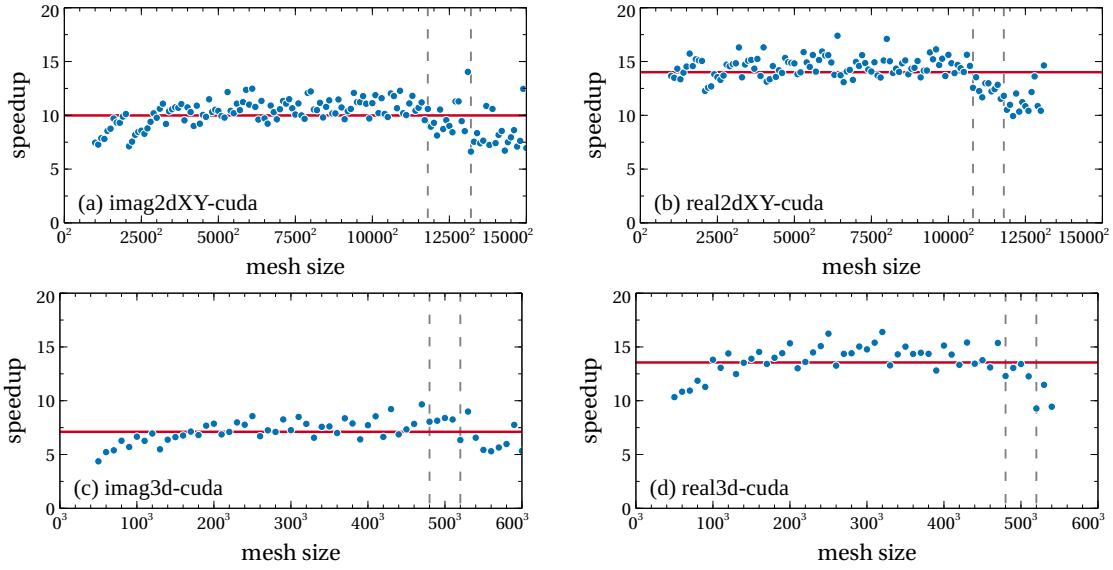


Figure 8.5: Speedup in the execution time of CUDA programs compared to the serial programs for all tested mesh sizes: (a) `imag2dXY-cuda`, (b) `real2dXY-cuda`, (c) `imag3d-cuda`, (d) `real3d-cuda`. Red horizontal lines represent the average speedup, while dashed vertical lines represent different values of POTMEM parameter (see text).

For small mesh sizes, the GPU remains underutilized, resulting in a smaller speedup. For large mesh sizes, where the GPU memory usage approaches the limit, we also see declining speedup. This is due to the inevitable use of POTMEM parameter, which keeps some arrays in the host memory when they cannot fit in the memory of the GPU. Overall, the CUDA implementation shows execution times similar to the OpenMP implementation for imaginary-time propagation, and slightly lower for real-time propagation. We stress that these results strongly depend on the type of GPU used and that the obtained speedups may be even better for newer-generation GPUs (e.g., *Kepler*-, *Maxwell*-, *Pascal*- and *Volta*-based GPUs).

Hybrid OpenMP/CUDA implementation was tested on a range of mesh sizes, similarly to the OpenMP- and CUDA-based implementations. Since OpenMP implementation showed that all cores should be used for 2D and 3D programs, we kept the number of OpenMP threads fixed at 16 in our hybrid algorithm. The parameters governing the amount of data offloaded to GPU were optimized using our GA method from Section 8.1. An illustrative subset of the results for 3D programs is shown in Table 8.9, which gives execution times for a single iteration of the main loop, for different mesh sizes. A comparison of execution times of hybrid and pure OpenMP programs with  $N_{\text{th}} = 16$  threads is given in Figure 8.6(a), while Figure 8.6(b) compares the performance of hybrid and pure CUDA programs, for all tested mesh sizes.

Table 8.9: Wall-clock execution times of a single iteration of the main time-propagation loop of single-node 3D hybrid programs (in milliseconds) for different mesh sizes. Offload parameters are optimized using the GA optimization method.

Program	$80^3$	$200^3$	$320^3$	$440^3$	$560^3$	$680^3$	$800^3$
<code>imag3d-hetero</code>	6.3	67.1	271.4	678.9	1493.3	2750.9	4461.6
<code>real3d-hetero</code>	8.0	94.7	374.1	922.8	1985.2	4020.2	6676.6

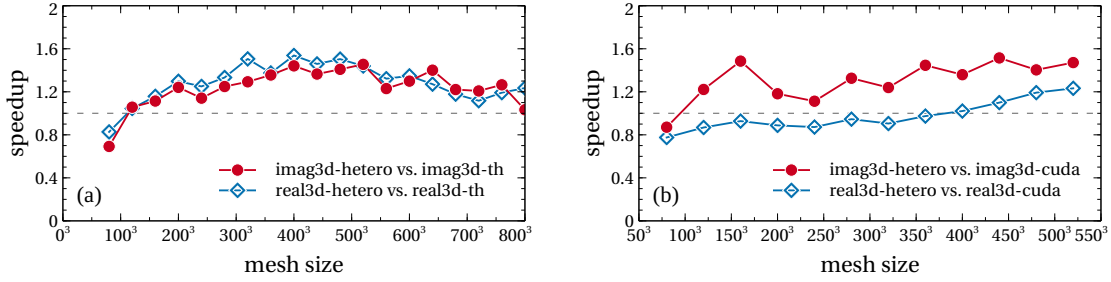


Figure 8.6: Speedup in the execution time of 3D hybrid programs compared to: (a) OpenMP programs with  $N_{th} = 16$  threads, (b) CUDA programs. Dashed horizontal lines correspond to a speedup value  $S = 1$ .

From Figure 8.6 we see that the hybrid implementation outperforms the OpenMP one for all mesh sizes except for the smallest one. The same applies to the comparison of imaginary-time hybrid and CUDA programs, while real-time hybrid program outperforms the corresponding CUDA program only for mesh sizes larger than  $400 \times 400 \times 400$ . Although one would expect that the speedup of the optimized hybrid algorithm is always equal to or larger than one, we see that this is not the case for all mesh sizes in Figure 8.6. This is due to the fact that the hybrid FFT algorithm is always employed in hybrid programs, meaning that even if no or all data are offloaded to GPU, the splitting of FFT will still take place. Therefore, in these limiting cases we do not obtain pure OpenMP or CUDA algorithms that would yield maximal performance. Splitting the computation of FFT along one direction in `calcpsidd2` function disables some of the potential optimizations that libraries like FFTW and cuFFT exploit, and introduces data copies between host and device which cannot be fully offset using CUDA streams. The hybrid algorithm can compensate for this if the amount of offloaded data to GPU is sufficiently large, which is not the case for the smallest mesh size.

Given that pure CUDA programs performing imaginary-time propagation demonstrated smaller speedup than the corresponding OpenMP programs (Figure 8.5 vs. Figure 8.3), it is expected that CUDA portion of the hybrid imaginary-time propagation programs would yield smaller improvement, which is evident when we compare hybrid implementation with the OpenMP one in Figure 8.6(a), except for the largest mesh sizes. We can reach the same conclusion if we consider that, in real-time propagation, CUDA implementation shows the highest speedup of all single-node programs, and therefore the corresponding hybrid programs show better speedup than imaginary-time propagation programs parallelized with OpenMP. On the other hand, when compared with the CUDA implementation in Figure 8.6(b), the situation is reversed, and we see much better performance of the hybrid imaginary-time propagation programs, while the hybrid real-time propagation programs achieve speedup larger than one only for mesh sizes larger than  $400 \times 400 \times 400$ . This result can be attributed to the fact that in real-time propagation the amount of data copied between host and device is larger, given that some of the copied arrays are complex-valued.

Also, we can observe that the speedup in Figure 8.6(a) declines with mesh size. This is due to the memory saturation of the GPU device, as only a small portion of the data can be offloaded when the mesh is large. The amount of data offloaded to GPU for all tested mesh sizes is given in Table 8.10, where we can see that the total amount of data processed by the GPU declines with increasing mesh size. As a result, the computation is unbalanced, given that the host has to work

with a much larger portion of the data. This is reflected in the optimization step, where our GA quickly converges to the highest possible mesh size that can be offloaded to the device, implying that a greater amount of device memory would be required to get a balanced computation.

Table 8.10: Optimal fraction of total data offloaded to GPU for different mesh sizes in hybrid 3D programs, obtained using the GA optimization method.

Mesh size	GPU portion of data [%]							
	imag3d-hetero				real3d-hetero			
	PS <sub>1</sub>	PS <sub>2</sub>	PS <sub>3</sub>	PS <sub>4</sub>	PS <sub>1</sub>	PS <sub>2</sub>	PS <sub>3</sub>	PS <sub>4</sub>
80 × 80 × 80	20	3	60	33	23	3	40	10
120 × 120 × 120	12	10	25	28	23	10	25	27
160 × 160 × 160	15	7	31	24	44	19	34	23
200 × 200 × 200	38	27	37	28	44	26	46	26
240 × 240 × 240	38	28	40	30	38	28	47	25
280 × 280 × 280	40	31	30	33	46	22	37	25
320 × 320 × 320	33	38	29	36	45	28	69	44
360 × 360 × 360	41	33	27	36	49	33	43	31
400 × 400 × 400	38	25	30	38	46	24	47	21
440 × 440 × 440	31	35	29	39	45	35	43	32
480 × 480 × 480	16	19	22	41	45	33	62	35
520 × 520 × 520	40	31	28	46	45	25	40	37
560 × 560 × 560	38	41	30	43	36	24	35	32
600 × 600 × 600	37	27	30	43	29	29	29	28
640 × 640 × 640	36	40	30	43	24	22	24	23
680 × 680 × 680	26	38	28	38	20	16	20	20
720 × 720 × 720	23	26	25	32	17	13	17	17
760 × 760 × 760	21	22	25	27	14	14	14	13
800 × 800 × 800	19	23	23	21	11	8	12	12

To make an optimal choice of programs to be used on a single node for a particular hardware platform, one has to perform detailed tests using the methodology presented in Section 8.2. The results presented in this section apply to hardware comparable to the one available at PARADOX supercomputing facility, but one can expect similar behavior for all modern types of CPU and GPU.

## 8.4 Performance test results and modeling of MPI programs

Following on from the previous section, here we present the results obtained for MPI-parallelized programs, executed on a computer cluster consisting of varying number of computing nodes, and compare them with the corresponding single-node programs.

MPI programs are highly dependent on the configuration of the cluster, mainly on the speed of interconnect, but also on the distribution of processes and threads, NUMA configuration, MPI configuration, etc. Getting the best performance out of the programs requires some experimentation with several different configurations. The results presented are obtained without extensive tuning

of the cluster or MPI runtime, with the aim to show the base performance.

Strong scaling of MPI-based implementations is tested in a similar way to the single-node testing, by varying number of cluster nodes  $N_p$  from 2 to 32. On each node, a single MPI process is launched, which then uses the node's resources, either by further spawning 16 OpenMP threads in case of OpenMP/MPI and Hybrid/MPI implementations, or by invoking the CUDA kernels utilizing the node's GPU in case of CUDA/MPI implementation. In this test, the baseline used for comparison was the equivalent single-node implementation, i.e., OpenMP/MPI programs are compared with OpenMP programs executing with  $N_{th} = 16$  threads, CUDA/MPI programs are compared to the single-node pure CUDA programs, and Hybrid/MPI programs is compared to the single-node hybrid programs.

Mesh size we use in this test is  $480 \times 480 \times 250$ . Since the values of  $N_x$  and  $N_y$  parameters must be divisible by the number of processes  $N_p$ , this mesh cannot be distributed over 28 processes. In this case we use slightly modified mesh, taking into account the multiple of 28 that is closest to 480, yielding a mesh size of  $476 \times 476 \times 250$ . Data is never distributed along  $z$  direction, so no such requirement exists for  $N_z$ . The resulting mesh has slightly fewer spatial points, and thus less work per process, potentially allowing for a better performance, but in our tests the smaller mesh did not have a significant impact on the measured execution time.

Note that strong scaling tests, where the mesh size is fixed, inevitably lead to saturation and decrease in measured speedup and efficiency values. Whatever the chosen mesh size, increase in the number of cluster nodes used will eventually yield insufficient amount of work per MPI process, such that communication will start to dominate over computation. Therefore, in our strong scaling tests we can expect that execution times initially decrease with increasing values of  $N_p$ , but eventually performance of MPI-based programs will decline. This issue is addressed by considering weak scaling tests.

The illustrative portion of obtained execution times of a single iteration of the main loop for the three MPI-parallelized implementations are shown in Table 8.11, together with the execution times of the corresponding baseline single-node program. Columns  $N_p = 4$ ,  $N_p = 8$ ,  $N_p = 16$ ,  $N_p = 24$  and  $N_p = 32$  correspond to the number of cluster nodes used, while the last column shows the obtained speedup  $S(32)$  with  $N_p = 32$  nodes compared to baseline single-node programs.

Table 8.11: Wall-clock execution times of a single iteration of the main time-propagation loop of OpenMP/MPI, CUDA/MPI and Hybrid/MPI programs (in milliseconds) for different number of MPI processes  $N_p$  and speedup  $S(32)$  in strong scaling tests. The speedup is calculated w.r.t. the baseline execution times of the corresponding single-node programs (OpenMP, CUDA and hybrid, respectively), given in the second column.

Program	Baseline	$N_p = 4$	$N_p = 8$	$N_p = 16$	$N_p = 24$	$N_p = 32$	$S(32)$
imag3d-mpi	1124	541	262	134	89	64	17.5
real3d-mpi	2140	700	358	207	155	98	21.8
imag3d-mpicuda	579	438	210	103	71	59	9.8
real3d-mpicuda	800	609	291	142	95	79	10.1
imag3d-mpihetero	489	299	162	99	84	81	6.0
real3d-mpihetero	613	407	255	154	135	101	6.0

The maximal speedup of OpenMP/MPI implementation ranges from 17 for imaginary-time propagation to 22 for real-time propagation programs for  $N_p = 32$ . The complete measurement results for speedup and efficiency are depicted in Figure 8.7(a) and 8.7(b), where we see that the speedup grows linearly with the number of nodes used, while the efficiency remains mostly constant in the range between 40% and 60%, thus making the use of OpenMP/MPI programs highly advantageous for simulations with large mesh size. In general, we can expect even better efficiency for larger mesh sizes.

Similar behavior is observed for CUDA/MPI implementation. The obtained speedup with  $N_p = 32$  nodes ranges from 9 to 10, with the slightly lower efficiency, between 30% and 40%, as shown in Figure 8.7(c) and 8.7(d). Even though the efficiency is lower for this implementation, the speedup still grows linearly and the execution times are lower than for the OpenMP/MPI implementation. This makes CUDA/MPI programs ideal choice for use on GPU-enabled computer clusters. Additional benefit of using CUDA/MPI programs is their low CPU usage (using only one CPU core per cluster node), allowing for the possibility that the same cluster nodes are used for other CPU-intensive simulations in a time-sharing fashion.

As we see in Figure 8.7(e) and 8.7(f), the linear growth of speedup is also present for the Hybrid/MPI implementation, however, as the amount of work per-process shrinks, the efficiency drastically drops down to 20% with  $N_p = 32$  nodes. For the mesh size used in this test, we observe that the Hybrid/MPI implementation performs very well on  $N_p < 16$  processes, providing the lowest execution times, but with  $N_p \geq 16$  processes its execution times become larger than for the CUDA/MPI implementation (especially for real-time propagation programs), and eventually even for the OpenMP/MPI implementation, as illustrated in Figure 8.8. This is again due to the insufficient amount of work each computing node performs after data are distributed among the MPI processes, which can be seen from the amount of data offloaded to GPU for all tested values of  $N_p$  in Table 8.12. Similar saturation in performance will eventually happen for the OpenMP/MPI and CUDA/MPI implementations, just for larger values of  $N_p$ .

We thus conclude that Hybrid/MPI implementation has the best performance of the three MPI-based implementations if the amount of work per process remains high enough to justify the use of hybrid algorithm. Otherwise, either of the other two MPI-based implementations should be considered first. The energy efficiency of this and other MPI-based implementations was not explored due to the difficulty in making precise measurements. If the energy consumption is not an issue, Hybrid/MPI implementation will yield the best performance, providing the cluster has powerful GPUs installed and the mesh size used is large enough.

