



ZIB


2018 Annual Report

**PHOTONIC
RESONANCES
BOOST
QUANTUM
TECHNOLOGY
PAGE 28**

**OPTIMIZATION
MEETS SUPER-
COMPUTING
PAGE 56**

**FAULT
TOLERANCE IN
DISTRIBUTED
SYSTEMS
PAGE 64**

**THE DEEP
PARALLELISM
CHALLENGE
PAGE 74**



ZIB

WWW.ZIB.DE

PREFACE

According to our web page, “ZIB is an interdisciplinary research institute for application-oriented mathematics and data-intensive high-performance computing. Its research focuses on modeling, simulation, and optimization (MSO) with scientific cooperation partners from academia and industry.” In our understanding, application-oriented mathematics does not mean just solving mathematical problems that appear in applications. Rather, ZIB pursues an integrated approach that combines efforts to describe the full application problem in mathematical terms and to develop methods for its mathematical, as well as practical, solution, including the required algorithmic and computational methods. This includes, in particular, the translation back into the application fields in concerted efforts with the application partners.

This way of doing research has made ZIB strong; an internationally renowned research institute with quite a unique profile. Since the definition of its new research strategy and development plan in 2015, the profile of ZIB has started to change, mainly in order to take up new challenges and utilize the related opportunities.

One of the main challenges is the exponentially growing amount of data acquired by digital technology, promising unprecedented insights into nature, economics, and society. Today’s data wave is already revolutionizing many areas of research and technology; some experts even predict that machine learning, especially deep learning, will replace mathematical modeling and analysis as the main driving force generating insights in scientific research. We believe that taming the growing complexity of technological, cultural, and social processes demands the seamless integration of both, advanced data analysis and MSO techniques. The challenge of raising application-oriented mathematics to a

new level thus requires novel forms of data-augmented MSO to enable more precise quantitative mathematical modeling of real-world processes, gaining deeper insights, improving predictions, and thus supporting decision-making.

In 2018, several crucial milestones of ZIB’s development plan were reached: the Berlin-based MATH+ Cluster of Excellence was successful in German’s most important funding line for research, the Excellence Strategy of the German federal and state governments. MATH+, a joint project of the three Berlin universities, FU, HU, and TU, with WIAS and ZIB, will receive about €7.5 million funding per year for at least seven years. In its core domain, application-oriented mathematics, ten interdisciplinary research projects will be funded at ZIB that all are designed to take up the challenges described above (for more details, please see the article “Berlin Mathematics is Shaping the Future”). In part, MATH+ aims at continuing the success story of the MATHEON research center, but with a stronger focus on data-based research.

This will be complemented by the newly established Berlin Center for Machine Learning (BCML), for which a Berlin-based consortium including ZIB could secure funding by the German Federal Ministry of Education and Research (BMBF) in another nation-wide competition (for more details please see the article “Berlin Center for Machine Learning Established”).

In addition to the successful acquirement of MATH+ and BCML, the CRC 1114 and TRR 154 DFG Collaborative Research Centers, as well as the Berlin Big Data Center (BBDC), all with essential contribution by ZIB, were successfully evaluated in 2018 and will continue to be funded for four further years. ZIB’s participation in these centers, especially in MATH+, and the research funded through it, contributes perfectly to the core research of our institute and to the ambitious strategic aims set out in 2015. In this sense, 2018 was a pivotal step in ZIB’s development.

This annual report provides insights into a variety of other success stories and

gives a general overview of ZIB’s organization and key factors for its successful development. In particular, six feature articles highlight aspects of our work: “How Mathematical Intelligence Can Help with Your Knee Pain” provides an insight into how advanced data analysis and modern MSO techniques can be combined to help address widespread disease. In “Photonic Resonances Boost Quantum Technology,” we report on recent research results at ZIB that provide insight into how quantum technology may be made possible based on nano-optical devices. “Optimization Meets Supercomputing” sheds some lights on how hard combinatorial structures can be tackled by parallel algorithms. The article “Digital Cultural Assets – digiS, or the Berlin Marvel” features the research agenda of the “Research and Competence Centre Digitalization Berlin” (digiS) that, after six years of successful activity, was made permanent and institutionally anchored at ZIB, in itself a small cultural-political miracle. “The Deep Parallelism

Challenge” reports on recent research of ZIB and its partners on why mastering the multiple levels of parallelism exposed by modern HPC systems is today the only way to achieve a sustained high application performance. And “Fault Tolerance in Distributed Systems” shows how researchers at ZIB study fault-tolerance approaches on different scales, ranging from the microlevel, byte-granular replication, to the macro-level, e.g., check-pointing in exascale computing.

In summary, ZIB continues to be a place for excellent research and first-rate scientific services and infrastructure. Last year was a fantastic year for ZIB with key success stories like the acquisition of the MATH+ Cluster of Excellence, which is set to have a positive impact on ZIB’s development for a decade and beyond, as well as lots of other very positive developments that make us confident that our institute has a bright future.

**Berlin, May 2019
Christof Schütte
President of ZIB**

28

PHOTONIC
RESONANCES
BOOST QUANTUM
TECHNOLOGY

36

HOW MATHEMATICAL
INTELLIGENCE
CAN HELP WITH
YOUR KNEE PAIN

4

ZUSE
INSTITUTE
BERLIN

Preface | Executive Summary |
Organization | ZIB Structure |
ZIB in Numbers | Berlin
Mathematics Is Shaping the
Future | Berlin Center for
Machine Learning Established |
Economic Situation | Spin-Offs |
Number of Employees

CONTENTS

46

DIGITAL CULTURAL
ASSETS – DIGIS OR
THE BERLIN MARVEL

74

THE DEEP
PARALLELISM
CHALLENGE

64

FAULT TOLERANCE
IN DISTRIBUTED
SYSTEMS

56

OPTIMIZATION MEETS
SUPERCOMPUTING

84

ZIB PUBLICATIONS
REFERENCES
IMPRINT

EXECUTIVE SUMMARY

MATHEMATICS FOR LIFE AND MATERIALS SCIENCES

The year 2018 was marked by the preparations for the MATH+ Cluster of Excellence, the continuation of the DFG Collaborative Research Center 1114, and the associated priority shifts regarding the division's research activity. In recent years, ZIB's Mathematics for Life and Materials Science division has concentrated its efforts on taking an internationally leading role in developing innovative application-oriented mathematics, moving beyond traditional approaches to modeling, simulation, and data analysis for key technologies. Until recently, most of the division's research had its roots in numerical mathematics and scientific computing, as well as scientific visualization.

Today, many of the great challenges originate from the growing complexity of technology and society, which

requires raising application-oriented mathematics to a new level. Nowadays, progress in the field requires an increased involvement of novel mathematical concepts, for example computational topology in shape analysis for medical imaging or geometry-based methods in the computational identification of nonlinear manifolds that govern the effective dynamics of complex systems. In addition to the need for new mathematics, addressing the challenges seems to require new forms of seamlessly merging classical model- and simulation-based approaches to machine learning and artificial intelligence. In 2018, the scientific divisions of ZIB invested considerable effort into new projects exploring these directions of research. Most of these new projects will now be supported by MATH+ and





CRC 1114, but also by the new Berlin Center for Machine Learning funded by the German Ministry of Education and Research, based on an application of a scientific consortium including ZIB. The feature articles “How mathematical intelligence can help with your knee pain” and “Photonic resonances boost quantum technology” present some of these activities.

The respective shift in our research orientation aims to establish new interactions between mathematics and new application fields, together with top-class Berlin experts and institutions, in areas such as social sciences and the humanities, revealing opportunities for cross-fertilization through joint interdisciplinary research endeavors and data-augmented approaches.

MATHEMATICAL OPTIMIZATION AND SCIENTIFIC INFORMATION

The integration of data-driven methods into optimization approaches and the increased use of supercomputing were the major research trends in the optimization department, with many activities in data analysis, hierarchical modeling, and parallelization. This trend resulted in several developments at ZIB, both strategically as well as research-wise. Strategically, the formation of a new academic partnership agreement with the Montreal-based Institute for Data Valorization (IVADO) aims at the joint development of cutting-edge expertise in computing and data science. Research-wise, notable progress in several real-life applications can be reported, for example the gas power plant monitoring tool KWPT going live, which was developed at ZIB within the Research Campus MODAL and is now integrated into the productive environment at OGE, the largest German

gas network operator. The developments were also made possible by the initiation of Plan4res, a Horizon 2020 project on the development of synergistic multi-energy models for European optimal energy system management, as well as by the publication of the new Handbook of Optimization in the Railway Industry.

The new 6.0 version of the SCIP open-source optimization suite for solving mixed-integer programming problems boosts performance by a factor of two. It includes the UG framework for massive parallelization of branch-and-bound codes on supercomputers with up to 80,000 cores, a new multi-criteria extension termed PolySCIP, and it comes along with MIPLIB version 6, the new academic and industrial standard test set for mixed-integer programming. SCIPJack, a UG-based parallel solver for Steiner tree problems, won first, sec-

ond, and third prizes at the 2018 PACE competition, complementing a GOR Thesis Award, an ATMOS Best Paper Award, and second place at the MOPTA Modeling Competition.