We now model the speedup and strong scaling efficiency of MPI-based programs. In general, the execution time of a program can be expressed as

$$T(N_p) = \alpha + \beta L + \gamma \frac{M}{N_p} + \beta V \frac{M}{N_p^2}, \quad (8.10)$$

where  $\alpha$  represents the average time to perform serial portion of the code,  $L$  is the communication latency associated with one MPI message,  $\beta$  is the frequency of MPI messaging,  $\gamma$  is the average time to process one spatial point,  $V$  is related to data transfer speed (throughput), and  $M$  is the mesh size ( $N_x \times N_y \times N_z$ ). The communication overhead of one all-to-all message passing instance is equal to  $L + VM/N_p^2$ , where each of  $N_p$  processes communicates its  $M/N_p$  part of the mesh evenly to all other processes, leading to a message size of  $M/N_p^2$ . Taking into account that in the strong

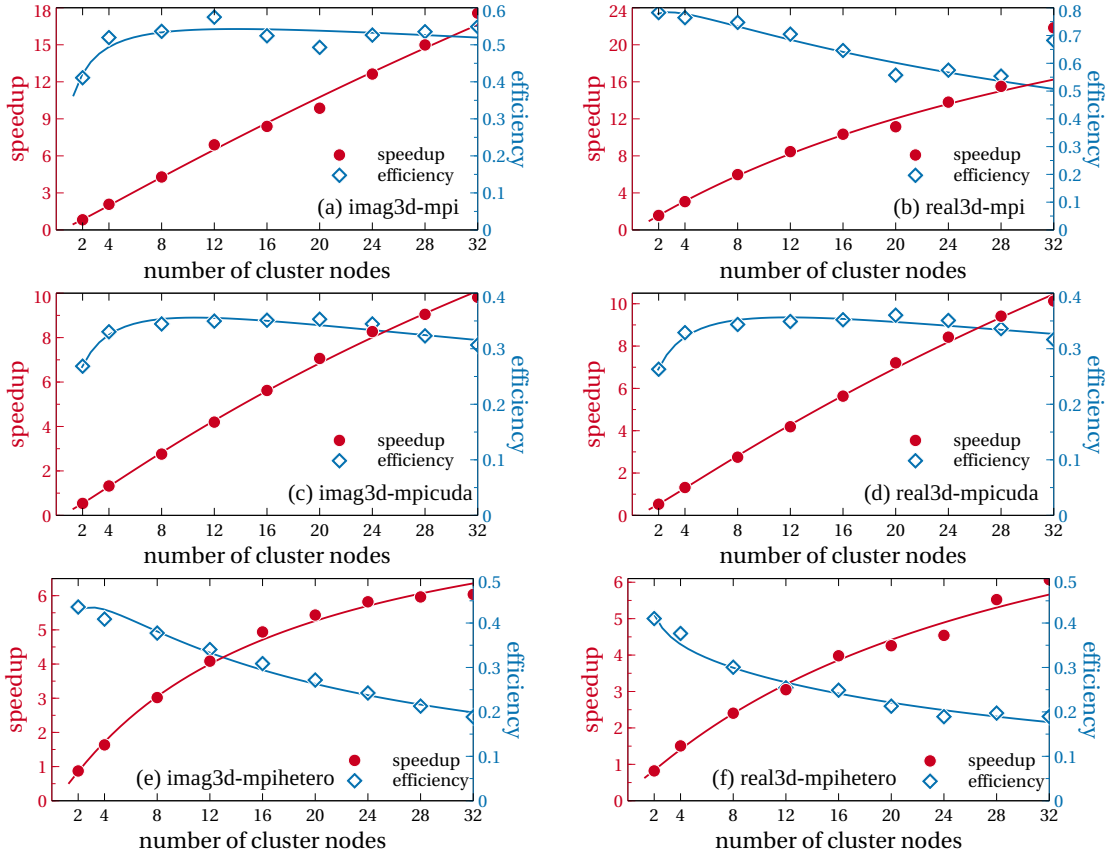


Figure 8.7: Speedup in the execution time and strong scaling efficiency of MPI-based programs compared to single-node runs: (a) imag3d-mpi, (b) real3d-mpi, (c) imag3d-mpicuda, (d) real3d-mpicuda, (e) imag3d-mpihetero, (f) real3d-mpihetero. Solid lines represent fits to measured data, where fit model functions are given in the text and obtained fit parameters are listed in Table 8.13. Note that the model parameters are fitted only to the speedup data, and then used to plot the efficiency model curve.

scaling tests the mesh size is fixed, the above model can be simplified to

$$T(N_p) = T(1) \left( a + \frac{b}{N_p} + \frac{c}{N_p^2} \right), \quad (8.11)$$

while the speedup can be modeled by

$$S(N_p) = \frac{T(1)}{T(N_p)} = \frac{1}{a + b/N_p + c/N_p^2}. \quad (8.12)$$

This model is fitted to the obtained strong scaling measurement data and the results are presented in Table 8.13, while the corresponding model curves are shown as solid lines in Figure 8.7. As we see, the proposed model agrees very well with the experimental data. The only exception is the speedup and efficiency of real-time OpenMP/MPI program (real3d-mpi) on  $N_p = 32$  processes, where FFTW library creates an optimal transform plan that works very well with the given mesh size. However, changing the mesh size even slightly gives the performance comparable to the model prediction. We note that the value of the model parameter  $c$  is negative for real-time Hybrid/MPI program (real3d-mpihetero), which should not be the case since  $c$  is parameter is related to data



Table 8.12: Optimal fraction of total data offloaded to GPU per process for different values of  $N_p$  in Hybrid/MPI programs, obtained using the GA optimization method in strong scaling tests.

$N_p$	GPU portion of data [%]							
	imag3d-mpihetero				real3d-mpihetero			
	PS <sub>1</sub>	PS <sub>2</sub>	PS <sub>3</sub>	PS <sub>4</sub>	PS <sub>1</sub>	PS <sub>2</sub>	PS <sub>3</sub>	PS <sub>4</sub>
1	30	17.5	34.4	36.7	45.2	32.5	55	35
2	25	10	17.5	16.7	20.8	11.5	23.8	16
4	10	5.2	9.2	9.4	10	5.2	15	6.7
8	5	2.9	5.8	6.7	6.3	2.1	6.7	5
12	3.3	1.7	4.2	4.2	4	1.3	5	3.3
16	2.9	1.3	2.9	2.1	3.1	0.4	2.9	2.1
20	1.7	0.6	2.5	1	2.5	0.4	2.5	1.7
24	1.7	1.3	2.1	1.7	2.1	1.3	2.5	0.8
28	1.7	0.4	2.1	1.3	1.7	0.4	1.7	1.3
32	0.8	0.4	1.7	0.8	1.5	0.4	2.1	0.8




Figure 8.8: Speedup in the execution time of Hybrid/MPI programs in strong scaling tests compared to the other two MPI-based programs: (a) OpenMP/MPI, (b) CUDA/MPI.

transfer speed. However, due to uncertainty of the optimization choice by the GA method and the fact that the fraction of total data offloaded to GPU gradually decreases (Table 8.12), leading from real hybrid algorithm to almost pure OpenMP/MPI one with hybrid FFT, it is not surprising that the obtained value differs from expected. Taking into account relatively large fit error  $\Delta c/c = 60\%$  for this parameter, we can still use expression (8.12) to model performance of Hybrid/MPI programs.

Next, we test weak scaling of MPI-based implementations. The same number of cluster nodes (and thus MPI processes) was used as in previous tests, while the starting mesh, which corresponds to a unit workload, had a size of  $480 \times 480 \times 480$ , amounting to 110,595,000 spatial points. This number of spatial points was kept constant per process, similarly to the weak test of OpenMP programs. There is an exception to this scheme, the case when programs are executed on  $N_p = 28$  nodes, as the scaled number of spatial points cannot be evenly distributed among 28 processes. In this case we compare the weak scalability by scaling up a starting mesh of  $476 \times 476 \times 476$ , which has 107,850,176 spatial points, just slightly less than for other values of  $N_p$ . The base of comparison was the execution time of a single iteration of the main loop with the corresponding MPI-based program running as a single process ( $N_p = 1$ ). While we used this configuration as the baseline, we do not recommend running MPI-based programs with  $N_p = 1$  processes, because no

Table 8.13: Values of obtained strong scaling model fit parameters  $a$ ,  $b$  and  $c$ , as well as the estimated fit errors for MPI-parallelized programs.

Program	$a$	$\Delta a$	$b$	$\Delta b$	$c$	$\Delta c$
imag3d-mpi	0.008	0.006	1.6	0.2	1.4	0.4
real3d-mpi	0.026	0.006	1.1	0.1	0.2	0.3
imag3d-mpicuda	0.025	0.008	2.3	0.2	2.8	0.6
real3d-mpicuda	0.020	0.008	2.3	0.2	2.8	0.6
imag3d-mpihetero	0.104	0.009	1.7	0.2	0.9	0.5
real3d-mpihetero	0.09	0.02	2.7	0.3	-1.0	0.6

special handling of such case has been implemented. This means that MPI-parallelized programs are always transposing the data, which is unnecessary with  $N_p = 1$  as all data is local to the single process. Instead, we recommend using single-node variants of the programs outside cluster environment.

Table 8.14 lists mesh sizes used and execution times obtained for  $N_p = 1, 2, 4, 8, 12, 16, 24$  and 32 cluster nodes. We also tested weak scaling on  $N_p = 20$  and 28, but have excluded them from the table for brevity. A complete comparison is shown in Figure 8.9. The ordering of mesh dimensions had a slightly greater impact than with single-node OpenMP programs, varying up to 10% for the OpenMP/MPI programs, and less for the other two MPI-based implementations. This means that for meshes of the same size, thus implying equal work, execution times were mostly lower for those with larger values of  $N_x$  (e.g., the execution time for a  $960 \times 480 \times 960$  mesh is lower than for a  $480 \times 960 \times 960$  one). However, no distinct pattern emerges that would give us a clue as to which order is the most favorable. Upon inspection, we find that the difference in execution times is due to the different FFT plans employed by the FFTW and cuFFT libraries. The communication time remains mostly the same, as is expected since the amount of data exchanged is the same.

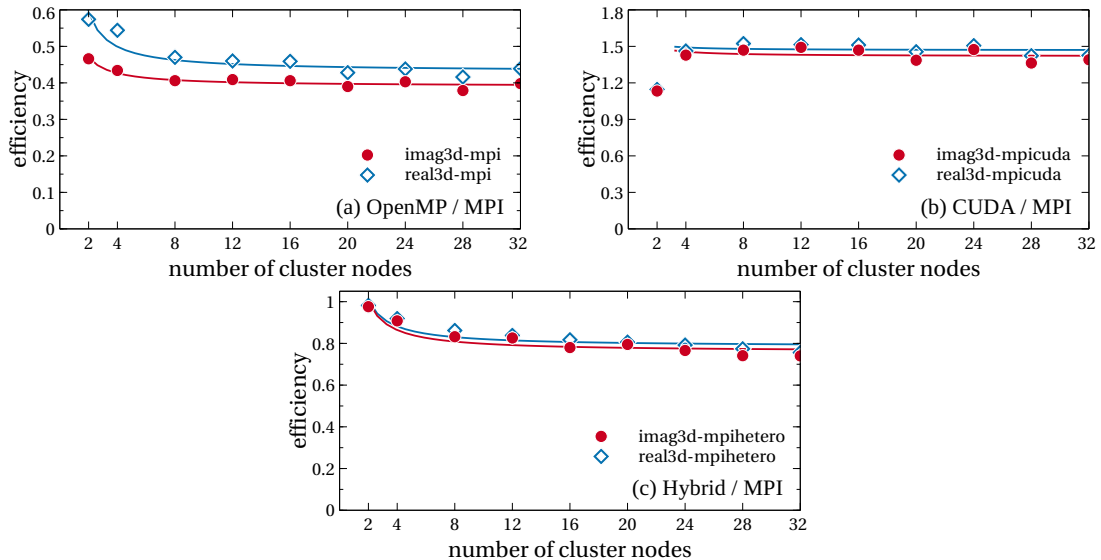


Figure 8.9: Weak scaling efficiency of three MPI-based implementations, averaged over all mesh sizes tested for each value of  $N_p$ . Solid lines represent fits to measured data, where fit model functions are given in the text.

Table 8.14: Weak scaling of MPI-based programs. Wall-clock execution times  $T_W$  are given in milliseconds, efficiencies  $E_W$  in percents.

Mesh size	imag3d-mpi		real3d-mpi		imag3d-mpicuda		real3d-mpicuda		imag3d-mpihetero		real3d-mpihetero	
	$T_W$	$E_W$	$T_W$	$E_W$	$T_W$	$E_W$	$T_W$	$E_W$	$T_W$	$E_W$	$T_W$	$E_W$
$N_p = 1$												
$480 \times 480 \times 480$	1758	100	2956	100	4338	100	6346	100	1971	100	2877	100
$N_p = 2$												
$480 \times 480 \times 960$	3858	45.6	4910	60.2	3785	114.6	5430	116.9	2044	96.4	2878	100
$480 \times 960 \times 480$	3704	47.5	5035	58.7	3825	113.4	5573	113.9	1997	98.7	2928	98.3
$960 \times 480 \times 480$	3758	46.8	5506	53.7	3871	112.1	5609	113.1	2019	97.6	2980	96.5
$N_p = 4$												
$480 \times 960 \times 960$	4267	41.2	6033	49.2	2995	144.8	4278	148.3	2149	91.7	3009	95.6
$960 \times 480 \times 960$	3877	45.4	5223	56.6	3042	142.6	4318	146.9	2218	88.9	3274	87.9
$960 \times 960 \times 480$	4000	44	5063	58.4	3078	140.9	4424	143.4	2148	91.8	3106	92.6
$N_p = 8$												
$960 \times 960 \times 960$	4327	40.6	6285	47	2952	146.9	4167	152.3	2369	83.2	3338	86.2
$N_p = 12$												
$1440 \times 960 \times 960$	4400	40	6729	43.9	2902	149.5	4254	149.2	2335	84.4	3504	82.1
$960 \times 1440 \times 960$	4268	41.2	6263	47.2	2953	146.9	4167	152.3	2412	81.7	3440	83.6
$960 \times 960 \times 1440$	4234	40.2	6299	46.9	2865	151.4	4153	152.8	2413	81.7	3351	85.9
$N_p = 16$												
$1920 \times 960 \times 960$	4356	40.4	6419	46	3001	144.5	4278	148.3	2392	82.4	3434	83.8
$960 \times 1920 \times 960$	4266	41.2	6376	46.4	2970	146.1	4170	152.2	2552	77.2	3525	81.6
$960 \times 960 \times 1920$	4234	40.2	6508	45.4	2888	150.2	4148	153	2639	74.7	3606	79.8
$N_p = 24$												
$1440 \times 1920 \times 960$	4677	40	6895	42.9	2924	148.4	4266	148.8	2552	77.3	3770	76.3
$1440 \times 960 \times 1920$	4796	38.3	7076	41.8	2984	145.4	4235	149.8	2611	75.5	3699	77.8
$1920 \times 1440 \times 960$	4470	40.6	6439	45.9	3007	144.2	4286	148.1	2546	77.4	3504	82.1
$1920 \times 960 \times 1440$	4450	40.8	6425	46	2984	145.3	4251	149.3	2551	77.3	3550	81.1
$960 \times 1440 \times 1920$	4364	41.1	6846	43.2	2879	150.7	4105	154.6	2682	73.5	3725	77.2
$960 \times 1920 \times 1440$	4346	41	6806	43.4	2878	150.7	4133	153.5	2501	78.8	3562	80.8
$N_p = 32$												
$1920 \times 1920 \times 960$	4452	39.5	6704	44.1	3183	136.3	4518	140.5	2585	76.3	3715	77.5
$1920 \times 960 \times 1920$	4520	38.9	6802	43.4	3163	137.2	4480	141.7	2632	74.9	3769	76.3
$960 \times 1920 \times 1920$	4301	40.9	6663	44.4	3014	143.9	4351	145.9	2776	71	3933	73.0

The results of weak scaling tests warrant further discussion as the execution times and plotted efficiencies appear misleading. On one hand, OpenMP/MPI programs show poor scaling results, achieving only 40% efficiency, while on the other hand the CUDA/MPI show efficiency well above 100%, a seemingly impossible result. The lower efficiency of the OpenMP/MPI programs is due to the very good performance of the baseline run. Here, the FFTW creates a very good FFT and transpose plans, which exploit the fact that all data are local, and thus have very good execution time. With  $N_p = 2$  nodes the execution time increases significantly, and efficiency drops to about 45-55%. However, adding more work and nodes has much smaller impact on efficiency, which only drops to 40-45% with  $N_p = 32$  nodes. If we use  $N_p = 2$  as the baseline, we see that the efficiency remains above 80%, a very good result.

In contrast to the OpenMP/MPI version, CUDA/MPI version performs very badly when executed on one node. Due to message sizes, MPI implementation used (Open MPI) relies on asynchronous copies through host memory [94] to perform the required communication. This results in bad performance due to a lack of overlap when copying data, as there is only one process involved. With  $N_p = 2$  processes the situation improves, as multiple streams are used, much like how we employ streams to overlap computation and data transfers in implementation of the hybrid algorithm. Between  $N_p = 4$  and 16 processes we get the best performance, and further increasing the number of processes only slightly reduces efficiency due to the size and amount of messages passed. If we compare the weak scaling efficiency with the best result as the baseline (obtained with  $N_p = 16$  processes), we see that the efficiency is lower than for the OpenMP/MPI version, due to the different transpose routine used and divided multidimensional FFT that has to be employed. However, we stress that, in terms of absolute execution times, CUDA/MPI is faster than OpenMP/MPI version for  $N_p \geq 4$ .

The Hybrid/MPI version behaves as expected and demonstrates weak scaling efficiency of about 70-75% with  $N_p = 32$  nodes. Since data are transposed only in host memory, this version does not suffer the penalty of memory copies like the CUDA/MPI version, and therefore achieves much better efficiency with small number of cluster nodes. In terms of absolute execution times, Hybrid/MPI version is the fastest of all three implementations for  $N_p \geq 2$  and is therefore the algorithm of choice for distributed memory systems. Of course, it assumes previous optimization of offloading parameters using the GA method, which itself is time-consuming and has to be taken into consideration when making the choice. Table 8.15 provides data on the optimal fraction of total data offloaded to GPU in our weak scaling tests. As the number of processes increases, one expects that the amount of data offloaded to GPU remains constant, which is the case, as we see from the table. The only exceptions are when the value of  $N_y$  is greater than  $N_x$  where we see a larger fraction of data offloaded to GPU in PS<sub>2</sub> and PS<sub>4</sub>, and when the value of  $N_y$  is smaller than  $N_x$  where we see a smaller fraction of data offloaded to GPU in PS<sub>2</sub> and PS<sub>4</sub>.

Table 8.15: Optimal fraction of total data offloaded to GPU in weak scaling tests for different mesh sizes in Hybrid/MPI programs, obtained using the GA optimization method.

Mesh size	imag3d-mpihetero				real3d-mpihetero			
	PS <sub>1</sub>	PS <sub>2</sub>	PS <sub>3</sub>	PS <sub>4</sub>	PS <sub>1</sub>	PS <sub>2</sub>	PS <sub>3</sub>	PS <sub>4</sub>
$N_p = 1$								
480 × 480 × 480	35	20	31	44	46	30	40	35
$N_p = 2$								
480 × 480 × 960	46	33	33	47	38	29	37	38
480 × 960 × 480	50	67	17	47	45	60	19	40
960 × 480 × 480	50	15	70	34	43	15	67	35
$N_p = 4$								
480 × 960 × 960	40	65	18	47	40	63	17	37
960 × 480 × 960	50	10	67	42	40	12	75	33
960 × 960 × 480	33	29	33	40	45	22	33	40
$N_p = 8$								
960 × 960 × 960	53	25	33	47	40	27	40	28
$N_p = 12$								
1440 × 960 × 960	37	20	53	45	42	18	60	63
960 × 1440 × 960	40	38	27	47	45	35	27	35
960 × 960 × 1440	50	25	38	45	45	25	40	40
$N_p = 16$								
1920 × 960 × 960	47	13	67	27	42	13	80	30
960 × 1920 × 960	47	47	35	37	47	53	20	25
960 × 960 × 1920	50	20	33	33	47	20	40	27
$N_p = 24$								
1440 × 1920 × 960	47	40	30	53	47	33	25	63
1440 × 960 × 1920	50	17	60	45	47	13	60	70
1920 × 1440 × 960	40	20	50	40	45	20	53	27
1920 × 960 × 1440	53	13	60	45	45	10	80	25
960 × 1440 × 1920	50	30	27	47	45	30	27	27
960 × 1920 × 1440	40	50	18	45	45	50	20	30
$N_p = 32$								
1920 × 1920 × 960	47	27	33	30	47	17	40	27
1920 × 960 × 1920	50	13	60	40	47	10	80	40
960 × 1920 × 1920	47	47	40	33	47	40	23	27

To model the obtained weak scaling results, we start from equation (8.10) for the execution time of a single iteration of the main loop. In weak scaling tests, the mesh size is increased proportionally to the number of MPI processes  $N_p$ , and therefore the execution time is given by

$$T_W(N_p) = \alpha + \beta L + \gamma \frac{MN_p}{N_p} + \beta V \frac{MN_p}{N_p^2} = \alpha + \beta L + \gamma M + \beta V \frac{M}{N_p}. \quad (8.13)$$

The weak scaling efficiency

$$E_W(N_p) = \frac{T_W(1)}{T_W(N_p)}, \quad (8.14)$$

can thus be modeled by

$$E_W(N_p) = \frac{1}{a + b/N_p}. \quad (8.15)$$

We fitted this model to data presented in Figure 8.9 and the obtained fit parameters are given in Table 8.16. Note that the model is fitted to CUDA/MPI scaling data only for  $N_p \geq 4$ . The agreement between the model and measurement data is very good, although we note that the parameter  $b$  has negative values, contrary to its expected relation to the communication cost. Therefore, we conclude that the above expression can be successfully used for modeling of the performance of all three MPI-based implementations.