Open access to scientific research continues to be the main focus of the Scientific Information department at ZIB. The Research and Competence Center Digitalization Berlin, digiS, is now permanently funded in order to provide consulting, technical support, and coordination of projects to digitize cultural heritage. The DeepGreen project aims to transfer scientific publications, which can be made freely available at the end of their embargo period, into open access repositories. After successful completion of its first period, DeepGreen is now building a Germany-wide open access service.

PARALLEL AND DISTRIBUTED COMPUTING

The beginning of 2018 was marked by the technical preparations and contract negotiations for the new HLRN-IV supercomputer to be operated by ZIB. In March 2018, Berlin's Permanent Secretary of Higher Education and Research, Steffen Krach, ZIB's leadership and representatives of the HPC infrastructure vendor Bull/Atos signed the respective contract. With almost a quarter of a million processor cores, the main component of the supercomputer will be delivered in autumn 2019 and will then boost the application performance by a factor of five compared to the current system. Like its predecessors, HLRN-IV will be operated for the North German HLRN Consortium at two sites, Zuse Institute Berlin and our new partner in Lower Saxony, the University of Göttingen.

Regarding research, ZIB's Parallel and Distributed Computing division has again demonstrated its international competitiveness. Several new projects and scientific awards are proof of this. The first example is the Best Paper Award for DM-HEOM. In this paper it is shown how programmers can design their code in such a way that both portability and efficiency can be achieved at the same time. For this subject, the term "performance portability" was coined by the HPC community. It is currently one of the hottest research fields in HPC. See our feature article "The Deep Parallelism Challenge."

ZIB researchers won another Best Paper Award in the field of distributed data analysis at the fourth IEEE International Conference on High-Performance Big Data, Deep Learning, and Cloud Computing. In the respective article, we demonstrate how I/O performance can be improved by co-locating interrelated input data and near-optimal load balancing. The results have been used to accelerate the analysis pipeline of data received from earth observation satellites. Such geographically distributed systems must be able to survive many types of unexpected failures, a topic which is discussed in the feature article "Fault Tolerance in Distributed Systems."

ADMINISTRATIVE BODIES

The bodies of ZIB are the President and the Board of Directors (*Verwaltungsrat*).

President of ZIB

PROF. DR. CHRISTOF SCHÜTTE

Vice President

PROF. DR. MARTIN SKUTELLA

The Board of Directors was composed in 2018 as follows:

PROF. DR. PETER FRENSCH

Vice President, Humboldt-Universität zu Berlin (Chairman)

PROF. DR. CHRISTIAN THOMSEN

President, Technische Universität Berlin (Vice Chairman)

DR.-ING. ANDREA BÖR

Provost, Freie Universität Berlin

DR. JUTTA KOCH-UNTERSEHER

Der Regierende Bürgermeister von Berlin Senatskanzlei – Wissenschaft und Forschung

DR. JÜRGEN VARNHORN

Senatsverwaltung für Wirtschaft, Energie und Betriebe

PROF. DR. MANFRED HENNECKE

Bundesanstalt für Materialforschung und -prüfung (BAM)

THOMAS FREDERKING

Helmholtz-Zentrum Berlin für Materialien und Energie (HZB)

PROF. DR. HEIKE GRASSMANN

Max-Delbrück-Centrum für Molekulare Medizin (MDC)

The Board of Directors met on June 5, 2018, and December 12, 2018.

SCIENTIFIC ADVISORY BOARD

The Scientific Advisory Board advises ZIB on scientific and technical issues, supports ZIB's work, and facilitates ZIB's cooperation and partnership with universities, research institutions, and industry.

The Board of Directors appointed the following members to the Scientific Advisory Board:

PROF. DR. JÖRG-RÜDIGER SACK

Carleton University, Ottawa, Canada

PROF. DR. CECILIA CLEMENTI

Rice University, Houston, Texas, USA

PROF. DR. MICHAEL DELLNITZ

Universität Paderborn, Germany

PROF. DR. ROLF KRAUSE

Università della svizzera italiana, Lugano, Switzerland

PROF. DR. ALFRED K. LOUIS

Universität des Saarlandes, Saarbrücken, Germany

LUDGER D. SAX

Grid Optimization Europe GmbH

PROF. DR. REINHARD SCHNEIDER

Université du Luxembourg, Luxembourg

PROF. DR. DOROTHEA WAGNER

Karlsruher Institut für Technologie (KIT), Karlsruhe, Germany

The Scientific Advisory Board met on June 25 and 26, 2018, at ZIB.