Table 8.16: Values of obtained weak scaling model fit parameters  $a$  and  $b$ , as well as the estimated fit errors for MPI-parallelized programs.

Program	$a$	$\Delta a$	$b$	$\Delta b$
imag3d-mpi	2.56	0.06	-0.9	0.3
real3d-mpi	2.32	0.05	-1.3	0.2
imag3d-mpicuda	0.70	0.02	-0.1	0.2
real3d-mpicuda	0.68	0.02	-0.1	0.2
imag3d-mpihetero	1.32	0.03	-0.6	0.2
real3d-mpihetero	1.27	0.03	-0.6	0.2

To summarize, in all test results of MPI-based programs, we note an expected declining of the scaling efficiency. This is due to the introduction of distributed transposes of data, creating overhead that negatively impacts scaling efficiency. It is most evident in both strong and weak scaling tests of the CUDA/MPI version, as the transpose algorithm is inferior to the one provided by FFTW, used in OpenMP/MPI implementation. In our tests, all three MPI versions of programs failed to achieve actual speedup ( $S(N_p) > 1$ ) on less than  $N_p = 4$  nodes, due to the introduction of these transpose routines. We therefore recommend using MPI versions only on  $N_p = 4$  or more cluster nodes.

## 8.5 Selecting optimal algorithm

In previous sections we presented results of detailed performance tests and associated models for all developed programs. Here we provide general guidelines for obtaining the best performance from each implementation.

We note that the extensive testing performed shows that the best performance can be achieved by evenly distributing the workload among the MPI processes and OpenMP threads, and by using mesh sizes which are optimal for FFT. In particular, the single-node OpenMP programs have the best performance if the number of spatial points in all directions, controlled by parameters  $N_x$ ,  $N_y$  and  $N_z$ , is divisible by the number of OpenMP threads used. Similarly, the OpenMP/MPI implementation achieves the best performance if  $N_x$  and  $N_y$  are divisible by a product of the number of MPI processes and the number of OpenMP threads used.

On the other hand, CUDA implementation works best if all the data for a given mesh size can fit into the GPU memory with the `POTMEM` parameter set to 2. In the case of a small mesh size, CUDA programs may not be able to saturate all Streaming Multiprocessors (SM) of a given GPU, and in this case the OpenMP programs may be considered first. For CUDA/MPI programs, the best performance is achieved if  $N_x$  and  $N_y$  are divisible by a product of the number of MPI processes and the number of SMs in the GPU used. Note that CUDA-based implementations write output files (e.g., density profiles) by transferring data from GPU memory to host memory, where a single thread writes to a file. Therefore, these implementations may not be suitable for simulations that require the output to be written frequently, after a small number of time propagation steps.

Best performance of hybrid implementations can be obtained by following the same guidelines. Furthermore, we recommend using these programs with the optimization methods described in Section 8.1 to optimally divide the work between host and device. If manual tuning of offload parameters is required, we recommend that the guidelines above are followed for both CPU and GPU portions of the mesh. In the case of large disparity in the performance of host or device, hybrid versions will not provide the lowest execution times, and in such cases the pure CPU or GPU implementations with or without MPI could be better suited.

In addition to the guidelines for mesh sizes presented above, all programs benefit from the mesh size which is also optimal for FFT. The best FFT performance on CPU with FFTW library is obtained if  $N_x$ ,  $N_y$  and  $N_z$  can be expressed as  $2^a 3^b 5^c 7^d 11^e 13^f$ , where  $e$  and  $f$  are either 0 or 1, and the other exponents are non-negative integer numbers [95]. Similarly, for cuFFT the best performance is achieved for transform sizes of the form  $2^a 3^b 5^c 7^d$  [96]. In hybrid implementations, the same applies to the host and device portions of  $N_x$  and  $N_y$ , i.e., `cpuNx`, `cpuNy`, `gpuNx` and `gpuNy`.

## Chapter 9

# Demonstration of usability of developed programs

The programs developed as part of this thesis can be used to model and study a variety of systems. In this chapter we demonstrate the versatility of developed programs on a simulation of a BEC. We use the MPI-based programs to simulate the effects of a moving obstacle in an oblate atomic BEC. The obstacle, a repulsive Gaussian laser beam, moves through the condensate and sheds quantum vortices, elementary excitation of a superfluid [10, 11]. The vortices appear only when the obstacle is moving above some critical speed, consistent with the Landau's criterion of superfluidity. At low obstacle velocities above a critical value, vortex dipoles emerge, and as the speed is increased further, individual vortices and rotating vortex pairs are also formed.

In Section 9.1 we document how our programs can be used to simulate the experiment reported in Ref. [97], and compare experimental and numerical results. The experiment measured critical velocity for the emergence of vortex dipoles and rotating vortex pairs in a BEC of sodium atoms ( $^{23}\text{Na}$ ), which do not exhibit the dipolar interaction. Therefore, in this section we also demonstrate how our programs can be modified to switch off the dipolar interaction term when solving GPE. The agreement of the results obtained numerically and experimental observations provides an external check of the correctness of our algorithms and their implementations.

Next, in Section 9.2, we investigate the formation of vortices in a dipolar BEC of dysprosium atoms ( $^{164}\text{Dy}$ ). Our simulations follow the same methodology as the experiment with sodium atoms, only with atomic species exhibiting strong dipolar interaction. We study effects of the dipolar interaction on the critical velocity for the emergence of vortices, as well as the interplay between contact and dipolar interaction. The visualization extensions presented in Chapter 7 have proven to be indispensable during these simulations, as they allow much easier study of the results and control of the simulation.

### 9.1 Formation of vortices in BEC

BEC is a superfluid quantum liquid and one of its hallmarks are quantized vortices, which appear as elementary excitations of the system. Quantization of vortices is connected to excitation spectrum and, according to Landau's criterion, leads to the existence of a minimal velocity an obstacle moving



through the superfluid has to have in order to generate such elementary excitations. This critical velocity  $v_c$  can be experimentally measured and, in principle, depends on the experimental protocol used.

Reference [97] reports the study of vortex formation in an oblate BEC of  $N_{\text{at}} = 3.2 \times 10^6$  sodium atoms. The trapping potential frequencies used are  $(\omega_x, \omega_y, \omega_z) = 2\pi \times (9, 9, 400)$  Hz, and the  $s$ -wave scattering length determining the strength of contact interaction was  $a_s = 51.9 a_0$ , where  $a_0 = 0.0529$  nm is the Bohr radius. The moving obstacle is realized by a repulsive Gaussian laser beam corresponding to an additional potential of the form

$$V_B(\mathbf{r}; t) = V_0 \exp \left\{ -2 \frac{[y - y_0(t)]^2 + z^2}{\sigma^2} \right\}, \quad (9.1)$$

where  $V_0$  represents the strength of the beam,  $\sigma$  is the  $1/e^2$  beam waist, while  $y_0(t) = vt$  determines the center of the beam  $(0, y_0(t), 0)$ . In our simulations, as in the experiment, the initial position of the beam is at the condensate center and then it moves with the velocity  $v$  along  $y$  direction for  $24 \mu\text{m}$ . Afterwards, the beam is switched off linearly during 0.5 s. Illustration of the experimental setup is given in Figure 9.1. Using imaginary-time propagation, we calculate the ground state of the system with the trapping potential  $V(\mathbf{r}) + V_B(\mathbf{r}; t)$ , where harmonic part  $V(\mathbf{r})$  is defined by Eq. (2.2), with the frequencies given above. The experiment measured critical velocity for the emergence of vortex dipoles, pairs of vortices of the opposite sign. We numerically addressed this setup and calculated this critical velocity for the values of the parameters  $\sigma = 1.3l$  and  $V_0 = 250 \hbar\tilde{\omega}$ , where  $l = 6.98777 \mu\text{m}$  is the harmonic oscillator length for the referent frequency  $\tilde{\omega} = \omega_x$ .




Figure 9.1: Illustration of the experimental setup used in Ref. [97] to study vortex formation in a BEC of sodium atoms. A repulsive Gaussian laser beam is initially at the center of the condensate and moves along  $y$  direction with a constant velocity  $v$  until it reaches its final position. The beam is then switched off during a period of 0.5 s. For sufficiently large velocity vortices and vortex pairs are generated.

Before it was possible to run simulations with the parameters described above, we had to make changes to the programs to model the experiment precisely. First, we had to exclude the dipolar interaction term from the propagation of the wave function, since sodium atoms have negligible

dipole moment. While we can set the corresponding input parameters to zero, this would still result in unnecessary computation. A better approach is to remove (or comment out) the portions of the code relating to the calculation of the dipolar interaction term (calling of `calcpsidd2` function and, optionally, the associated allocations and deallocations), and adjust the `calcnu` function not to include the dipolar interaction term. This is a very simple change, resulting in modification of a few lines of the code, however, the performance impact is significant, due to the code calling the FFT functions no longer being part of the main loop. A single iteration of the main loop now performs roughly 30-40% faster.

Final change was related to the implementation of time-dependent potential  $V(\mathbf{r}) + V_B(\mathbf{r}; t)$ . Initialization of  $V(\mathbf{r})$  remains the same, i.e., it is initialized before the main loop and stored in an array, while calculation of  $V_B(\mathbf{r}; t)$  is implemented in a separate function that is called in each iteration within the main loop. Two distinct phases of  $V_B(\mathbf{r}; t)$  exist: one relating to the movement of the beam, and the other relating to the beam shutdown during 0.5 s. Each phase is implemented as a separate function, and the active phase is determined by the current time, i.e., iteration number.

Figure 9.2 shows several snapshots of a typical dynamical evolution of the system. Imaginary-time propagation yields ground state of the system, which is shown in Figure 9.2(a), where we plot integrated 2D density profile in  $y$ - $z$  plane. To model local inhomogeneities always present in the experiment, we add uniformly distributed random noise to the wave function of the order of 10%. Such modified ground state represents initial state of the system, and the beam starts to move along  $y$  direction at time  $t = 0$  with the speed  $v = 1.26$  mm/s. Figures 9.2(b) and 9.2(c) show 2D density profile of the system at times  $t = 52.17$  ms and  $t = 358.1$  ms, respectively. In Figure 9.2(b) the beam already reached its final position and its switch-off started. We can observe that several vortices (vortex dipoles) are generated and that the used speed exceeds the critical velocity. Generated vortices are stable and can be seen after extended period of time, Figure 9.2(c). To confirm that we indeed have vortices and not just localized density minima, we plot the  $x = 0$  slice of the phase of the wave function in Figure 9.3. We observe characteristic braiding and jumps of the phase in the vicinity of density minima, which is a well-known hallmark of a vortex.




Figure 9.2: Time evolution of 2D density profile of BEC of sodium atoms with an obstacle moving at speed  $v = 1.26$  mm/s  $> v_c$ , when several vortices are generated. Each panel shows integrated 2D density profile in  $y$ - $z$  plane at different time  $t$ . All lengths are expressed in units of  $\mu\text{m}$ , and particle density is given in units  $N_{\text{at}}/l^2$ .




Figure 9.3: Time evolution of the wave function phase of BEC of sodium atoms with an obstacle moving at speed  $v = 1.26 \text{ mm/s} > v_c$ , when several vortices are generated. Each panel shows  $x = 0$  slice of the phase of the wave function in  $y$ - $z$  plane at different time  $t$ . All lengths are expressed in units of  $\mu\text{m}$ .

When the obstacle is moving with an under-critical velocity, no vortices are generated, as in Figure 9.4. The velocity  $v = 0.87 \text{ mm/s}$  is just slightly lower than the the critical one, and we can see precursors of vortices at the beam edge in Figures 9.4(a) and 9.4(b). At higher velocities, vortices emerge, as we have seen in Figure 9.2. At even higher velocities, Figure 9.5, rotating vortex pairs are generated, in addition to individual vortices.




Figure 9.4: Time evolution of 2D density profile of BEC of sodium atoms with an obstacle moving at speed  $v = 0.87 \text{ mm/s} \lesssim v_c$ , when no vortices are generated. Each panel shows integrated 2D density profile in  $y$ - $z$  plane at different time  $t$ . All lengths are expressed in units of  $\mu\text{m}$ , and particle density is given in units  $N_{\text{at}}/l^2$ .

The obtained results are in good agreement with experimental findings of Ref. [97] and show that the programs developed within this thesis can be successfully used to model BEC systems with contact interaction, even for the most complex setup, when vortices are generated due to a moving obstacle.




Figure 9.5: Time evolution of 2D density profile of BEC of sodium atoms with an obstacle moving at speed  $v = 1.4 \text{ mm/s} \gg v_c$ , when rotating vortex pairs are also generated. Each panel shows integrated 2D density profile in  $y$ - $z$  plane at different time  $t$ . A rotating vortex pair is highlighted in panel (c). All lengths are expressed in units of  $\mu\text{m}$ , and particle density is given in units  $N_{\text{at}}/l^2$ .

## 9.2 Effects of dipolar interactions on vortex formation in BEC

To further test our programs, we model vortex formation in a BEC of dysprosium atoms ( $^{164}\text{Dy}$ ), which exhibit strong magnetic dipole moment. In particular, we study effects of the dipolar interaction strength for varying  $s$ -wave scattering lengths on the critical velocity for the emergence of vortices, which was not investigated experimentally. Dysprosium atoms have the largest magnetic dipole moment ( $m = 10\mu_B$ ) available in ultracold atom experiments, corresponding the characteristic dipole-dipole interaction length  $a_{\text{dd}} = 132 a_0$ . External magnetic field can be used to align all atomic dipoles in the same direction, as well as to tune their strength up to a maximal value given above. The same applies to the  $s$ -wave scattering length, which determines the contact interaction strength and which can also be tuned using the Feshbach resonance technique. Following Ref. [98], we take the same range of possible values for  $a_s$  as for  $a_{\text{dd}}$ , i.e., up to a maximal value of  $132 a_0$ . With this we show how the developed programs can be used to model and theoretically address new physical phenomena, before they are studied experimentally. Such approach is essential for the design of many upcoming experiments, since otherwise it would be extremely difficult to predict what would be the relevant range of physical quantities to be measured, thus making it very challenging to perform the experiments. Having the results of detailed numerical simulations for the particular system enables experiments to target appropriate range of all relevant quantities, and to focus on discovering new phenomena.

The experiment follows the same methodology as in previous section and uses the same parameters where applicable. Namely, the trap frequencies remain the same, i.e.,  $(\omega_x, \omega_y, \omega_z) = 2\pi \times (9, 9, 400) \text{ Hz}$ , as well as strength of the beam  $V_0 = 250 \hbar\bar{\omega}$ . Due to the strong dipolar interactions affecting the stability of the BEC, we had to use a much smaller number of atoms,  $N_{\text{at}} = 8 \times 10^4$ . This resulted in a much smaller BEC, so we had to reduce the  $1/e^2$  beam waist to  $\sigma = l$ , where  $l = 2.617 \mu\text{m}$  is harmonic oscillator length corresponding to dysprosium atoms. Also, we positioned the beam outside of the condensate, and move it all the way to the other side, as illustrated in Figure 9.6. The potential  $V_{\text{B}}(\mathbf{r}; t)$  has the same form (9.1) as in previous section, just

the center of the beam is now given by  $y_0(t) = y_{00} + v_t$ , where  $y_{00} = -15$  in units of  $l$ .




Figure 9.6: Illustration of the setup we use to study vortex formation in a dipolar BEC of dysprosium atoms. A repulsive Gaussian laser beam is initially outside of the condensate and moves along  $y$  direction with a constant velocity  $v$  until it reaches its final position on the other side. For sufficiently large velocity vortices are generated.

Modifications to the programs presented in the previous section can be reused for this numerical experiment and we only need to re-enable previously disabled calculation of the dipolar term in the main time propagation loop. As before, the ground state of the system is obtained using imaginary-time program, which then serves as the initial state of the real-time propagation program, with uniformly distributed random noise of 10% added. To calculate critical velocity  $v_c$  for the emergence of vortices, we search for the minimal speed of the laser beam for which two vortices appear, as illustrated in Figure 9.6.

Figure 9.7 shows several 2D density profiles of a typical dynamical evolution of the system for velocity  $v = v_c = 0.16$  mm/s. The ground state, obtained through imaginary-time propagation is shown in Figure 9.7(a). The beam, initially outside of the condensate, moves along the  $y$  direction, as seen in Figures 9.7(b) through 9.7(f). In Figure 9.7(c) the precursors of the vortices form on the edges of the beam, which then separate from the beam if  $v \geq v_c$ , a situation seen in Figure 9.7(d). In Figures 9.7(e) and 9.7(f) we see that the generated vortices are stable, and survive for long propagation times.

Figure 9.8 shows results of our numerical study of the dipolar interaction effects on the critical velocity. We see that for large values of the  $s$ -wave scattering length, i.e., for the contact interaction comparable or larger than the dipole interaction, effects of decreasing  $a_{dd}$  are very small and probably could not be experimentally measured. On the other hand, when contact interaction is tuned down so that the dipole interaction starts to dominate the behavior of the system, critical velocity depends much stronger on  $a_{dd}$  and could be easily measured in future experiments.

The above study represents an example on how our programs can be used to verify and compare results of current experiments, as well as to theoretically investigate new phenomena and plan future experiments.






Figure 9.7: Time evolution of 2D density profile of BEC of dysprosium atoms for  $a_s = 66a_0$  and  $a_{dd} = 44a_0$ , with an obstacle moving at speed  $v = v_c = 0.16$  mm/s, when two individual vortices are generated. Each panel shows integrated 2D density profile in  $y$ - $z$  plane at different time  $t$ . All lengths are expressed in units of  $\mu\text{m}$ , and particle density is given in units  $N_{\text{at}}/l^2$ .




Figure 9.8: Critical velocity  $v_c$  for the emergence of vortices generated by a moving obstacle as a function of the characteristic dipolar interaction length  $a_{dd}$  (in units of  $a_0$ ) for varying values of the  $s$ -wave scattering length  $a_s$ .

## Chapter 10

# Conclusions and future work

The main contribution of this thesis is the development of parallel algorithms for solving nonlinear differential equations of Gross-Pitaevskii type with a convolution integral kernel, as well as six implementations of the algorithms. The algorithms are based on the split-step, semi-implicit Crank-Nicolson method, while the convolution integral is solved using Fourier transform. Several parallelization approaches were considered, targeting both shared memory and distributed memory systems, with an emphasis on heterogeneous computing platforms. Algorithm for shared-memory architectures based on multi-core processing units is presented in Chapter 3 and its extension to specialized accelerator architecture in the form of GPUs is given in Chapter 4. Combination of the two approaches is given in Chapter 5, where we propose a hybrid algorithm targeting heterogeneous computer architectures, which fully harnesses the power of modern heterogeneous parallel resources of a single computer. While the (discrete) Fourier transform can be computed using the specialized FFT libraries that target specific hardware architecture, in a hybrid environment this is not easily achievable as no FFT library for heterogeneous architectures exists. Therefore, as part of the hybrid implementation, we have developed a heterogeneous algorithm for discrete Fourier transform, which combines the specialized multi-core and GPU FFT libraries. Our heterogeneous FFT algorithm remains general enough that it can easily be used outside the programs we developed, potentially simplifying the development of future heterogeneous algorithms which rely on Fourier transform.

The shared memory algorithms were further extended using MPI, allowing them to be executed in a distributed computing environment, e.g., on a computer cluster. This work is presented in Chapter 6. The benefit of distributed memory programs is twofold: they can be used to speed up the computation of simulations that can be executed on a single computer, and enable the simulations of higher resolution, which due to large memory requirements cannot fit in the memory of a single computer. Three implementations are given, targeting different computer cluster installations, from the ones based entirely on multi-core CPUs or GPUs, to the heterogeneous ones where the high-performing CPU and GPU are installed on the same computing node.

Two additional important contributions of this thesis are presented in Chapter 8. The first is the parameter optimization method for heterogeneous algorithms, which we developed to optimize the performance of our hybrid implementations. We have studied three approaches, based on brute-force search, gradient descent and genetic algorithm, and found the optimizer based on a genetic algorithm to be best suited to the problem. This method allows us to quickly find the optimal division of work between CPU host and GPU device, and we note that it can be adapted

to work with other heterogeneous algorithms. The second contribution relates to the assessment of the efficiency of all developed programs through extensive performance evaluation. Testing results confirm that optimization of all programs described throughout the thesis yield highly efficient implementations. This chapter also includes detailed performance models for programs at different parallelization approaches. The developed methodology can be directly used in the evaluation of the expected performance of our programs on different computing resources.

The larger volume of data produced by the higher-resolution simulations could negatively impact user's ability to analyze the data and extract new information. We have addressed this by providing extensive integration with the VisIt visualization system, presented in Chapter 7. *In situ* visualization enables higher level of interaction with the simulation, through modification of the simulation's parameters while the simulation is running, visualization of its current state and control of its course. More traditional, *post-mortem* visualization is also possible, using standardized visualization file format. The usefulness of the developed integrations can be seen in Chapter 9, in which we demonstrate the usability of the developed programs on the example of study of vortex formation in BECs with a moving obstacle. Here, the visualization of high-resolution simulations was integral to the observations and measurements that were made.

In the future, we plan to address several important topics that include multi-component BEC systems, fast-rotating BECs, as well as other types of interaction, such as spin-orbit coupling. The main challenge for the development of multi-component programs is their increased memory requirement and therefore a distributed memory approach seems suitable to address this problem. Fast-rotating BECs are described by a Hamiltonian that includes additional, first-order partial derivatives w.r.t. spatial coordinates, which will require substantial changes in the implementation of Crank-Nicolson method in all approaches. Finally, different forms of the spin-orbit coupling, which represents a hot topic in the physics of ultracold atoms, will require combining of multi-component algorithm with the improved Crank-Nicolson method to allow study of novel physical systems that exhibit this type of interaction.



# Bibliography

- [1] S. Bose, “Plancks gesetz und lichtquantenhypothese,” *Zeitschrift für Physik*, vol. 26, pp. 178–181, 1924.
- [2] A. Einstein, “Quantentheorie des einatomigen idealen gases,” *Sitzungsber. Kgl. Preuss. Akad. Wiss.*, vol. 1924, p. 261, 1924.
- [3] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell, “Observation of Bose–Einstein condensation in a dilute atomic vapor,” *Science*, vol. 269, pp. 198–201, 1995.
- [4] C. C. Bradley, C. A. Sackett, J. J. Tollett, and R. G. Hulet, “Evidence of Bose–Einstein condensation in an atomic gas with attractive interactions,” *Phys. Rev. Lett.*, vol. 75, pp. 1687–1690, 1995.
- [5] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle, “Bose–Einstein condensation in a gas of sodium atoms,” *Phys. Rev. Lett.*, vol. 75, pp. 3969–3973, 1995.
- [6] O. Morsch and M. Oberthaler, “Dynamics of Bose–Einstein condensates in optical lattices,” *Rev. Mod. Phys.*, vol. 78, pp. 179–215, 2006.
- [7] I. Bloch, “Quantum coherence and entanglement with ultracold atoms in optical lattices,” *Nature*, vol. 453, pp. 1016–1022, 2008.
- [8] R. P. Feynman, “Simulating Physics with Computers,” *Int. J. Theor. Phys.*, vol. 21, pp. 467–488, 1982.
- [9] A. Griesmaier, J. Werner, S. Hensler, J. Stuhler, and T. Pfau, “Bose–Einstein condensation of chromium,” *Phys. Rev. Lett.*, vol. 94, p. 160401, 2005.
- [10] L. P. Pitaevskii and S. Stringari, *Bose–Einstein Condensation*. Oxford University Press, 2003.
- [11] C. J. Pethick and H. Smith, *Bose-Einstein Condensation in Dilute Gases*. Cambridge University Press, 2008.
- [12] P. Muruganandam and S. Adhikari, “Fortran programs for the time-dependent Gross–Pitaevskii equation in a fully anisotropic trap,” *Comput. Phys. Commun.*, vol. 180, pp. 1888–1912, 2009.
- [13] D. Vudragović, I. Vidanović, A. Balaž, P. Muruganandam, and S. K. Adhikari, “C programs for solving the time-dependent Gross–Pitaevskii equation in a fully anisotropic trap,” *Comput. Phys. Commun.*, vol. 183, pp. 2021–2025, 2012.

- [14] B. Satarić, V. Slavnić, A. Belić, A. Balaž, P. Muruganandam, and S. K. Adhikari, “Hybrid OpenMP/MPI programs for solving the time-dependent Gross–Pitaevskii equation in a fully anisotropic trap,” *Comput. Phys. Commun.*, vol. 200, pp. 411–417, 2016.
- [15] L. E. Young-S., D. Vudragović, P. Muruganandam, S. K. Adhikari, and A. Balaž, “OpenMP Fortran and C programs for solving the time-dependent Gross–Pitaevskii equation in an anisotropic trap,” *Comput. Phys. Commun.*, vol. 204, pp. 209–213, 2016.
- [16] R. K. Kumar, L. E. Young-S., D. Vudragović, A. Balaž, P. Muruganandam, and S. Adhikari, “Fortran and C programs for the time-dependent dipolar Gross–Pitaevskii equation in an anisotropic trap,” *Comput. Phys. Commun.*, vol. 195, pp. 117–128, 2015.
- [17] V. Lončar, A. Balaž, A. Bogojević, S. Škrbić, P. Muruganandam, and S. K. Adhikari, “CUDA programs for solving the time-dependent dipolar Gross–Pitaevskii equation in an anisotropic trap,” *Comput. Phys. Commun.*, vol. 200, pp. 406–410, 2016.
- [18] V. Lončar, L. E. Young-S., S. Škrbić, P. Muruganandam, S. K. Adhikari, and A. Balaž, “OpenMP, OpenMP/MPI, and CUDA/MPI C programs for solving the time-dependent dipolar Gross–Pitaevskii equation,” *Comput. Phys. Commun.*, vol. 209, pp. 190–196, 2016.
- [19] P. Wittek and L. Calderaro, “Extended computational kernels in a massively parallel implementation of the Trotter–Suzuki approximation,” *Comput. Phys. Commun.*, vol. 197, pp. 339–340, 2015.
- [20] L. J. O’Riordan, T. Morgan, and N. Crowley, “GPUE: Phaseengineering release.” <https://github.com/mlxd/GPUE>, 2016. DOI:10.5281/zenodo.57968.
- [21] R. Caplan, “NLSEmagic: Nonlinear Schrödinger equation multi-dimensional Matlab-based GPU-accelerated integrators using compact high-order schemes,” *Comput. Phys. Commun.*, vol. 184, pp. 1250–1271, 2013.
- [22] X. Antoine and R. Duboscq, “GPELab, a Matlab toolbox to solve Gross–Pitaevskii equations I: Computation of stationary solutions,” *Comput. Phys. Commun.*, vol. 185, pp. 2969–2991, 2014.
- [23] X. Antoine and R. Duboscq, “GPELab, a Matlab toolbox to solve Gross–Pitaevskii equations II: Dynamics and stochastic simulations,” *Comput. Phys. Commun.*, vol. 193, pp. 95–117, 2015.
- [24] M. Caliori and S. Rainer, “GSGPEs: A MATLAB code for computing the ground state of systems of Gross–Pitaevskii equations,” *Comput. Phys. Commun.*, vol. 184, pp. 812–823, 2013.
- [25] Ž. Marojević, E. Göklü, and C. Lämmerzahl, “ATUS-PRO: A FEM-based solver for the time-dependent and stationary Gross–Pitaevskii equation,” *Comput. Phys. Commun.*, vol. 202, pp. 216–232, 2016.
- [26] G. Vergez, I. Danaila, S. Auliac, and F. Hecht, “A finite-element toolbox for the stationary Gross–Pitaevskii equation with rotation,” *Comput. Phys. Commun.*, vol. 209, pp. 144–162, 2016.

- [27] M. Suzuki, “Decomposition formulas of exponential operators and lie exponentials with some applications to quantum mechanics and statistical physics,” *J. Math. Phys.*, vol. 26, pp. 601–612, 1985.
- [28] H. A. van der Vorst, “Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems,” *SIAM J. Sci. Stat. Comp.*, vol. 13, pp. 631–644, 1992.
- [29] P. Wittek, “Comparing three numerical solvers of the Gross–Pitaevskii equation.” <http://peterwittek.com/gpe-comparison.html>, 2016. Accessed: 2017-04-10.
- [30] J. Crank and P. Nicolson, “A practical method for numerical evaluation of solutions of partial differential equations of the heat-conduction type,” *Math. Proc. Cambridge Philos. Soc.*, vol. 43, p. 50, 1947.
- [31] W. F. Ames, *Numerical Methods for Partial Differential Equations*. Elsevier Science, 1992.
- [32] R. Dautray and J.-L. Lions, *Mathematical Analysis and Numerical Methods for Science and Technology*, vol. 6, ch. XX, pp. 45–74. Springer, 1993.
- [33] G. C. Wick, “Properties of Bethe–Salpeter wave functions,” *Phys. Rev.*, vol. 96, pp. 1124–1134, 1954.
- [34] G. Blelloch, “Scans as primitive parallel operations,” *IEEE Trans. Comput.*, vol. 38, pp. 1526–1538, 1989.
- [35] G. E. Blelloch, “Prefix sums and their applications,” tech. rep., School of Computer Science, Carnegie Mellon University, 1990.
- [36] D. Vudragović, I. Vidanović, A. Balaž, P. Muruganandam, and S. K. Adhikari, “GP-SCL package.” [http://cpc.cs.qub.ac.uk/summaries/AEDU\\_v2\\_0.html](http://cpc.cs.qub.ac.uk/summaries/AEDU_v2_0.html), 2012. Accessed: 2017-04-10.
- [37] V. Lončar, L. E. Young-S., S. Škrbić, P. Muruganandam, S. K. Adhikari, and A. Balaž, “Repository of DBEC-GP family of programs.” <https://git.ipb.ac.rs/vloncar/DBEC-GP>, 2016. Accessed: 2017-04-10.
- [38] International Organization for Standardization (ISO), “The ANSI C standard (C99),” tech. rep., ISO/IEC, 1999.
- [39] M. Frigo and S. G. Johnson, “FFTW – Fastest Fourier Transform in the West.” <http://www.fftw.org/>, 2016. Accessed: 2017-04-10.
- [40] R. L. Burden and J. D. Faires, *Numerical Analysis*, ch. 4, pp. 185–190. Brooks Cole Pub Co, 2010.
- [41] S. Pop, A. Cohen, C. Bastoul, S. Girbal, G.-A. Silber, and N. Vasilache, “GRAPHITE: Polyhedral analyses and optimizations for GCC,” in *Proceedings of the 2006 GCC Developers Summit*, p. 2006, 2006.
- [42] UPC consortium, “Unified parallel C.” <https://upc-lang.org/>, 2013. Accessed: 2017-04-10.

- [43] D. Bonachea and G. Funck, “UPC language and library specifications, version 1.3,” tech. rep., Lawrence Berkeley National Laboratory, 2013.
- [44] Intel, “Intel Cilk Plus.” <https://www.cilkplus.org/>, 2010. Accessed: 2017-04-10.
- [45] OpenMP Architecture Review Board, *OpenMP Application Programming Interface*, 4.5 ed., 2015.
- [46] V. Lončar, L. E. Young-S., S. Škrbić, P. Muruganandam, S. K. Adhikari, and A. Balaž, “DBEC-OMP-CUDA-MPI package.” <https://data.mendeley.com/datasets/j3z9z379m8/2>, 2016.
- [47] Nvidia Corporation, “CUDA home page.” [http://www.nvidia.com/object/cuda\\_home\\_new.html](http://www.nvidia.com/object/cuda_home_new.html), 2007. Accessed: 2017-04-10.
- [48] Khronos Group, “OpenCL home page.” <https://www.khronos.org/opencl/>, 2009. Accessed: 2017-04-10.
- [49] A. Danalis, G. Marin, C. McCurdy, J. S. Meredith, P. C. Roth, K. Spafford, V. Tipparaju, and J. S. Vetter, “The scalable heterogeneous computing (SHOC) benchmark suite,” in *Proceedings of the 3rd Workshop on General-Purpose Computation on Graphics Processing Units*, ACM Press, 2010.
- [50] M. J. Flynn, “Some computer organizations and their effectiveness,” *IEEE T. Comput.*, vol. C-21, pp. 948–960, 1972.
- [51] D. Kirk and W. mei Hwu, *Programming Massively Parallel Processors*. Elsevier LTD, Oxford, 2016.
- [52] J. Cheng, M. Grossman, and T. McKercher, *Professional CUDA C Programming*. John Wiley & Sons Inc, 2014.
- [53] N. Wilt, *The CUDA Handbook*. Pearson Education, 2013.
- [54] Nvidia Corporation, *cuFFT Library User’s Guide*, 7.5 ed., 2017.
- [55] M. Harris, “CUDA Pro Tip: Write Flexible Kernels with Grid-Stride Loops.” <https://devblogs.nvidia.com/paralleforall/cuda-pro-tip-write-flexible-kernels-grid-stride-loops/>, 2013. Accessed: 2017-04-10.
- [56] Stack Overflow discussion, “atomicAdd() for double on GPU.” <http://stackoverflow.com/questions/16077464/>, 2013. Accessed: 2017-04-10.
- [57] M. Harris, “Optimizing parallel reduction in CUDA.” [http://developer.download.nvidia.com/compute/cuda/1.1-Beta/x86\\_website/projects/reduction/doc/reduction.pdf](http://developer.download.nvidia.com/compute/cuda/1.1-Beta/x86_website/projects/reduction/doc/reduction.pdf), 2007. Accessed: 2017-04-10.
- [58] Stack Overflow discussion, “How can I add up two 2d (pitched) arrays using nested for loops?.” <http://stackoverflow.com/questions/6137218/>, 2011. Accessed: 2017-04-10.
- [59] V. Lončar, A. Balaž, A. Bogojević, S. Škrbić, P. Muruganandam, and S. K. Adhikari, “DBEC-GP-CUDA package.” <https://data.mendeley.com/datasets/s858p8zsgp/1>, 2016.

- [60] J. Lee, M. Samadi, Y. Park, and S. Mahlke, “Transparent CPU–GPU collaboration for data-parallel kernels on heterogeneous systems,” in *Proceedings of the 22nd International Conference on Parallel Architectures and Compilation Techniques*, PACT ’13, (Piscataway, NJ, USA), pp. 245–256, IEEE Press, 2013.
- [61] M. Papadrakakis, G. Stavroulakis, and A. Karatarakis, “A new era in scientific computing: Domain decomposition methods in hybrid CPU–GPU architectures,” *Comput. Methods Appl. Mech. Eng.*, vol. 200, pp. 1490–1508, 2011.
- [62] Y. Li, J. R. Diamond, X. Wang, H. Lin, Y. Yang, and Z. Han, “Large-scale fast Fourier transform on a heterogeneous multi-core system,” *Int. J. High Perform. Comput. Appl.*, vol. 26, pp. 148–158, 2012.
- [63] Y. Ogata, T. Endo, N. Maruyama, and S. Matsuoka, “An efficient, model-based CPU-GPU heterogeneous FFT library,” in *2008 IEEE International Symposium on Parallel and Distributed Processing*, IEEE, 2008.
- [64] S. Chen and X. Li, “A hybrid GPU/CPU FFT library for large FFT problems,” in *2013 IEEE 32nd International Performance Computing and Communications Conference (IPCCC)*, IEEE, 2013.
- [65] E. Chu and A. George, *Inside the FFT Black Box*. CRC Press, 1999.
- [66] Miroway, “Comparison between NVIDIA GeForce and Tesla GPUs.” <https://www.microway.com/knowledge-center-articles/comparison-of-nvidia-geforce-gpus-and-nvidia-tesla-gpus/>, 2016. Accessed: 2017-04-10.
- [67] Message Passing Interface Forum, “MPI: A Message-Passing Interface Standard Version 3.1,” tech. rep., University of Tennessee, Knoxville, Tennessee, 2015.
- [68] O. Ayala and L.-P. Wang, “Parallel implementation and scalability analysis of 3D Fast Fourier Transform using 2D domain decomposition,” *Parallel Comput.*, vol. 39, pp. 58–77, 2013.
- [69] N. Li and S. Laizet, “2DECOMP&FFT – a highly scalable 2D decomposition library and FFT interface,” in *Cray User Group 2010 conference, Edinburgh*, 2010.
- [70] Nvidia Corporation, “MPI solutions for GPUs.” <https://developer.nvidia.com/mpi-solutions-gpus>, 2017. Accessed: 2017-04-10.
- [71] Open MPI project, “Open MPI home page.” <https://www.open-mpi.org/>. Accessed: 2017-04-10.
- [72] U. Ayachit, *The ParaView Guide: A Parallel Visualization Application*. Kitware, Incorporated, 2015.
- [73] H. Childs, E. Brugger, B. Whitlock, J. Meredith, S. Ahern, D. Pugmire, K. Biagas, M. Miller, C. Harrison, G. H. Weber, H. Krishnan, T. Fogal, A. Sanderson, C. Garth, E. W. Bethel, D. Camp, O. Rübel, M. Durant, J. M. Favre, and P. Navrátil, “VisIt: An end-user tool for visualizing and analyzing very large data,” in *High Performance Visualization—Enabling Extreme-Scale Scientific Insight*, pp. 357–372, CRC Press, 2012.