THE STATUTES

The Statutes, adopted by the Board of Directors at its meeting on June 30, 2005, define the functions and procedures of ZIB's bodies, determine ZIB's research and development mission and its service tasks, and frame upon the composition of the Scientific Advisory Board and its role.

ORGANIZATION

SCIENTIFIC ADVISORY BOARD

CHAIRMAN PROF. DR. JÖRG-RÜDIGER SACK | *Ottawa, Canada*
PROF. DR. CECILIA CLEMENTI | *Houston, Texas, USA*
PROF. DR. MICHAEL DELLNITZ | *Paderborn, Germany*
PROF. DR. ROLF KRAUSE | *Lugano, Switzerland*
PROF. DR. ALFRED K. LOUIS | *Saarbrücken, Germany*
LUDGER D. SAX | *Essen, Germany*
PROF. DR. REINHARD SCHNEIDER | *Luxemburg, Luxemburg*
PROF. DR. DOROTHEA WAGNER | *Karlsruhe, Germany*

BOARD OF DIRECTORS

CHAIRMAN: PROF. DR. PETER FRENSCH
Humboldt-Universität zu Berlin (HUB)

PRESIDENT

PROF. DR. CHRISTOF SCHÜTTE

VICE PRESIDENT

PROF. DR. MARTIN SKUTELLA

MATHEMATICS FOR LIFE AND MATERIAL SCIENCES

Prof. Dr. Christof Schütte

MATHEMATICAL OPTIMIZATION AND SCIENTIFIC INFORMATION


Prof. Dr. Martin Skutella

PARALLEL AND DISTRIBUTED COMPUTING

Prof. Dr. Alexander Reinefeld

ADMINISTRATION

Annerose Steinke



ZIB STRUCTURE

ZIB is structured into four divisions: three scientific divisions and ZIB's administration.


Each of the scientific divisions is composed of two departments that are further subdivided into research groups (darker bluish color) and research service groups (lighter bluish color).

LEGEND

- SCIENTIFIC DIVISIONS AND DEPARTMENTS
- RESEARCH GROUPS
- RESEARCH SERVICE GROUPS
- CORE FACILITY

ZIB IN NUMBERS


SEMINARS
GIVEN BY ZIB SCIENTISTS AT UNIVERSITIES



LECTURES
GIVEN BY ZIB SCIENTISTS AT UNIVERSITIES



14,000



DOWNLOADS OF
MIP SOLVER SCIP

PROMOTION OF YOUNG SCIENTISTS:

DISSERTATIONS 

MASTERS 

3

PROFESSORSHIPS
OFFERED TO ZIB
RESEARCHERS

156




PEER-REVIEWED PUBLICATIONS IN
INTERNATIONAL SCIENTIFIC JOURNALS



53 DISTINGUISHED

152 INVITED

82 PB DATA ARCHIVE AT ZIB
TOTAL CAPACITY ON
17,500 TAPES



47 OUTREACH EVENTS FOR
SCHOOL CLASSES AND
THE GENERAL PUBLIC

€8,132,000

€6,425,000

PROJECT-RELATED PUBLIC
THIRD-PARTY FUNDS



€1,707,000

INDUSTRIAL THIRD-PARTY PROJECTS



AWARDS


5,210

INTERNATIONAL GUESTS
AT ZIB IN 2018



72

CONFERENCES AND
WORKSHOPS AT ZIB



**IN JUNE 2018:
ZIB SUPERCOMPUTER IS
NO. 273 IN TOP 500 LIST**

BERLIN MATHEMATICS IS SHAPING THE FUTURE

The Berlin-based Cluster of Excellence MATH+ was successful in the Excellence Strategy research competition run by the German federal and state governments. Starting in January 2019, MATH+ will be funded for seven years with up to 7.5 million euros per annum. After the first seven years it may be possible to extend the funding for a further seven years. The decision of the German Research Foundation to fund MATH+ confirms Berlin's outstanding position as an internationally excellent location for research in mathematics and its applications.

MATH+ is a cross-institutional and transdisciplinary Cluster of Excellence where researchers explore and further develop new approaches in application-oriented mathematics. Emphasis is placed on merging mathematical modeling with mathematical techniques for analyzing the ever-growing amounts of data in life and material sciences, in energy and network research, and in the humanities and social sciences. The aim is to boost not only scientific progress, but also technological innovation and the comprehensive understanding of social processes.

MATH+ is a joint project of the three major universities in Berlin – Freie Universität Berlin, Humboldt-Universität zu Berlin, and Technische Universität Berlin – the Weierstrass Institute for Applied Analysis and Stochastics, and ZIB. It continues the success stories of the renowned


MATHEON Research Center and of the Berlin Mathematical School, which has been supported by the Excellence Initiative since 2006.