- [74] Lawrence Livermore National Laboratory, *Silo User's Guide*, 2014. LLNL-SM-654357.
- [75] W. Schroeder, K. Martin, and B. Lorensen, *Visualization Toolkit: An Object-Oriented Approach to 3D Graphics, 4th Edition*. Kitware, 2006.
- [76] The Qt Company, "Qt application framework." [www.qt.io](http://www.qt.io). Accessed: 2017-04-10.
- [77] J. Snyman, *Practical Mathematical Optimization*. Springer, 2005.
- [78] C. Darken, J. Chang, and J. Moody, "Learning rate schedules for faster stochastic gradient search," in *Neural Networks for Signal Processing II Proceedings of the 1992 IEEE Workshop*, IEEE, 1992.
- [79] T. Schaul, S. Zhang, and Y. LeCun, "No more pesky learning rates," in *Proceedings of the 30th International Conference on Machine Learning* (S. Dasgupta and D. McAllester, eds.), vol. 28 of *Proceedings of Machine Learning Research*, (Atlanta, Georgia, USA), pp. 343–351, PMLR, 2013.
- [80] J. C. Spall, "Implementation of the simultaneous perturbation algorithm for stochastic optimization," *IEEE T. Aero. Elec. Sys.*, vol. 34, pp. 817–823, 1998.
- [81] J. C. Spall, "Stochastic optimization and the simultaneous perturbation method," in *Proceedings of the 31st conference on Winter simulation Simulation—a bridge to the future - WSC 1999*, ACM Press, 1999.
- [82] J. H. Holland, *Adaptation in Natural and Artificial Systems*. MIT University Press Group Ltd, 1992.
- [83] M. Mitchell, *An Introduction to Genetic Algorithms*. MIT Press, 1998.
- [84] M. Gen and R. Cheng, *Genetic Algorithms and Engineering Optimization*. John Wiley & Sons Inc, 1999.
- [85] A. Kumar, "Encoding schemes in genetic algorithm," *Int. J. Adv. Res. Inf. Technol. Eng.*, vol. 2, pp. 1–7, 2013.
- [86] L. J. Eshelman, R. A. Caruana, and J. D. Schaffer, "Biases in the crossover landscape," in *Proceedings of the Third International Conference on Genetic Algorithms*, (San Francisco, CA, USA), pp. 10–19, Morgan Kaufmann Publishers Inc., 1989.
- [87] V. M. Spears and K. A. D. Jong, "On the virtues of parameterized uniform crossover," in *Proceedings of the Fourth International Conference on Genetic Algorithms*, pp. 230–236, 1991.
- [88] J. N. Richter, *On mutation and crossover in the theory of evolutionary algorithms*. PhD thesis, Montana State University, 2010.
- [89] D. R. White and S. Poulding, "A rigorous evaluation of crossover and mutation in genetic programming," in *Lecture Notes in Computer Science*, vol. 5481, pp. 220–231, Springer Berlin Heidelberg, 2009.

- [90] A. P. A. da Silva and D. M. Falcao, “Fundamentals of genetic algorithms,” in *Modern Heuristic Optimization Techniques: Theory and Applications to Power Systems* (K. Y. Lee and M. A. El-Sharkawi, eds.), ch. 2, pp. 25–42, John Wiley & Sons Inc, 2008.
- [91] S. Luke and L. Spector, “A comparison of crossover and mutation in genetic programming,” in *Genetic Programming 1997: Proceedings of the Second Annual Conference* (J. R. e. a. Koza, ed.), pp. 240–248, Morgan Kaufmann, 1997.
- [92] S. Luke and L. Spector, “A revised comparison of crossover and mutation in genetic programming,” in *Genetic Programming 1998: Proceedings of the Third Annual Conference* (J. R. e. a. Koza, ed.), pp. 208–213, Morgan Kaufmann, 1998.
- [93] G. M. Amdahl, “Validity of the single processor approach to achieving large scale computing capabilities,” in *Proceedings of the April 18-20, 1967, Spring Joint Computer Conference - AFIPS '67 (Spring)*, ACM Press, 1967.
- [94] Open MPI project, “FAQ: Running CUDA-aware Open MPI.” <https://www.open-mpi.org/faq/?category=runcuda>, 2016. Accessed: 2017-04-10.
- [95] M. Frigo and S. G. Johnson, “FFTW 3 documentation, version 3.3.6.” [http://www.fftw.org/fftw3\\_doc/](http://www.fftw.org/fftw3_doc/), 2017. Accessed: 2017-04-10.
- [96] See reference [54], Section 2.12.
- [97] W. J. Kwon, G. Moon, S. W. Seo, and Y. Shin, “Critical velocity for vortex shedding in a Bose–Einstein condensate,” *Phys. Rev. A*, vol. 91, p. 053615, 2015.
- [98] H. Kadau, M. Schmitt, M. Wenzel, C. Wink, T. Maier, I. Ferrier-Barbut, and T. Pfau, “Observing the Rosensweig instability of a quantum ferrofluid,” *Nature*, vol. 530, pp. 194–197, 2016.

# Prošireni sažetak

Algoritmi za rešavanje parcijalnih diferencijalnih jednačina, posebno paralelni numerički algoritmi, danas predstavljaju veoma značajnu oblast istraživanja, posebno uzimajući u obzir da su praktično svi današnji računari zasnovani na sistemima sa više procesorskih jedinica. Višejezgarni procesori predstavljaju osnovni nivo paralelizma koji postoji u svim modernim računarima, a grafički akceleratori (engl. Graphics Processing Unit, GPU) i druge vrste koprocesora (npr. Intel Xeon Phi kartice) su sledeći nivo u ovoj hijerarhiji. Računarski klasteri, koji se sastoje od velikog broja računara povezanih brzom mrežnom infrastrukturom, uz odgovarajuća softverska rešenja koja omogućavaju da se ceo sistem koristi na paralelan način (npr. pomoću Message Passing Interface, MPI paradigme), kompletiraju ovu veoma heterogenu hijerarhiju. Efikasno korišćenje velikih računarskih sistema zahteva domen-specifična znanja i razvoj paralelnih algoritama koji uzimaju u obzir hardverske osobine sistema, kao i mogućnosti paralelizacije numeričkih metoda koji se koriste za rešavanje datog problema.

Jedna od oblasti gde su ovakvi algoritmi od velikog značaja je fizika ultrahladnih atoma, koja proučava osobine materije na veoma niskim temperaturama, reda nanokelvina. Na tim temperaturama gasovi nekih elemenata ili jedinjenja doživljavaju fazni prelaz pod nazivom Boze-Ajnštajn kondenzacija i prelaze u novo stanje, koje se odlikuje koherentnim kvantnim kolektivnim ponašanjem. Ovakvi sistemi su od velikog interesa jer pružaju ranije nezamislivu kontrolu nad eksperimentalnim parametrima i omogućavaju primene u kvantnom računarstvu, kao i simuliranje drugih kvantnih sistema (Fajnmanov kvantni simulator), opisanih sličnim jednačinama. U poslednjoj deceniji su posebno proučavani sistemi u kojima se, pored sveprisutne kontaktne interakcije, pojavljuje i dipol-dipol interakcija, a eksperimentalna realizacija Boze-Ajnštajn kondenzacije sa atomima velikog dipolnog momenta (erbijum, disprozijum) u poslednjih par godina je otvorila novu istraživačku oblast u kojoj se proučavaju dominantni efekti anizotropne dipolne interakcije.

Osnovni opis ovakvih fizičkih sistema daje teorija srednjeg polja. Dok za uobičajene kvantne sisteme možemo da koristimo linearnu parcijalnu diferencijalnu Šredingerovu jednačinu, u slučaju Boze-Ajnštajn-kondenzovanih sistema teorija srednjeg polja daje nelinearnu Šredingerovu jednačinu, odnosno Gros-Pitaevski jednačinu (GPJ). U prisustvu dipol-dipol interakcije ona sadrži nelinearni integralni kernel konvolucionog tipa, i rešavanje ovakvih jednačina zahteva primenu jednog ili više od nekoliko poznatih numeričkih metoda. U poglavlju 1 dat je kratak uvod u oblast, kao i formulacija osnovnih problema.

Ova teza se bavi razvojem i implementacijom paralelnih algoritama za rešavanje Gros-Pitaevski jednačine sa dipolnom interakcijom, bazirane na CUDA (na GPU) i OpenMP paralelizaciji (na višejezgarim procesorima) u okviru jednog računara, MPI paralelizaciji na nivou računarskog klastera, kao i kombinacije ovih pristupa. Razvijeni algoritmi omogućavaju ne samo rešavanje vremenski za-



visne GPJ, nego i vremenski nezavisne jednačine, prelaskom na propagaciju u imaginarnom vremenu. Ovakvi algoritmi omogućavaju efikasno rešavanje GPJ u velikom broju primena u fizici ultrahladnih atoma, nelinearnoj optici, primenjenoj fizici, astrofizici, kao i u kvantitativnim finansijama, u kojima se takođe koriste jednačine sličnog tipa. Tokom izrade disertacije su razvijena softverska rešenja u programskom jeziku C i detaljno izmerena poboljšanja koja ona donose u odnosu na druge pristupe. Takođe je razmotren problem *in situ* vizuelizacije podataka dobijenih tokom izvršavanja numeričkih simulacija i rešavanja vremenski zavisne i vremenski nezavisne GPJ. Sve implementacije opisane u ovoj tezi su otvorenog koda, i dostupne su za preuzimanje [37].

Za rešavanje GPJ i jednačina sličnog tipa koristi se veliki broj metoda, a kratak pregled relevantne literature je dat u poglavlju 2. Analitički se koriste različite aproksimativne metode, kao što je Tomas-Fermi aproksimacija, koja zanemaruje kinetičku energiju sistema, odnosno Laplasijan u GPJ. Koriste se i varijacioni metodi, kod kojih se za rešavanje GPJ pretpostavi funkcionalni oblik talasne funkcije, a onda se na osnovu ekstremizacionog uslova nalazi optimalno rešenje. Oba ova pristupa su popularna jer omogućavaju analitički tretman problema, ali, sa druge strane, ne daju mogućnost kontrole niti ocene greške takve aproksimacije.

Perturbativni razvoj (na primer, aproksimacija Bogoljubova, Hartri-Fok-Bogoljubova, Hartri-Fok-Bogoljubov-Popova) predstavlja sledeći nivo u analitičkom tretmanu ovakvih sistema i pruža mogućnost ocene greške, ali uz izuzetnu matematičku kompleksnost rešenja, koja u mnogim slučajevima praktično onemogućava korišćenje ovog pristupa. Postoje i malobrojni neperturbativni metodi, ali je njihova primenljivost veoma sužena na mali broj konkretnih problema (npr. solitonska rešenja u jednodimenzionalnom slučaju).

Pored analitičkih metoda, veoma značajnu ulogu imaju i numerički metodi za rešavanje GPJ, koji u velikom broju slučajeva predstavljaju jedinu mogućnost za proučavanje datih fizičkih sistema. U ove pristupe spadaju spektralne metode, metode podeljenog koraka, kao i metode konačne razlike elemenata. U slučaju GPJ sa dipolnom interakcijom, konvolucionni integral se obično rešava višestrukim korišćenjem diskretne Furijeove transformacije. Svaki od ovih metoda ima prednosti i mane kada se razmatraju njihova stabilnost i preciznost, a jedan od pristupa koji garantuje apsolutnu stabilnost i odlično skaliranje grešaka diskretizacije je semi-implicitni Krenk-Nikolson metod podeljenog koraka (KN). Ranije je slučaj kontaktne interakcije razvijen serijski algoritam koji implementira ovaj metod [12], kao i njegova paralelizovana verzija pomoću OpenMP pristupa [13]. Za slučaj GPJ sa dipol-dipol interakcijom ranije je razvijen samo serijski algoritam [16]. Ovo predstavlja odlično rešenje za jednodimenzionalne probleme, dok u dve i tri prostorne dimenzije vreme izvršavanja simulacija za sisteme u razumnoj diskretizaciji veoma brzo postaje suviše veliko da bi ovi algoritmi mogli praktično da se koriste. Zbog toga je jedan od otvorenih problema razvoj efikasnih paralelnih algoritama na heterogenim platformama koji će omogućiti brzo rešavanje GPJ za sisteme sa finom diskretizacijom u dve i tri prostorne dimenzije. Pošto je izračunavanje konvolucionog interakcionog integrala bazirano na korišćenju diskretne Furijeove transformacije, od velike koristi je upotreba GPU kartica, koje mogu značajno da ubrzaju izvršavanje ovakvih algoritama. Takođe, u trodimenzionalnim slučajevima velika količina podataka koja je neophodna za opis sistema često zahteva distribuiranje izračunavanja na više računara, odnosno primenu MPI paralelizacije. Svi ovi problemi su predmet aktivnog istraživanja u ovoj oblasti.

U poglavlju 3 predstavljen je algoritam za rešavanje GPJ sa dipolnim članom namenjen sistemima sa deljenom memorijom. Algoritam se oslanja na KN metod podeljenog koraka i omogućava

efikasne implementacije na sistemima gde je celokupan sadržaj memorije vidljiv svim računskim resursima. Opis algoritma započinjemo definicijom diskretizacione šeme, koja definiše način na koji su prostor i vreme diskretizovani pomoću konačnog broja tačaka. Razlikujemo tri slučaja, u zavisnosti od broja prostornih dimenzija problema koji rešavamo. Srž algoritma je petlja u kojoj jedna iteracija odgovara jednom vremenskom koraku propagacije. Unutar petlje, vrednosti talasne funkcije se računaju u odnosu na različite delove Hamiltonijana, u skladu sa šemom podeljenog koraka opisanom u poglavlju 2. Svaki korak glavne petlje podeljen je na više manjih koraka. U prvom od njih se računa vrednost dipolnog člana GPJ i na egzaktan način propagira vrednost talasne funkcije u odnosu na članove Hamiltonijana koji ne sadrže prostorne izvode, dok se u preostalim koracima talasna funkcija propagira u odnosu na delove Hamiltonijana sa prostornim izvodima pomoću KN metoda. Vrednost dipolnog člana se računa korišćenjem diskretne Furijeove transformacije. Prilikom propagacije u imaginarnom vremenu, neophodan je dodatni korak normalizacije talasne funkcije u svakoj iteraciji glavne petlje.

S obzirom na to da se u svakoj iteraciji više puta računa vrednost talasne funkcije u svim tačkama diskretizacione šeme, lako se može zaključiti da je algoritam računski izuzetno zahtevan. Ovo je dodatno izraženo u dve i tri dimenzije, gde ukupan broj diskretizacionih tačaka brzo raste povećanjem rezolucije, te se javlja potreba za paralelizacijom.

Osim algoritma, u poglavlju 3 su opisane i serijska i paralelna implementacija. Serijska implementacija je napisana u programskom jeziku C, sadrži 10 programa, i oslanja se na prethodne implementacije u programskim jezicima Fortran i C [12, 13, 16]. Implementirani su zasebni programi za rešavanje GPJ u jednoj, dve i tri dimenzije, u realnom i imaginarnom vremenu. Za računanje vrednosti dipolnog člana GPJ korišćena je FFTW biblioteka. Radi preglednosti, koraci u glavnoj petlji su implementirani u zasebnim funkcijama, što olakšava buduće izmene programa u skladu sa potrebama eksperimenata i drugih primena. Takođe, da bi se smanjila količina numeričkih operacija u svakoj iteraciji, vrednosti koje se ne menjaju (potencijal zamke i dipol-dipol interakcije, koeficijenti KN šeme, itd.) se izračunavaju samo jednom i smeštaju u memoriju, a zatim se po potrebi koriste tokom propagacije vrednosti talasne funkcije.

Paralelna implementacija algoritma za sisteme sa deljenom memorijom je zasnovana na OpenMP tehnologiji. Ova tehnologija omogućava paralelizaciju korišćenjem niti koje se mogu izvršavati paralelno na višejezgarnim procesorima. Nitima se ne manipuliše direktno, već se putem instrukcija u programskom kodu označavaju regioni koji će biti paralelizovani, na osnovu kojih kompajler generiše paralelni kod. Paralelizacija petlji u kojima pojedinačne iteracije nemaju međusobnih zavisnosti na ovaj način postaje veoma jednostavna. Međutim, u slučaju da se u petlji javljaju rekurzivne relacije, kao što je slučaj u funkcijama koje implementiraju KN šemu, paralelizacija je moguća jedino uz obimnu izmenu algoritma, od koje smo zbog kompleksnosti odustali. Ovo znači da u jednoj dimenziji nije moguće ostvariti paralelizaciju svakog koraka u iteraciji, što nije velik problem imajući u vidu da broj tačaka u diskretizacionoj šemi nije velik u ovim programima. U programima koji rešavaju GPJ u dve ili tri dimenzije problem rekurzivnih relacija efektivno nestaje, jer se rekurzije javljaju samo u unutrašnjoj petlji što ostavlja mogućnost da se paralelizuje spoljašnja petlja. U ovom pristupu neophodno je jedino obezbediti zaseban prostor u memoriji za privremene promenljive koje svaka nit koristi. Posledica ovakvog pristupa je veće zauzeće memorije zbog privremenih promenljivih, međutim ukupno povećanje nije veliko, a svakako je opravdano značajno boljim performansama koje ova implementacija ima u odnosu na serijsku, što je pokazano

u poglavlju 8. Osim petlji u funkcijama, neophodno je paralelizovati i računanje vrednosti dipolnog člana GPJ, za koji se koristi FFTW biblioteka. Ova biblioteka već podržava OpenMP, čime je paralelizacija ovog koraka značajno olakšana.

Algoritam za sisteme sa deljenom memorijom se može prilagoditi i implementirati na grafičkim procesorima (GPU), što je tema poglavlja 4. Iako algoritam nije značajno izmenjen, zbog različitog programskog modela koji GPU koristi implementacija je značajno kompleksnija. Naša implemetacija za grafičke procesore je zasnovana na CUDA tehnologiji, koju koriste Nvidia grafički procesori. GPU arhitektura se fundamentalno razlikuje od tradicionalne arhitekture koje poseduju centralni procesori računara, i njeno razumevanje je neophodno da bi se iskoristio pun potencijal grafičkih procesora. Dok su centralni procesori dizajnirani da izvršavaju raznorodne zadatke korišćenjem relativno malog broja jezgara, grafički procesori su dizajnirani isključivo za *floating-point* operacije, za šta koriste velik broj GPU niti. Niti se izvršavaju po modelu *Single Instruction Multiple Thread*, gde sve izvršavaju istu instrukciju nad različitim podacima, uz značajno manju kontrolu nad međusobnom sinhronizacijom. Grafički procesori poseduju svoju RAM memoriju, odvojenu od glavne memorije računara, i svi podaci kojima GPU efikasno pristupa moraju biti dostupni u njegovoj memoriji. Iako novi grafički procesori imaju i do 24 GB memorije, to je ipak osetno manje od glavne memorije najmodernijih računara, koji često imaju i po 128 GB. Iz tog razloga, optimalno korišćenje GPU memorije je ključno za njihovu efikasnu primenu.