**Transforming the
World through
Mathematics!**

Berlin Mathematics Research Center

MATH+



ZIB's research strengths in application-oriented mathematics including its trans-disciplinary context and the translation to industry and society as a whole is one of the cornerstones of MATH+. At ZIB, MATH+ will start ten new research projects funded through the Cluster of Excellence:


- “The spatio-temporal modelling of mechanisms underlying pain relief via the μ -opioid receptor” (Schütte, Weber) in cooperation with FU Berlin, Berlin’s university hospital Charité (Stein), and the Max Delbrück Center for Molecular Medicine in the Helmholtz Association (Lohse),
- “Space-time stochastic models for neurotransmission processes” (Schütte, Winkelmann) in cooperation with the NeuroCure Cluster of Excellence (Sigrist),
- “Discrete-continuous shortest path problems in flight planning” – a collaboration between ZIB’s Numerical Mathematics (Weiser) and Mathematical Optimization (Borndörfer) departments with Lufthansa Systems,
- “Simulation and optimization of integrated solar fuel and photovoltaics devices” (Burger) together with the Helmholtz Zentrum Berlin (Rech),
- “Spline models for shape trajectory analysis” (Hege, v. Tycowicz) in cooperation with Charité (Hennemuth) and the German Archeological Institute (Fless),
- “Model-based 4D reconstruction of subcellular structures” (Prohaska, Weiser) in cooperation with Freie Universität Berlin (v. Kleist) and the NeuroCure Cluster of Excellence (Hiesinger),
- “Learning reduced descriptions for large-scale agent-based mobility models” (Schütte) in cooperation with Freie Universität Berlin (Klus) and the Global Climate Forum (Jäger, Wolf),
- “Data-driven modeling of the romanization process of Northern Africa” (Djurdjvac-Conrad, Schütte) together with the German Archeological Institute (Ducke, Fless),
- “Morphological scoring of disease states – learning from large anatomical databases” (Zachow), a transfer unit together with Charité (Duda), and
- “Demand modelling and control for e-commerce using RKHS transfer operator approaches” (Sullivan), a transfer unit in cooperation with Freie Universität Berlin (Klus) and Zalando AG.

BERLIN CENTER FOR MACHINE LEARNING ESTABLISHED

BERLIN-BASED
CONSORTIUM,
INCLUDING ZIB,
SUCCESSFUL
IN NATIONWIDE
COMPETITION
FOR FUNDING

In recent years, machine learning (ML) and artificial intelligence have attracted a lot of attention, have been in the public eye due to several spectacular success stories, and facilitated a wide range of new products and services on global markets. Slowly, this disruptive development is also penetrating the German economy as well. In the sciences, too, ML is

gaining more and more popularity and is being used increasingly in interdisciplinary research projects. Due to its great scientific and technological importance, there are currently a number of publicly funded major projects and centers around the world that will provide the basis for new ML procedures for future developments in science and industry.




In addition, companies such as Facebook, Amazon, Google, Microsoft, Tesla, IBM, Uber, etc., invest billions in ML basic research, especially in deep learning, in order to gain sustainable innovation leadership in this new sector. ZIB has increased research activities regarding method development in and application of ML for years. At present, more than 40% of the projects at ZIB are (partly) involved with research in ML. However, the breadth and depth of the development of ML-based research seems so immense that it goes well beyond our capacity to even keep track of the progress in science and technology. Therefore, ZIB became part of a larger Berlin-based consortium led by TU Berlin, with the aim of establishing a larger ML research center that would allow bundling and fostering activities. As one of the first results of this initiative, the Berlin Center for Machine Learning (BCML) was established in September 2018. It was the result of a national competition in which the German Federal Ministry of Education and Research (BMBF) awarded funding for five competence centers all over Germany. The BCML intends to combine the synergy effects of Berlin's extraordinarily rich landscape of science in the ML field with the following goals:

- Advancing theoretical and algorithmic foundations of ML.
- Developing novel scientific and technical ML applications.
- Striving to make ML methods for science and industry significantly simpler and more usable.
- Jointly exploring new interdisciplinary research fields.

The BCML will also implement structures to create open platforms for knowledge and technology exchange in ML, both for industry and academia. Our methodological approach is based on the interaction between ML and applications, in particular in biomedicine, communication, and digital humanities.