CUDA grafički procesori se programiraju pomoću ekstenzije programskog jezika C, nazvanog CUDA C. CUDA pravi razliku između koda koji se izvršava na glavnom procesoru i koji je pisan u neizmenjenom programskom jeziku C, i koda koji se izvršava na grafičkom procesoru za koji je neophodno korišćenje proširenog jezika C i specijalizovanog kompajlera. Centralno mesto u CUDA programskom modelu zauzima koncept *kernel*-a, odnosno funkcija koje programer definiše, a koje se izvršavaju na grafičkom procesoru. Kerneli se pozivaju korišćenjem specijalne sintakse, i prilikom poziva CUDA kreira željeni broj niti koje izvršavaju telo funkcije. Unutar kernela je moguće napraviti razliku između pojedinačnih niti korišćenjem posebnih varijabli preko kojih se može izračunati njihov broj ili mesto u hijerarhiji izvršavanja.

Grafički procesori se, iako međusobno imaju sličnu hardversku arhitekturu, značajno razlikuju po svojim računskim kapacitetima. Da bi se omogućio razvoj programa koji ne zavise od računskog kapaciteta grafičkog procesora, CUDA model izvršavanja se oslanja na logičku hijerarhiju na dva nivoa. Na prvom nivou se nalazi skup svih niti koje će se izvršiti, nazvan *grid*, koji je podeljen u niz blokova iste veličine, dok drugi nivo čine niti unutar blokova. Veličina grida i blokova se definiše pri pozivu kernela. Prilikom izvršenja programa, CUDA određuje koliko blokova će se paralelno izvršavati, u zavisnosti od kapaciteta grafičkog procesora. Bitna posledica ovakvog modela je da ne postoji garancija o broju blokova koji se izvršavaju ili njihovom redosledu, što znači da sinhronizacija između njih nije moguća. Sinhronizacija niti je moguća jedino unutar jednog bloka. Ovo u značajnoj meri utiče na dizajn algoritama za grafičke procesore, koji se često moraju značajnije menjati da bi mogli da se efikasno implementiraju.

Implementacija našeg algoritma na GPU je zahtevala značajne izmene u odnosu na OpenMP paralelizaciju. S obzirom na to da se implementacija algoritma u značajnoj meri oslanja na korišćenje petlji bez međuzavisnosti njihovih iteracija, implementacija nije zahtevala veće izmene algoritma i upotrebu sinhronizacije. Jedino je pomoćni algoritam za numeričku integraciju, baziran na Simpsonovom pravilu, morao da se izmeni u značajnijoj meri. S druge strane, implementacija

predstavlja veliki tehnički izazov jer zahteva korišćenje CUDA tipova podataka za reprezentaciju kompleksnih brojeva, specifičan način (de)alokacije memorije i pristup indeksima nizova, aritmetiku sa pokazivačima, izmenu biblioteke za Furijeovu transformaciju, itd. Takođe, posebna pažnja je bila posvećena optimizaciji upotrebe memorije, odnosno smanjenju ukupne memorije koju program zauzima na grafičkom procesoru, što je postignuto reorganizacijom računskih operacija i naprednom upotrebom pokazivača. Na ovaj način, u CUDA implementaciji smo uspeli da smanjimo upotrebu memorije i do 50% u odnosu na OpenMP implementaciju.

Moderni računarski sistemi često poseduju snažan višezgarni centralni procesor povezan sa grafičkom karticom sličnih računskih performansi. Korišćenjem programa iz poglavlja 3 ili 4 se može koristiti samo jedan od ovih resursa. Bolji pristup je razvoj hibridnog algoritma i implementacije koji će objediniti oba resursa u jedan heterogeni sistem, što se razmatra u poglavlju 5. U ovom poglavlju je opisan hibridni algoritam, od načina podele podataka između dva resursa do podele obrade podataka na svakom od njih, a takođe je predstavljena implementacija i optimizacija dobijenog rešenja.

Prvi korak u dizajnu hibridnog algoritma je razvoj efikasne šeme podele podataka i posla između dva resursa. Imajući u vidu da su memorije centralnog i grafičkog procesora odvojene i različitog su kapaciteta, nije opravdano šemu bazirati na modelu koji zahteva da sadržaj obe memorije bude istovetan. Umesto toga, potrebno je logički podeliti podatke na deo koji će obrađivati grafički procesor i deo koji će obrađivati centralni procesor. U slučaju da je jednom resursu potreban deo podataka na kojima radi drugi resurs, ove podatke je potrebno preraspodeliti kopiranjem u odgovarajuću memoriju. Pritom treba omogućiti da se simultano računanje na oba resursa nesmetano nastavi, kao i da ne dođe do problema sa sinhronizacijom. Upravo takvu šemu koristi naš hibridni algoritam i ona predstavlja važan doprinos ove teze. Korisnost šeme je takođe vidljiva kroz hibridni algoritam za Furijeovu transformaciju, koji je razvijen za potrebe programa, ali se može lako iskoristiti i za druge hibridne algoritme koji zahtevaju ovu transformaciju. Ovaj algoritam je primenljiv na višedimenzionu Furijeovu transformaciju i funkcioniše tako što deli transformaciju na dva dela: deo koji se može izračunati na svakom resursu sa podacima koji se nalaze u lokalnoj memoriji, i deo za koji je potrebna distribucija podataka po dimenziji koja nije lokalna.

Hibridna implementacija osnovnog algoritma je zasnovana na delovima koda OpenMP i CUDA implementacija. Glavne funkcije za računanje vremenske propagacije talasne funkcije se mogu iskoristiti uz male izmene granica petlji. Međutim, ovako naivna implementacija bi bila veoma neefikasna, jer bi oba resursa bila neiskorišćena tokom preraspodele podataka. Implementacija se može unaprediti tako što se računске operacije nad jednim delom podataka podele u manje celine i povežu u jedan tok (engl. *stream*) koji ima definisan redosled izvršavanja. Upotrebom više tokova moguće je u potpunosti preklopiti raspodelu podataka i računanje nad njima, i time značajno povećati efikasnost hibridne implementacije. U ovom slučaju nastaje problem da za maksimalno iskorišćenje oba resursa moramo precizno da odredimo količinu podataka koju će svaki procesor obraditi, kao i broj tokova koji će biti korišćen, što nije jednostavan zadatak jer umnogome zavisi od performansi oba resursa. Rešenje ovog problema je predstavljeno u poglavlju 8, u formi metoda za automatsku optimizaciju parametara podele podataka.

Algoritmi predstavljeni u poglavljima 3, 4 i 5 namenjeni su izvršavanju na jednom računaru, gde mogu da koriste sve dostupne resurse. Sledeći korak je njihovo proširenje na sisteme sa distribuiranom memorijom, odnosno na računarske klastere, što je tema poglavlja 6. Karakteristika

ovakvih sistema je model memorije u kome nijedan računarski resurs ne poseduje kompletan skup podataka. Pristup podacima koji pripadaju nekom drugom resursu zahteva njihov prenos kroz mrežni interfejs, što je značajno sporije od rada sa lokalnim podacima. Algoritmi za ovakve sisteme treba da poseduju šemu distribucije podataka koja maksimalno iskorišćava lokalnost podataka, a podacima udaljenih resursa pristupa samo kada je to neophodno. U okviru poglavlja 6 predstavljeno je proširenje glavnog algoritma na sisteme sa distribuiranom memorijom, kao i tri implementacije. Algoritam koristi model razmene poruka da pristupi podacima na drugim računarima u klasteru, dok se za implementaciju koristi *Message Passing Interface* (MPI). Sve tri implementacije koriste MPI samo za razmenu podataka, dok se za paralelno računanje nad lokalnim podacima koriste algoritmi koje smo prethodno razvili. Tri implementacije su:

1. Implementacija koja radi samo sa centralnim procesorima, zasnovana na algoritmu za deljenu memoriju, odnosno njegovoj OpenMP paralelizaciji iz poglavlja 3 (OpenMP/MPI),
2. Implementacija koja radi samo sa grafičkim procesorima, a zasnovana je na implementaciji iz poglavlja 4 (CUDA/MPI),
3. Hibridna implementacija zasnovana na algoritmu predstavljenom u poglavlju 5 (Hibrid/MPI).

Sve implementacije koriste istu šemu distribucije podataka i imaju sličnu strukturu. Prilikom razvoja, fokus je stavljen na programe koji rešavaju GPJ u tri dimenzije, jer je potreba za njihovom paralelizacijom najveća. Korišćenje ovih implementacija na sistemima sa distribuiranom memorijom pomaže nam da ostvarimo dva cilja, tj. omogućava nam da upotrebom više računarskih resursa skratimo vreme izvršavanja simulacija za datu rezoluciju, dok nam takođe omogućavaju i da izvršavamo simulacije sa značajno većom rezolucijom, za koje kapacitet memorije jednog računara nije dovoljan i koje na drugi način ne bi ni mogle da budu izvršene.

Predstavljeni algoritam za distribuirane sisteme se oslanja na podelu podataka po jednoj dimenziji, što znači da svaki resurs (ili proces u MPI terminologiji) ima u svojoj lokalnoj memoriji podatke u preostale dve dimenzije. Inicijalno su podaci podeljeni između procesa po  $x$  koordinati, što omogućava da svaki proces u svojoj lokalnoj memoriji propagira talasnu funkciju u odnosu na delove Hamiltonijana koji sadrže prostorne izvode po  $y$  i  $z$  koordinati. Da bi talasnu funkciju propagirali u odnosu na deo Hamiltonijana sa prostornim izvodima po  $x$  koordinati neophodno je transponovati podatke, tako da se obezbedi lokalnost podataka po celokupnoj  $x$  koordinati u svakom procesu. Operacija transponovanja podataka zauzima centralno mesto u ovim programima i njihove performanse su direktno povezane. U tezi smo ispitivali dva načina transponovanja podataka: korišćenjem FFTW biblioteke koja ovu operaciju pruža kao svoju dodatnu funkcionalnost, i putem rutine koju smo sami razvili za ovu potrebu. Uzimajući u obzir da OpenMP i hibridna implementacija već koriste FFTW prilikom računanja vrednosti dipolnog člana GPJ, upotreba ove biblioteke za transponovanje podataka ne unosi dodatne zavisnosti u program, odnosno jedini preduslov je da je FFTW biblioteka kompajlirana sa podrškom za MPI. Međutim, CUDA implementacija se ne oslanja na ovu biblioteku, već koristi cuFFT biblioteku koja radi isključivo sa grafičkim procesorima, i koja ne poseduje ekvivalentnu operaciju za transponovanje, niti bilo kakvu podršku za MPI. Da bi rešili ovaj problem, razvili smo novu rutinu za ovu operaciju, koja koristi MPI tipove podataka i komunikaciju između svih procesa da bi ostvarila neophodnu razmenu podataka. Uz upotrebu MPI implementacije koja pristupa memoriji grafičkog procesora direktno, bez potrebe za prebacivanjem

u glavnu memoriju, ova rutina ima odlične performanse. Takođe, pošto dva navedena pristupa operaciji transponovanja kao rezultat imaju identičan raspored podataka u memoriji, u Hybrid/MPI implementaciji su oba pristupa omogućena i korisnik prilikom kompajliranja programa može da odabere koju implementaciju želi.

Osim operacije transponovanja koja nam je neophodna da bi izvršili propagaciju talasne funkcije u odnosu na odgovarajući deo Hamiltonijana, potreban nam je i algoritam za računanje vrednosti dipolnog člana GPJ. Da bi se njegova vrednost izračunala, neophodno je koristiti algoritam za diskretnu Furijeovu transformaciju koji radi na sistemima sa distribuiranom memorijom. U slučaju OpenMP/MPI implementacije, ovo je jednostavan zadatak, jer FFTW biblioteka ima podršku za MPI i koristi identičan model podele podataka između procesa. U Hybrid/MPI programima bilo je neophodno proširiti algoritam za Furijeovu transformaciju iz poglavlja 5. Osnovni koncept podele po dimenzijama i korišćenja tokova podataka ostaje nepromenjen, s tim što je u ovom slučaju neophodno i dodatno transponovati podatke i obezbediti odgovarajući raspored podataka u memoriji. Kao i u slučaju hibridnog algoritma za Furijeovu transformaciju na jednom računaru, i ovaj algoritam predstavlja dodatni doprinos teze koji se može primeniti izvan okvira programa za koje je razvijen. U slučaju CUDA/MPI implementacije može se primeniti ista ideja podele podataka, dodatno olakšana činjenicom da su svi podaci jednog procesa u memoriji grafičkog procesora, te nije neophodno korišćenje tokova podataka i njihovo prebacivanje u glavnu memoriju računara.

Korišćenje MPI programa omogućava simulacije eksperimenata značajno veće rezolucije, te se javlja problem ispisa vrednosti koje program računa. Sa jedne strane, ukoliko bi MPI procesi unutar programa ispisivali svoje lokalne podatke u zasebne fajlove, korisniku bi ostao nezavidan zadatak da te fajlove objedini na korektan način i analizira. Sa druge strane, tekstualni zapis koji programi iz poglavlja 3, 4 i 5 koriste postaje neadekvatan za simulacije visoke rezolucije usled ogromne veličine fajlova. MPI programi koje smo razvili ovaj problem rešavaju korišćenjem MPI *input/output* operacija. Umesto u zasebne fajlove, svi procesori pišu u jedan, binarni fajl koji zadržava preciznost podataka iz programa, i pritom zauzima manje prostora na disku od tekstualnog zapisa.

Sa povećanim obimom izlaznih podataka, javio se i problem njihove analize i vizuelizacije, kako nakon završetka izvršavanja simulacije, tako i tokom samog izvršavanja. Ovim problemom se bavi poglavlje 7. Vizuelizacija podataka tokom izvršavanja simulacije, tzv. *in situ* vizuelizacija, je veoma korisna jer omogućava značajno bržu analizu rezultata, a ukoliko dodatno pruža mogućnost uticaja na tok simulacije (na primer, kroz modifikaciju ulaznih parametara dok se simulacija izvršava), takođe značajno smanjuje i vreme koje je korisniku potrebno da putem simulacija ispita neki fizički fenomen. Pored vizuelizacije ove vrste, veoma je korisno i pojednostaviti vizuelizaciju podataka nakon završetka simulacije korišćenjem standardnog formata. Oba navedena slučaja korišćenja smo pokrili kroz integraciju razvijenih programa sa *VisIt* sistemom za analizu i vizuelizaciju podataka.

*VisIt* je program otvorenog koda koji omogućava vizuelizaciju podataka iz velikih simulacija, čije napredne funkcionalnosti predstavljaju poslednju reč tehnologije u oblasti analize i vizuelizacije podataka. Oslanja se na klijent-server arhitekturu, gde se server izvršava paralelno na sistemu sa distribuiranom memorijom, te se uz pomoć njega može prikazati veliki skup podataka poput onog iz naših MPI programa. *VisIt* može da pristupi podacima na više načina od kojih smo mi koristili dva pristupa: putem fajlova zapisanih u VTK formatu koje *VisIt* može da pročita i putem *libsim* biblioteke koja omogućava *in situ* vizuelizaciju.

Proširenje programa da pored dosadašnjeg tekstualnog ili binarnog zapisa svoje podatke zapisuje

i u VTK formatu je olakšano postojanjem više biblioteka za ovu namenu, dok se za osnovne potrebe može iskoristiti postojeći slobodan kod koji ne unosi dodatne zavisnosti u program. Na ovaj način smo omogućili programima da ispisuju profile gustine u tri dimenzije i integrisane dvodimenzione profile gustine talasne funkcije koje se kao VTK fajlovi mogu lako vizuelizovati ne samo u VisIt-u već i u drugim programima koji podržavaju ovaj format.

Integracija sa VisIt-om putem `libsim` biblioteke omogućava da se podacima u memoriji pristupi kao da su resursi koje VisIt može da prikaže na ekranu, bez potrebe za njihovim upisom na disk. Pošto putem ove integracije komponenta VisIt sistema zadužena za vizuelizaciju radi zajedno sa procesom koji poseduje podatke, nema dodatne potrebe za razmenom podataka između procesa, te se ovaj pristup izuzetno dobro skalira prilikom izvršavanja na velikom broju računara u klasteru. Putem ove biblioteke smo omogućili vizuelizaciju stanja svih relevantnih promenljivih u programu. Osim vizuelizacije, `libsim` je iskorišćen i za kontrolu same simulacije. VisIt i `libsim` poseduju više mehanizama za kontrolu toka simulacije. Na osnovnom nivou, VisIt omogućava da se definišu funkcije u programu koje se pozivaju pritiskom na dugme u korisničkom interfejsu VisIt-a, što smo iskoristili za osnovne komande pokretanja, zaustavljanja, prekida i ponovnog pokretanja simulacije. Naprednija funkcionalnost je moguća kroz razvoj dodatnog modula koji VisIt pretvori u korisnički interfejs koji se prikazuje korisniku. Kroz ovaj mehanizam moguća je naprednija kontrola simulacije putem prosljeđivanja parametra, ali potpuna kontrola trenutno ipak nije moguća, i predmet je aktivnog istraživanja za proširenje u budućoj verziji VisIt-a. I pored nedostataka, kroz ovaj mehanizam smo omogućili izmenu svih bitnih parametara i obezbedili kontrolu toka simulacije do nivoa pojedinačne iteracije glavne petlje. Osim dva navedena pristupa, moguće je napraviti i konzolni interfejs, putem kojeg se programu iz VisIt-a šalju tekstualne naredbe. Ovaj mehanizam je ograničen samo sposobnošću programa da tumači tekstualne naredbe, i uz dovoljno napredan parser moguće je ostvariti potpunu kontrolu nad simulacijom. Ovaj interfejs smo iskoristili da obezbedimo istu funkcionalnost koja je dostupna i preko specijalizovanog grafičkog interfejsa. Razvijeni sistem kontrole simulacije, kao i vizuelizacija podataka nakon završetka simulacije bili su od ogromne koristi tokom istraživanja predstavljenog u poglavlju 9.