The vision of the BCML is to further develop ML methods along the entire ML process through a close intermeshing of basic and applied research, while demonstrating the potential of using ML as the scientific process leading to new insights. The center achieves this through a synergetic, coordinated collaboration of leading ML experts from TU Berlin, FU Berlin, HU Berlin, and the University of Potsdam, as well as the Fraunhofer Heinrich Hertz Institute, the Max Delbrück Center for Molecular Medicine, the Max-Planck Institute for Molecular Genetics, the Max-Planck Institute for Human Development, the Charité, the German Heart Center, the German Research Center for Artificial Intelligence, and, last but not least, ZIB. The BCML will cooperate closely with the Berlin Big Data Center (BBDC) in which ZIB also participates.

In 2018, the total income of ZIB comprised 24.6 million euros. The main part of this was made available by the Federal State of Berlin as the basic financial stock of ZIB (10.9 million euros) including investments and Berlin's part of the budget of HLRN at ZIB. The second largest part of the budget resulted from third-party funds (8.1 million euros) acquired by ZIB from public funding agencies (mainly DFG and BMBF) and via industrial research projects. This was complemented by a variety of further grants, such as the HLRN budget made available by other German states or the research service budget of KOBV, summing up to almost 5.6 million euros in total.



ECONOMIC SITUATION IN 2018


ZIB INCOME

- 44% Core budget by State of Berlin
- 33% Third-party funds
- 23% Further grants

ECONOMIC SITUATION IN 2018


The Zuse Institute Berlin (ZIB) finances its scientific work via three main sources: the basic financial stock of the Federal State of Berlin and third-party funds from public sponsors and those of industrial cooperation contracts.

In 2018, ZIB raised third-party funding through a large number of projects. Project-related public third-party funds rose from 6.147 million euros in 2017 to 6.425 million euros in 2018, while industrial third-party projects declined from 1.748 million euros to 1.707 million euros. In total, 8.132 million euros in third-party funding marked a new record in ZIB's history – an increase for the seventh year in a row.




ZIB THIRD-PARTY FUNDS BY SOURCE

- 3% EU
- 12% DFG
- 24% Other public funds
- 21% Industry
- 40% BMBF incl. FC MODAL



ZIB THIRD- PARTY FUNDS IN EUROS

- INDUSTRY
- PUBLIC FUNDS



COMPUTING IN
TECHNOLOGY GMBH
(CIT)

1992 | www.cit-wulkow.de

Mathematical modeling and development of numerical software for technical chemistry

RISK-CONSULTING
PROF. DR. WEYER GMBH

1994 | www.risk-consulting.de

Database marketing for insurance companies

INTRANETZ GMBH

1996 | www.intranetz.de

Software development for logistics, database publishing, and e-government

AKTUARDATA GMBH

1998 | www.aktuardata.de

Development and distribution of risk-evaluation systems in health insurance

VISAGE IMAGING GMBH

(Originating from the ZIB spin-off
Visual Concepts GmbH)

1999 | www.visageimaging.com

Advanced visualization solutions for diagnostic imaging

ATESIO GMBH

2000 | www.atesio.de

Development of software and consulting for planning, configuration, and optimization of telecommunication networks

BIT-SIDE GMBH

2000

Telecommunication applications and visualization

DRES. LÖBEL, BORNDÖRFER &
WEIDER GBR / LBW OPTIMIZATION
GMBH

2000 | www.lbw-optimization.com

Optimization and consulting in public transport

LBW Optimization GmbH was founded in 2017 and is a spin-off of LBW GbR

LENNE 3D GMBH

2005 | www.lenne3d.com

3-D landscape visualization, software development, and services

JCMWAVE GMBH

2005 | www.jcmwave.com

Simulation software for optical components

ONSCALE SOLUTIONS GMBH

2006 | www.onscale.de

Software development, consulting, and services for parallel and distributed storage and computing systems

LAUBWERK GMBH

2009 | www.laubwerk.com

Construction of digital plant models

1000SHAPES GMBH

2010 | www.1000shapes.com

Statistical shape analysis

TASK – Berthold Gleixner Heinz Koch
GbR

2010

Distribution, services, and consulting for ZIB's optimization suite

QUOBYTE INC.

2013 | www.quobyte.com

Quobyte develops carrier-grade storage software that runs on off-the-shelf hardware

KEYLIGHT GMBH

2015 | www.keylight.de

Keylight develops scalable real-time Web services and intuitive apps. The focus is on proximity, marketing, iBeacon, and Eddystone for interactive business models

DOLOPHARM BIOSCIENCES UG

2017

A specialty pharmaceutical company focused on the clinical and commercial development of new products in pain management that meet the needs of acute and chronic care practitioners and their patients

SPIN- OFFS

NUMBER OF EMPLOYEES

In 2018, 221 people were employed at ZIB; of these, 152 positions were financed by third-party funds.