Performanse svih implementacija su izmerene u detaljnim testovima predstavljenim u poglavlju 8. Kao osnovna metrika je korišćeno vreme izvršavanja, odnosno mereno je vreme izvršavanja jedne iteracije glavne petlje, usrednjeno nad 1000 iteracija, za različite veličine diskretizacione šeme. U opštem slučaju, vreme trajanja jedne simulacije najviše zavisi od veličine diskretizacione šeme, koja kontroliše broj tačaka a samim tim i broj iteracija u petljama glavnih funkcija. Naravno, vreme izvršavanja simulacija zavisi linearno od ukupnog broja iteracija glavne petlje, koja kontroliše fizičko vreme propagacije i koja je specifična za eksperiment koji se simulira. Imajući ovo u vidu, predložena metrika nam omogućava da steknemo sliku o performansama programa i predvidimo njegovo vreme izvršavanja za buduće simulacije. Svi testovi su izvršeni na PARADOX-IV klasteru Laboratorije za primenu računara u nauci, u okviru Nacionalnog centra izuzetnih vrednosti za izučavanje kompleksnih sistema Instituta za fiziku u Beogradu. PARADOX klaster se sastoji od preko 100 računara, svaki sa dva Sandy Bridge Xeon procesora (ukupno 16 jezgara po računaru), 32 GB RAM memorije i Nvidia Tesla M2090 grafičkom karticom.

Testirano je skaliranje programa koji se izvršavaju na jednom računaru u odnosu na serijsku implementaciju, kao i skaliranje MPI implementacija u odnosu na paralelizacije na jednom računaru. Proučavano je i *jako* i *slabo* skaliranje, gde se jako skaliranje odnosi na slučaj kada je veličina

problema (diskretizacione šeme) konstantna, a povećava se broj paralelnih resursa, dok je kod slabog skaliranja količina računskih operacija po pojedinačnom resursu konstantna, a povećava se broj paralelnih resursa, što je praćeno i povećanjem diskretizacione šeme na odgovarajući način. Na osnovu dobijenih rezultata, za sve programe je izračunato ubrzanje, kao odnos vremena koje jedan računarski resurs provede u jednoj iteraciji glavne petlje u odnosu na vreme za istu operaciju u paralelnom programu, kao i efikasnost skaliranja, koja predstavlja ostvareno skaliranje u odnosu na idealno teorijsko linearno skaliranje. Osim predstavljenih rezultata, razvijeni su i modeli performansi većine programa.

Pre ocenjivanja samih performansi, proučavali smo optimizaciju hibridnih implementacija što je poseban problem koji se pojavljuje samo kod hibridnih algoritama. Za razliku od ostalih algoritama kod kojih se optimizacija završava na programskom nivou, tokom implementacije, performanse hibridnih implementacija dodatno, na značajan način, zavise od iskorišćenosti centralnog i grafičkog procesora, koju određuju parametri podele podataka. Izbor optimalnih vrednosti ovih parametara nije lak zadatak, jer zahteva duboko razumevanje performansi hardvera koji se koristi, i u opštem slučaju veoma je teško ručno izabrati optimalne parametre. To je zahtevalo razvoj metoda koji će biti u stanju da izabere optimalne parametre na osnovu pokretanja programa i merenja njegovog vremena izvršavanja, bez intervencije korisnika. Ispitivana su tri pristupa, od naivnog pretraživanja najbolje konfiguracije putem iscrpne pretrage, preko metoda zasnovanog na algoritmu opadajućeg gradijenta, do metaheurističke metode zasnovane na genetskom algoritmu.

U programima koji rešavaju GPJ u tri prostorne dimenzije postoji ukupno 33 parametara koji određuju podelu i obradu podataka na oba resursa, što kao posledicu ima veoma velik prostor pretraživanja. Problem se može olakšati ako parametre grupišemo u četiri grupe čiji se izbor prirodno nameće, na osnovu tačaka u algoritmu na kojim je neophodno da se oba resursa sinhronizuju. Svaki skup parametara je relevantan samo do tačke naredne sinhronizacije, te se može pretražiti zasebno.

I pored navedene podele problema u manje grupe, iscrpno pretraživanje nije odgovarajuće zbog velikog broja kombinacija koje treba proveriti. Ukoliko se umesto svake vrednosti svakog parametara isproba svaka  $n$ -ta vrednost, možemo značajno smanjiti vreme ovakve pretrage, ali uz rizik da odabrano rešenje ne bude zaista optimalno, odnosno da se u pretraživanju slučajno „preskoči” optimalna kombinacija parametara.

Drugi metod koji smo ispitivali, algoritam opadajućeg gradijenta, je zasnovan na računanju gradijenta funkcije koju minimizujemo, i koja u našem slučaju predstavlja vreme izvršavanja datog programa. Pošto ova funkcija ne može da se napiše u analitičkoj formi, njen gradijent možemo da izračunamo jedino pomoću numeričke aproksimacije. Ovo se pokazalo kao suviše problematično, jer je šum pri merenju vrednosti funkcije značajan, odnosno dva ili više pokretanja programa sa istim vrednostima parametara neće imati identična vremena izvršavanja, pri čemu razlike imaju veliki uticaj na izvršavanje algoritma pretrage. To znači da algoritam često može pretragu da usmeri u pogrešnom pravcu. Takođe, pošto ovakav algoritam u osnovi radi sa realnim brojevima, a vrednosti naših parametara mogu biti samo celobrojne, javlja se i problem izbora novih vrednosti parametara na osnovu rezultata jedne iteracije algoritma. Ove probleme smo pokušali da rešimo na više načina, međutim nijedan pokušaj nije uspeo da u potpunosti otkloni sve manjkavosti ovog pristupa, što nas je navelo na zaključak da algoritam opadajućeg gradijenta nije pogodan za problem koji smo pokušali da rešimo.

Treći pristup je baziran na genetskom algoritmu (GA) koji je dizajniran da oponaša prirodnu



selekciju, po uzoru na biološku evoluciju u prirodi. GA funkcioniše tako što kreira populaciju od individualnih probnih rešenja problema, koju potom oceni, modifikuje i kreira novu. Ovaj proces se ponavlja dok se ne dođe do željenog rešenja. Inicijalna populacija se obično formira slučajnim izborom, što omogućava ovom algoritmu da započne pretragu prostora iz više tačaka. Svako individualno rešenje u populaciji se ocenjuje uz pomoć funkcije koju zadaje korisnik. Na osnovu ocena, najbolje individue se izaberu za „reprodukciju”, i na osnovu njih se formira nova generacija, a proces se ponavlja. U našem slučaju, individue u populaciji predstavljaju jednu (validnu) kombinaciju parametara podele posla u hibridnom algoritmu, dok je funkcija kojom se ocenjuju vreme izvršavanja programa sa parametrima koje definišu individuu. Kvalitet GA zavisi od implementacije tri ključna operatora: selekcije, reprodukcije i mutacije. Svaki od operatora može biti implementiran na više načina, i za sva tri smo razvili i testirali više potencijalnih implementacija.

Selekcija diktira način na koji se biraju najbolje individue iz populacije. Ovaj izbor može biti proporcionalan njihovoj oceni, a može biti zasnovan i na nekoj drugoj metrici ili pak proizvoljan. Treba imati u vidu da je moguće da dođe do prerane konvergencije ka neoptimalnom rešenju ukoliko se biraju samo najbolje individue, odnosno potrebno je napraviti operator selekcije koji ne zanemaruje sva rešenja i daje evolutivnu mogućnost nekim od slabije ocenjenih individua. Ovo se najčešće radi tako što se relaksira selekcija u prvih nekoliko generacija, da bi neke individue sa slabijom ocenom nakon reprodukcije omogućile bolje pretraživanje prostora. Isprobali smo tri oblika selekcije i među njima utvrdili da oblik selekcije poznat pod nazivom *turnir* daje najbolje rezultate. Ovakva selekcija podrazumeva da se slučajno izabere određen broj jedinki iz populacije i iz ovog skupa odabere najbolja. Kontrolom veličine turnira se može smanjiti pritisak na izbor najboljih jedinki i sprečiti prerana konvergencija.

Reprodukcija je operator koji određuje način na koji će se parametri od dve izabrane individue kombinovati da se napravi nova jedinka, a da se pritom dobije validna kombinacija, odnosno skup parametara pomoću kojih se program može izvršiti. Mi smo ispitali više implementacija, međutim nijedna se nije pokazala kao značajno bolja od drugih.

Poslednji operator, mutacija, podrazumeva izmenu slučajno izabranog parametra u jedinki i ekvivalentan je slučajnom prolazu kroz prostor pretraživanja. Mutacija funkcioniše sa malom verovatnoćom, u našem slučaju do 5%. Uticaj mutacije u GA je diskutabilan i često osporavan, međutim u našem slučaju je bio koristan za sprečavanje konvergencije ka neoptimalnom rešenju u slučaju kada se koristi mala populacija.

Poređenjem tri metoda optimizacije na različitim veličinama diskretizacije šeme utvrdili smo da GA daje najbolje rezultate. Metod iscrpnog pretraživanja je uspevao da nađe skup parametara blizak optimalnom, ali predloženo rešenje obično nije bilo optimalno, usled prevelikog broja koraka koji bi ovaj metod morao da koristi da bi uspeo da ispita veći deo prostora pretrage. Metod zasnovan na algoritmu opadajućeg gradijenta je pokazao najgore performanse, iz razloga koje smo već naveli. GA je pronalazio odlična rešenja, veoma blizu najboljeg rešenja otkrivenog ručnom proverom, i jedini njegov nedostatak je to što osim za slučaj najmanje diskretizacije šeme, nije uspevao da konvergira ka egzaktno najboljem rešenju. Umesto toga, u poslednjoj generaciji GA nalazilo se nekoliko veoma sličnih rešenja koja su sva veoma blizu najbolje kombinacije parametara, što je dovoljno dobar rezultat za naše potrebe. Pošto se GA pokazao kao bolje rešenje od preostala dva pristupa, izabrali smo njega za optimizaciju hibridnih programa i pristupili testiranju.

Ubrzanje i jako skaliranje OpenMP programa je određeno poređenjem njihovog vremena izvrša-

vanja sa prethodno objavljenim serijskim programima. Pritom su testirane performanse sa različitim brojem OpenMP niti, od jedne do 16 niti. Imajući u vidu da je ova implementacija unapređena u odnosu na serijsku, najviše u delu gde se koristi Furijeova transformacija, već sa jednom niti se dobijaju bolja vremena izvršavanja, dok se povećanjem broja niti do 16 dobija dodatno ubrzanje od 11 do 12 puta, u zavisnosti od programa. Efikasnost skaliranja je veoma dobra, od 60-80%. Programi koji rade u jednoj prostornoj dimenziji ostvaruju manje ubrzanje, usled dela programa koji je ostao serijski. U ovom slučaju maksimalno ubrzanje je između 2 i 3, dostupno već sa 4 niti, i daljim povećanjem broja niti se ne dobija značajnije ubrzanje. Dobijeni rezultati su saglasni sa Amdalovim zakonom i mogu se modelirati na osnovu njega, što je detaljnije prikazano u poglavlju 8. Slabo skaliranje OpenMP programa je takođe testirano, pri čemu smo se fokusirali na programe koji rešavaju GPJ u tri prostorne dimenzije. Efikasnost ovog skaliranja je takođe veoma dobra, gde programi za propagaciju u realnom vremenu ostvaruju efikasnost od oko 75%, dok programi u imaginarnom vremenu ostvaruju nešto nižu efikasnost, oko 60%. Ova razlika je usled upotrebe kompleksnih brojeva u propagaciji u realnom vremenu, što zahteva više računskih operacija, dok se pri radu sa realnim brojevima puno iskorišćenje procesora ne ostvaruje usled ograničenog kapaciteta protoka memorije. Pored toga, u imaginarnom vremenu je obavezan i dodatan korak normalizacije talasne funkcije, što takođe negativno utiče na performanse ovih programa. Ovaj oblik skaliranja smo takođe modelirali na jednostavan način, funkcijom sa jednim parametrom, koji odgovara frakciji koda koji se izvršava paralelno.

CUDA programi su testirani malo drugačije, jer grafički procesor funkcioniše kao jedan, paralelan računarski resurs. Samim tim, nije moguće testiranje ove implementacije promenom broja paralelnih računskih resursa. Međutim, promenom veličine diskretizacione šeme možemo dobiti uvid u ponašanje grafičkog procesora i predvideti njegove performanse. U našim testovima, CUDA implementacija je pokazala ubrzanje od 10 do 14 za programe u dve prostorne dimenzije, u zavisnosti od tipa vremenske propagacije. Programi koji rešavaju GPJ u tri prostorne dimenzije su pokazali nešto manje ubrzanje, od 7 za programe u imaginarnom vremenu do 14 za programe u realnom vremenu. Najveća ubrzanja su ostvarena za diskretizacione šeme koje nisu ni suviše male, a ni previše velike. Rad sa malim diskretizacionim šemama ne zahteva dovoljno računskih operacija da bi se grafički procesor u potpunosti iskoristio, dok se u slučaju velikih šema, koje u potpunosti popunjavaju memoriju grafičkog procesora, takođe vidi pad performansi usled potrebe za prebacivanjem podataka iz glavne memorije računara.

Proučavanje performansi programa koji se izvršavaju na jednom računaru završeno je testom hibridne implementacije. S obzirom na to da su testovi OpenMP programa pokazali da je potrebno koristiti svih 16 niti u programima koji rešavaju GPJ u više od jedne prostorne dimenzije, u ovim testovima smo uvek koristili maksimalni broj niti (16), dok je optimalna raspodela podataka između centralnog i grafičkog procesora dobijena korišćenjem genetskog algoritma za zadatu veličinu diskretizacione šeme. Osim u slučaju malih diskretizacionih šema, ova implementacija pruža najbolje performanse od svih programa koji se izvršavaju na jednom računaru. Za male diskretizacione šeme, nakon podele posla, nijedan od resursa nema dovoljno podataka za obradu i ne može da bude u potpunosti iskorišćen, te se ne ostvaruje ubrzanje. Takođe, do pada performansi dolazi i za velike diskretizacione šeme, jer u ovom slučaju ograničenje dostupne memorije na grafičkom procesoru onemogućava dalje ubrzanje. Algoritam za optimizaciju parametara za velike diskretizacione šeme brzo konvergira ka modelu raspodele posla gde je memorija grafičkog procesora u potpunosti

iskorišćena, što je znak da nam za bolje performanse treba više memorije na GPU.

MPI programi su testirani na različitom broju računara u klasteru, od 2 do 32. Na svakom računaru je pokrenut jedan MPI proces, koji u slučaju OpenMP/MPI i Hybrid/MPI implementacije pokreće svojih 16 niti, dok u slučaju CUDA/MPI implementacije pristupa instaliranom grafičkom procesoru. Osnova za poređenje tri MPI implementacije je bila ekvivalentna implementacija za jedan računar, tj. OpenMP/MPI programi su poređeni sa OpenMP programima koji se izvršavaju sa 16 niti, CUDA/MPI programi su poređeni sa CUDA implementacijom koja koristi jednu grafičku karticu, dok su Hybrid/MPI programi upoređeni sa hibridnim programima namenjenim izvršavanju na jednom računaru. Performanse svakog MPI programa, pa tako i programa koje smo mi razvili, umnogome zavise od konkretnog klastera na kom se izvršavaju, najviše od brzine protoka podataka kroz mrežu. Pronalaženje najbolje konfiguracije klastera za konkretan problem zahteva dosta truda i može da stvori pogrešnu sliku o performansama programa. Rezultati koje smo predstavili su dobijeni bez specijalnog podešavanja klastera, sa ciljem da se pokažu osnovne performanse na osnovu kojih bi se, uz razvijene modele, moglo predvideti ponašanje programa na drugim klasterima. Zbog toga razvijena metodologija testiranja performansi programa predstavlja jedan od značajnih doprinosa ove teze.

Prilikom testiranja jake efikasnosti MPI programa za fiksnu veličinu problema, odnosno bez izmene veličine diskretizacione šeme, jasno je da će na dovoljno velikom broju računara performanse početi da opadaju, usled nedovoljne količine posla po jednom procesu i prevelike cene razmene podataka. Prema tome, iako povećanjem broja MPI procesa odnosno računara očekujemo da će vremena izvršavanja glavne petlje početi da se smanjuju, takođe očekujemo i da će za dovoljno veliki broju procesa doći do zasićenja, nakon čega će efikasnost početi da opada. Ovo znači da iako na malom broju procesa možemo da uočimo linearno skaliranje, modeli performansi ne mogu biti bazirani na linearnim funkcijama i moraju da uzmu u obzir i parametre vezane za brzinu mreže. Modeli koje smo razvili razmatraju ove detalje i pokazuju dobro slaganje sa eksperimentalnim rezultatima za performanse sve tri MPI implementacije.

U testovima jakog skaliranja, OpenMP/MPI programi su pokazali odlično ubrzanje, koje ide do 17 za programe u imaginarnom vremenu i do 22 za programe u realnom vremenu, koristeći 32 računara i ukupno 512 jezgara. Efikasnost ostaje uglavnom konstantna i kreće se u rasponu od 40% do 60%. Pritom, ubrzanje raste skoro linearno, što čini ovu implementaciju veoma pogodnom za simulacije sa veoma velikim diskretizacionim šemama. Slično ponašanje je uočeno i kod CUDA/MPI programa, gde je ubrzanje manje, od 9 do 10, uz manju efikasnost, od 30% do 40%. Iako je ubrzanje manje, treba imati u vidu da su apsolutna vremena izvršavanja CUDA/MPI programa manja od OpenMP/MPI programa, što ovu implementaciju čini veoma pogodnom za klustere koji sadrže grafičke procesore, naročito manje klustere. Na Hybrid/MPI programima je uočljiv pad efikasnosti sa povećanjem broja procesa, pošto je količina posla po samom procesu veoma mala, te dalja podela između centralnog i grafičkog procesora ne donosi bolje performanse. Međutim, ako posmatramo samo vremena izvršavanja, na manje od 16 procesa Hybrid/MPI implementacija ima najbolje performanse, što znači da je ova implementacija najpogodnija ako količina posla po procesu ostane dovoljno velika da opravda upotrebu hibridnog algoritma. U suprotnom, treba razmotriti bilo koju od druge dve implementacije.