1/1/2018			1/1/2019			
3	0	3	3	1	4	MANAGEMENT
15	100	115	22	96	118	SCIENTISTS
34	10	44	35	6	41	SERVICE PERSONNEL
8	9	17	9	7	16	KOBV HEADQUARTERS
0	52	52	0	42	42	STUDENTS
60	171	231	69	152	221	TOTAL
Permanent	Temporary	TOTAL	Permanent	Temporary	TOTAL	



Sven Burger | burger@zib.de | +49-30-84185-302

PHOTONIC RESONANCES BOOST

QUANTUM TECHNOLOGY



Modeling Light-Matter Interaction in Quantum Technology Devices

Prototypes for future technologies such as quantum computing and quantum cryptography are realized in semiconductor physics laboratories. Advanced modeling and simulation methods are essential for interpreting physical results and for creating an effective design of the experimental setups.

SINGLE PHOTONS ON DEMAND

Quantum networks for communication offer the possibility to transmit information without any chance for unwanted eavesdropping. This offers a great advantage for information societies, and is a contrast to classical communication channels. In quantum communications, single photons at a well-defined quantum state are used to transmit information. Single photon sources with high levels of brightness are essential to making quantum networks a reality.


WHY IS IT SUCH A CHALLENGE TO CREATE A SINGLE PHOTON?

For quantum technology applications, single photons need to have a very well defined wavelength and direction of propagation, and they need to be fired by an emitter (the single photon source) exactly at the right time. A technological solution for defining the wavelength is the use of so-called quantum dots as sources. These are tiny semiconductor crystals made from just a few thousand atoms. With sophisticated experimental techniques it is possible to engineer the surroundings of the quantum dots on a nanometer scale. The light sources become resonant with a well-designed environment. This effect is known as

the Purcell effect. When the resonance condition is reached, the quantum dots emit single photons just in the right way: in a well-specified direction, and with the right energy and timing.

Together with physicists from Technische Universität Berlin (TU Berlin), ZIB's Numerical Mathematics department is optimizing such environments through numerical simulations. A single photon source with a complex nanostructure is shown in Figure 1. The layout for this device has been optimized through simulations at ZIB and it has been experimentally realized at TU Berlin [1].

- 1 Artistic view of the probability density distribution of a photon in a single photon source based on a quantum dot emitter and various structured and unstructured semiconductor layers. The inset represents the circular symmetric device geometry (blue layers: Bragg-interference back reflector, green: semiconductor material hosting the quantum dot, red: adjacent fiber for transporting the emitted photons to distant experiments).





ON-CHIP MANIPULATION OF QUANTUM STATES

Quantum computers perform computations by using quantum-mechanical phenomena such as superposition and entanglement. A quantum computer is totally different from classical computers that rely on binary digital electronics based on transistors and capacitors. In quantum computations, quantum bits are used which can be in superpositions of states. Large-scale quantum computers are expected to solve problems like integer factorization much more efficiently than classical computers. However, the development of quantum computers is still in a start-up period.

Multimode quantum optical circuits are essential for creating photonic quantum computers and other quantum devices. In particular, interfacing single photon sources, gates, and detectors on a single chip enables large quantum networks to be set up. The on-chip integration avoids losses between different technology platforms such that stability and performance can be significantly increased.

Figure 2 shows an image (scanning electron microscope) of a single photon source integrated into a waveguide and interfaced to a manipulation section which splits the single-photon quantum states into two outgoing waveguides. In quantum mechanics, the simultaneous presence of a single photon in two waveguides is possible through the quantum superposition principle. This device has been realized experimentally by our project partners at TU Berlin [2]. The design of the optical properties has been created as part of a collaboration between TU Berlin and ZIB [2]. This includes the design of the quantum dot placement within the waveguides, as well as the design of the quantum splitter section and the interfacing to external detectors. Single-photon emission as well as the combined emitter-gate interface functionality have been demonstrated. As the approach is scalable, it paves the way for multimode, fully integrated quantum photonic devices.





RESONANCES ENLIGHTEN COMPLEX SCATTERERS


Resonance phenomena are omnipresent in physics. Storage and transfer of energy between different resonant states allow one to explore wave effects in atomic, molecular, and optical physics as well as in other fields of science. Optical resonators are scaled down to the wavelength scale and beyond by using modern nanotechnology, as demonstrated in experiments of our collaborators. Placing point-like sources in the vicinity of such nanoresonators or antennas enables exploration of new regimes of light-matter interaction. In addition to quantum technology applications, further examples are plasmon lasers and modification of chemical reaction rates.