U testovima slabog skaliranja korišćena je slična metodologija kao i u testovima OpenMP programa. Kao osnovu za poređenje koristili smo vreme izvršavanja jedne iteracije glavne petlje MPI

programa pokrenutog sa jednim procesom. Pokretanje MPI programa sa samo jednim procesom se ne preporučuje jer se operacija transponovanja podataka izvršava iako nema potrebe za njom. Ovo takođe može da utiče na rezultate testova, što je vidljivo u rezultatima OpenMP/MPI i CUDA/MPI implementacija. OpenMP/MPI programi, koji se oslanjaju na FFTW biblioteku za transponovanje podataka, pokazuju veoma dobre performanse kad se koristi samo jedan proces, jer FFTW prepoznaje ovaj slučaj i podatke transponuje na efikasniji način. Usled toga, prividno se dobija efikasnost od samo 40%. Međutim, ukoliko bi se kao osnova za poređenje koristilo vreme izvršavanja programa sa dva procesa, videlo bi se da efikasnost ostaje iznad 80% prilikom upotrebe 32 procesa, što je odličan rezultat. Sa druge strane, CUDA/MPI implementacija ostvaruje veoma loše performanse kad se koristi samo jedan proces, što je posledica upotrebe memorije centralnog procesora za operaciju transponovanja u ovom slučaju. Sa dva ili više procesa ovi programi ostvaruju mnogo bolje vreme izvršavanja, pa efikasnost može da bude čak i veća nego kod OpenMP/MPI programa, što nije očekivano jer je algoritam za transponovanje inferioran u odnosu na rutine koje pruža FFTW biblioteka. U ovom slučaju potrebno je uzeti apsolutno vreme izvršavanja na 16 procesa kao osnovu za poređenje da bi se dobila prava slika o efikasnosti skaliranja ove implementacije, koja je manja nego kod OpenMP/MPI implementacije. U testovima Hybrid/MPI implementacije se dobije efikasnost skaliranja od oko 75% i, usled upotrebe glavne memorije za transponovanje, ova implementacija nema probleme kao CUDA/MPI implementacija i pogodna je za izvršavanje na bilo kom broju procesa. Ako se poredi samo vremena izvršavanja, ova implementacija se nameće kao najbolje rešenje, jer količina posla po procesu ostaje dovoljno velika da opravda upotrebu hibridnog algoritma. Naravno, podrazumeva se da se pre izvršavanja simulacija optimizuju parametri podele podataka pomoću GA, i ovo optimizaciono vreme takođe treba uzeti u razmatranje prilikom izbora algoritma.

Da rezimiramo, u svim testovima MPI programa koji su obavljani, uočena je očekivana opadajuća efikasnost skaliranja, prouzrokovana potrebom za transponovanjem podataka u sistemu sa distribuiranom memorijom. Imajući sve rezultate u vidu, poglavlje 8 se završava savetima za izbor najboljeg algoritma za zadati problem.

Programi razvijeni u okviru ove teze se mogu koristiti za modeliranje i proučavanje raznih fizičkih sistema. U okviru poglavlja 9 demonstrirana je njihova upotrebljivost za proučavanje formiranja vorteksa u Boze-Ajnštajn kondenzatima. Korišćeni su MPI programi da bi se simulirali efekti prolaska prepreke kroz kondenzat. Prepreka, u formi odbojnog laserskog snopa, prolazi kroz kondenzat, pri čemu njeno kretanje dovodi do formiranja kvantnih vorteksa. Odvajanje vorteksa od prepreke se dešava jedino ukoliko se ona kreće brzinom većom od kritične, u skladu sa Landauovim kriterijumom za superfluidnost. Pošto se kretanje prepreke simulira kao promena potencijala u GPJ jednačini, programi su morali da budu izmenjeni tako da rade sa dinamičkim, vremenski zavisnim potencijalom da bi se ovo istraživanje omogućilo. Izmene programa su bile veoma jednostavne, što ide u prilog algoritmu podeljenog koraka koji nam omogućava da sve izmene budu lokalizovane u funkcijama koje su jednostavne za razumevanje i lake za paralelizaciju, i pri tom su odvojene od Krenk-Nikolson šeme.

Pomoću izmenjenih programa smo najpre simulirali rezultate ostvarene u skorašnjem eksperimentu [97], u kojem je merena kritična brzina na kojoj se formiraju rotirajući parovi vorteksa u Boze-Ajnštajn kondenzatu sačinjenom od atoma natrijuma. Pošto ovi atomi nemaju dipolni moment, programi su dodatno izmenjeni i uklonjeno je računanje vrednosti dipol-dipol interakcije.

Pokazana je saglasnost između rezultata dobijenih numeričkim simulacijama i u eksperimentu, što predstavlja jaku nezavisnu proveru ispravnosti naših algoritama i njihovih implementacija.

Osim navedene simulacije kondenzata sa atomima bez dipolnog momenta, takođe smo istraživali i formiranje vorteksa u kondenzatu od atoma disprozijuma, kod kojih se mora uzeti u obzir snažna dipol-dipol interakcija. Korišćena je ista metodologija za određivanje kritične brzine neophodne za formiranje vorteksa i proučen je uticaj jačine kontaktne i dipol-dipol interakcije na kritičnu brzinu. Zavisnost kritične brzine od jačine dipol-dipol interakcije do sada nije proučavana eksperimentalno, pa rezultati predstavljeni u poglavlju 9 predstavljaju originalni naučni doprinos koji će biti osnova za dodatna istraživanja u narednom periodu.

Budući pravci istraživanja uključuju razvoj paralelnih numeričkih algoritama za višekomponentne Boze-Ajnštajn-kondenzovane sisteme, za brzo rotirajuće sisteme, kao i za druge vrste interakcija između atoma. Glavni izazov u razvoju algoritama za višekomponentne sisteme predstavljaju njihovi povećani memorijski zahtevi, pa očekujemo da će algoritmi za računarske klustere sa distribuiranom memorijom biti osnova za dalje istraživanje u ovom pravcu. Sa druge strane, brzo rotirajući sistemi će zahtevati izmenu Krenk-Nikolson šeme, zbog prisustva dodatnih prostornih izvoda u Hamiltonijanu. Novi programi koji će omogućiti rešavanje GPJ uopštene na ovaj način, posebno uz implementaciju dodatnih tipova interakcije, biće od interesa za širu naučnu zajednicu koja u istraživanju koristi rešavanje ovakvih parcijalnih diferencijalnih jednačina.

# Short biography

Vladimir Lončar was born on 28 October 1985 in Novi Sad. He started his undergraduate studies in computer science at the Faculty of Sciences of the University of Novi Sad in 2004, and graduated in 2009. He continued his master studies at the same faculty and obtained a MSc degree in 2011, and immediately afterward he enrolled in PhD studies in computer science. Between 2012 and 2014 he actively participated in the development of the information system of the Faculty of Sciences in Novi Sad, where he was employed. Since 2015 he is employed at the Scientific Computing Laboratory, within the National Center of Excellence for the Study of Complex Systems of the Institute of Physics Belgrade, working on the national research project ON171017 "Modeling and numerical simulations of complex many-body systems". He coauthored two papers in peer-reviewed journals (category M21a), one conference paper (category M33) and two posters at international conferences (category M34). One of the journal papers is marked as a Highly Cited Paper in the Web of Science since its publishing in March 2016 and has a total of 21 citations.

# Kratka biografija

Vladimir Lončar je rođen 28. oktobra 1985. godine u Novom Sadu. Osnovne studije na Prirodno-matematičkom fakultetu Univerziteta u Novom Sadu, smer diplomirani informatičar - poslovna informatika, upisao je 2004. godine, a završio 2009. godine. Master studije na istom fakultetu, na smeru informacioni sistemi, završio je 2011. godine. Školske 2011/2012. godine je upisao doktorske studije informatike na istom fakultetu. Od 2012. do kraja 2014. godine je aktivno učestvovao u razvoju informacionog sistema Prirodno-matematičkog fakulteta u Novom Sadu, gde je bio i zaposlen. Od 2015. je zaposlen u Laboratoriji za primenu računara u nauci, u okviru Nacionalnog centra izuzetnih vrednosti za izučavanje kompleksnih sistema Instituta za fiziku u Beogradu, na projektu osnovnih istraživanja ON171017 „Modeliranje i numeričke simulacije složenih višestručnih sistema”. Objavio je dva rada kategorije M21a, jedno saopštenje kategorije M33 i dva saopštenja kategorije M34. Jedan od radova kategorije M21a je u Web of Science označen kao Highly Cited Paper za period od objavljivanja u martu 2016. godine do danas i ima ukupno 21 citat.

**University of Novi Sad**  
**Faculty of Sciences**  
**Key Words Documentation**

Accession number:  
NO

Identification number:  
INO

Document type: Monograph documentation  
DT

Type of record: Textual printed material  
TR

Contents code: Doctoral dissertation  
CC

Author: Vladimir Lončar  
AU

Advisor: Dr. Antun Balaž and Dr. Srđan Škrbić  
MN

Title: Hybrid parallel algorithms for solving nonlinear  
Schrödinger equation  
TI

Language of text: English  
LT

Language of abstract: Serbian/English  
LA

Country of publication: Serbia  
CP

Locality of publication: Vojvodina  
LP

Publication year: 2017  
PY

Publisher: Author's reprint  
PU

Publ. place: Novi Sad, Trg D. Obradovića 4  
PP

Physical description: 10/161/98/16/43/41/0  
(no. of chapters/pages/bib. refs/tables/figures/listings/appendices)  
PO

Scientific field: Informatics  
SF



Scientific discipline:

Computer Science

SD

Subject/Key words:

Numerical simulations, parallel algorithms, hybrid algorithms, nonlinear Schrödinger equation, GPU, OpenMP, MPI

SKW

UC

Holding data:

HD

Note:

N

Abstract:

Numerical methods and algorithms for solving of partial differential equations, especially parallel algorithms, are an important research topic, given the very broad applicability range in all areas of science. Rapid advances of computer technology open up new possibilities for development of faster algorithms and numerical simulations of higher resolution. This is achieved through parallelization at different levels that practically all current computers support.

In this thesis we develop parallel algorithms for solving one kind of partial differential equations known as nonlinear Schrödinger equation (NLSE) with a convolution integral kernel. Equations of this type arise in many fields of physics such as nonlinear optics, plasma physics and physics of ultracold atoms, as well as economics and quantitative finance. We focus on a special type of NLSE, the dipolar Gross-Pitaevskii equation (GPE), which characterizes the behavior of ultracold atoms in the state of Bose-Einstein condensation.

We present novel parallel algorithms for numerically solving GPE for a wide range of modern parallel computing platforms, from shared memory systems and dedicated hardware accelerators in the form of graphics processing units (GPUs), to heterogeneous computer clusters. For shared memory systems, we provide an algorithm and implementation targeting multi-core processors using OpenMP. We also extend the algorithm to GPUs using CUDA toolkit and combine the OpenMP and CUDA approaches into a hybrid, heterogeneous algorithm that is capable of utilizing all available resources on a single computer.

Given the inherent memory limitation a single computer has, we develop a distributed memory algorithm based on Message Passing Interface (MPI) and previous shared memory approaches. To maximize the performance of hybrid implementations, we optimize the parameters governing the distribution of data and workload using a genetic algorithm. Visualization of the increased volume of output data, enabled by the efficiency of newly developed algorithms, represents a challenge in itself. To address this, we integrate the implementations with the state-of-the-art visualization tool (VisIt), and use it to study two use-cases which demonstrate how the developed programs can be applied to simulate real-world systems.

AB

Accepted by Scientific Board on:

May 12, 2016

AS

Defended:

DE

Dissertation Defense Board:

(Degree/first and last name/title/faculty)

DB

President:

Dr. Dragan Mašulović, full professor,  
Faculty of Sciences,  
University of Novi Sad

Advisor:

Dr. Srđan Škrbić, associate professor,  
Faculty of Sciences,  
University of Novi Sad

Advisor:

Dr. Antun Balaž, research professor,  
Institute of Physics Belgrade

Member:

Dr. Nataša Krejić, full professor,  
Faculty of Sciences,  
University of Novi Sad

Member:

Dr. Miljko Satarić, academician, full professor  
Faculty of Technical Sciences,  
University of Novi Sad

**Univerzitet u Novom Sadu**  
**Prirodno-matematički fakultet**  
**Ključna dokumentacijska informacija**

Redni broj:  
RBR

Identifikacioni broj:  
IBR

Tip dokumentacije: Monografska dokumentacija  
TD

Tip zapisa: Tekstualni štampani materijal  
TZ

Vrsta rada: Doktorska disertacija  
VR

Autor: Vladimir Lončar  
AU

Mentor: dr Antun Balaž i dr Srđan Škrbić  
MN

Naslov rada: Hibridni paralelni algoritmi za rešavanje nelinearne Šredingerove jednačine  
NR

Jezik publikacije: Engleski  
JP

Jezik izvoda: Srpski/Engleski  
JI

Zemlja publikovanja: Srbija  
ZP

Uže geografsko područje: Vojvodina  
UGP

Godina: 2017  
GO

Izdavač: autorski reprint  
IZ

Mesto i adresa: Novi Sad, Trg D. Obradovića 4  
MA

Fizički opis rada: 10/161/98/16/43/41/0  
(broj poglavlja/strana/lit. citata/tabela/slika/listinga/priloga)  
FO

Naučna oblast: Informatika  
NO

Naučna disciplina: Računarske nauke  
ND  
Predmetna odrednica/Ključne reči: Numeričke simulacije, paralelni algoritmi, hibridni algoritmi, nelinearne parcijalne diferencijalne jednačine, GPU, OpenMP, MPI

PO  
UDK  
Čuva se:  
ČU

Važna napomena:  
VN

Izvod: Numerički metodi i algoritmi za rešavanje parcijalnih diferencijalnih jednačina, naročito paralelni algoritmi, predstavljaju izuzetno značajnu oblast istraživanja, uzimajući u obzir veoma široku primenljivost u svim oblastima nauke. Veliki napredak informacione tehnologije otvara nove mogućnosti za razvoj brzih algoritama i numeričkih simulacija visoke rezolucije. Ovo se ostvaruje kroz paralelizaciju na različitim nivoima koju poseduju praktično svi moderni računari.

U ovoj tezi razvijeni su paralelni algoritmi za rešavanje jedne vrste parcijalnih diferencijalnih jednačina poznate kao nelinearna Šredingerova jednačina sa integralnim konvolucionim kernelom. Jednačine ovog tipa se javljaju u raznim oblastima fizike poput nelinearne optike, fizike plazme i fizike ultrahladnih atoma, kao i u ekonomiji i kvantitativnim finansijama. Teza se bavi posebnim oblikom nelinearne Šredingerove jednačine, Gros-Pitaevski jednačinom sa dipol-dipol interakcionim članom, koja karakteriše ponašanje ultrahladnih atoma u stanju Boze-Ajnštajn kondenzacije.

U tezi su predstavljeni novi paralelni algoritmi za numeričko rešavanje Gros-Pitaevski jednačine za širok spektar modernih računarskih platformi, od sistema sa deljenom memorijom i specijalizovanih hardverskih akceleratora u obliku grafičkih procesora, do heterogenih računarskih klastera. Za sisteme sa deljenom memorijom, razvijen je algoritam i implementacija namenjena višejezgarnim centralnim procesorima korišćenjem OpenMP tehnologije. Ovaj algoritam je proširen tako da radi i u okruženju grafičkih procesora korišćenjem CUDA alata, a takođe je razvijen i predstavljen hibridni, heterogeni algoritam koji kombinuje OpenMP i CUDA pristupe i koji je u stanju da iskoristi sve raspoložive resurse jednog računara.

Imajući u vidu inherentna ograničenja raspoložive memorije koju pojedinačan računar poseduje, razvijen je i algoritam za sisteme sa distribuiranom memorijom zasnovan na Message Passing Interface tehnologiji i prethodnim algoritmima za sisteme sa deljenom memorijom. Da bi se maksimalizovale performanse razvijenih hibridnih implementacija, parametri koji određuju raspodelu podataka i računskog opterećenja su optimizovani korišćenjem genetskog algoritma. Poseban izazov je vizualizacija povećane količine izlaznih podataka, koji nastaju kao rezultat efikasnosti novorazvijenih algoritama. Ovo je u tezi rešeno kroz integraciju implementacija sa najsavremenijim alatom za vizualizaciju (VisIt), što je omogućilo proučavanje dva primera koji pokazuju kako razvijeni programi mogu da se iskoriste za simulacije realnih sistema.

IZ

Datum prihvatanja teme od strane

NN veća:

12. maj 2016.

DP

Datum odbrane:

DO

Članovi komisije:

(Naučni stepen/ime i prezime/zvanje/fakultet)

KO

Predsednik:

dr Dragan Mašulović, redovni profesor,  
Prirodno-matematički fakultet,  
Univerzitet u Novom Sadu

Mentor:

dr Srđan Škrbić, vanredni profesor,  
Prirodno-matematički fakultet,  
Univerzitet u Novom Sadu

Mentor:


dr Antun Balaž, naučni savetnik,  
Institut za fiziku u Beogradu

Član:

dr Nataša Krejić, redovni profesor,  
Prirodno-matematički fakultet,  
Univerzitet u Novom Sadu

Član:

dr Miljko Satarić, akademik, redovni profesor  
Fakultet tehničkih nauka,  
Univerzitet u Novom Sadu



PRIRODNO-MATEMATIČKI FAKULTET  
Univerzitet u Novom Sadu

FACULTY OF SCIENCES  
University of Novi Sad

TRG DOSITEJA OBRADOVIĆA 3, 21000 NOVI SAD, SRBIJA (SERBIA)  
tel +381.21.455.630 fax +381.21.455.662 e-mail dekanpmf@uns.ac.rs web www.pmf.uns.ac.rs  
PIB 101635863 MB 08104620

Broj: 0603-708/13  
Datum: 03.11.2017.

Na osnovu člana 29. st. 1 Zakona o opštem upravnom postupku "Sl. glasnik RS" br. 18/2016, u skladu sa članom 112. Zakona o visokom obrazovanju "Sl. glasnik RS" br. 88/2017, uvida u matične knjige studenata doktorskih studija Prirodno-matematičkog fakulteta Univerziteta u Novom Sadu i zahteva Lončar Milenko Vladimira, iz Novog Sada izdaje se

## U V E R E N J E

O STEČENOM VISOKOM OBRAZOVANJU TREĆEG STEPENA  
DOKTORSKIH STUDIJA

Lončar (Milenko) Vladimir

rođen 28.10.1985. godine u Novom Sadu, opština Novi Sad, država Republika Srbija, završio je visoko obrazovanje trećeg stepena-doktorskih studija, na studijskom programu Doktorske studije Informatika, Departmana za matematiku i informatiku Prirodno-matematičkog fakulteta Univerziteta u Novom Sadu, dana 17.10.2017. godine, sa prosečnom ocenom 9.92 (devet i 92/100), u toku studija i postignutim ukupnim brojem ESPB bodova 180.00 (slovima: sto osamdeset i 00/100) i stekao akademski naziv Doktor nauka-računarske nauke. Naslov doktorske disertacije je: "Hibridni paralelni algoritmi za rešavanje nelinearne Šredingenove jednačine"

Uverenje se izdaje radi lične upotrebe i zamenjuje diplomu do izdavanja iste.

Na osnovu člana 19. stav 1. tačka 7. Zakona o republičkim administrativnim taksama ("Sl. glasnik RS" broj 43/2003, 51/2003 - ispr., 61/2005, 101/2005 - dr. zakon, 5/2009 i 54/2009) ovo uverenje je oslobođeno takse.

Novi Sad, 03.11.2017.

Prodekan za nastavu  
Prirodno-matematičkog fakulteta

dr Tatjana Pivac, vanredni profesor