Theoretical models of light-matter interaction are needed to understand and optimize the performance of related

photonic devices. Maxwell's equations can be solved directly to obtain solutions for the electromagnetic field. In order to understand the interaction of emitters with nanoresonators, it is essential to precisely describe the coupling of the emitter to specific resonance modes. In collaboration with the Humboldt-Universität zu Berlin and JCMwave GmbH, the Computational Nano Optics group at ZIB has developed a theory for modeling dispersive light-matter interaction based on Riesz projections [3]. Figure 3 shows schematically how the coupling of a light source to specific resonances is described: for each resonance mode, the coupling is equal to exactly one contour integration around the complex resonance frequency. In addition, the coupling to the nonresonant


background can be quantified exactly using our theory.

We apply the method to compute modal decay rates of emitters embedded into optical nanoantennas. Figure 4 shows exemplarily how complex the light field emitted by a localized source in a dielectric cone antenna looks like. With the Riesz-projection-based theory, the field can be analyzed and all modal components and the nonresonant background can be quantified in simulations. Future applications include exploiting the Purcell effect also for perovskite quantum dots with enhanced two-photon-pumped photoluminescence, as recently demonstrated in a collaboration between Helmholtz-Zentrum Berlin, Lund University, and ZIB [4].






- 2 Electron microscope image of an on-chip beam splitter realized at TU Berlin [2]. A quantum dot (in the green region) integrated in a waveguide (red) emits a single photon. The photon is split in the (numerically optimized) manipulation section and leaves the device in a quantum superposition through two separate waveguides (right-hand side).
- 3 Schematic description of the numerical approach to describing interaction of an emitter with a resonant nanostructure. Red crosses indicate resonance frequencies of the nanostructure. A source emits at a real-value frequency level. For computing the coupling to specific resonances, contour integrations around these are performed (Riesz projections [3], dashed blue lines). For computing the coupling to a nonresonant background and to more distant resonances, a further contour integration can be performed (solid blue line).
- 4 Simulation of the light field in a dielectric nanoantenna. Light is emitted from an embedded center (indicated by an arrow). The complex field distribution can be quantitatively analyzed and interpreted using the developed Riesz projection method.



In collaboration with Humboldt-Universität zu Berlin, we have investigated model systems of strong coupling [5]. We have especially investigated plasmonic setups involving metal nanoparticles which can confine light fields on scales much smaller than the wavelength of light. The Riesz projection method developed for analyzing the involved light-matter interaction has allowed for understanding and designing experimental setups, such that the strong-coupling regime can be reached.

SINGLE MOLECULES COUPLED TO PLASMONIC RESONATORS

When emitters are placed in an optical cavity on resonance, the environment changes their coupling to light. In the weak-coupling regime, as it is for quantum dots in typical single photon sources, light extraction is enhanced. But more profound effects arise in the strong-coupling regime. Here, mixed light-matter states can form. It has been proposed to use these as key components in quantum information systems as well for as for ultraefficient switches and lasers. Strongly confined nano-optical resonances make it possible to reach the strong-coupling regime in experimental configurations. Figure 5 shows a sketch of a typical arrangement where a plasmonic nanoresonator (gold sphere) traps light in a thin gap above a substrate. When a single molecule is placed in this gap, its radiation properties change, the emitted spectrum is modified, and eventually it can be possible to form a mixed light-matter state, which is a sign of strong coupling.



- 5 *Strong coupling: a single molecule is placed in the vicinity of a metal sphere. When the molecule emits light at the wavelength of a plasmonic resonance of the sphere, the environment changes the properties of the emitted photons.*
- 6 *Geometry and finite element mesh for a plasmonic nanoantenna used in a recent software benchmark.*


RESONANCE IN A BOX

The field of computational photonics is highly competitive. ZIB has been active in this field for more than two decades. We actively participate in research cooperations for benchmarking our numerical methods and software implementations with competing approaches and realizations. Such benchmarks are of the highest importance for qualification and acceptability of new algorithms and modeling approaches, and for assessing the stability and efficiency of an algorithm. Therefore, the sustainable development of algorithms also depends on benchmarking activities.

In a recent benchmarking process, we compared the Riesz-projection-based method for the computation of photonic resonances with standard methods for approaching this simulation task [6]. Participants of the benchmark are groups from the University of Bordeaux, Université Paris-Saclay, Delft University

of Technology, Nankai University, University of Stuttgart, Institut Fresnel in Marseille, and ZIB. Figure 6 shows the geometry underlying one of the benchmark problems: a silver cube with rounded corners placed on a thin polymer spacer above a gold substrate. A strongly localized photonic resonance is present in the polymer spacer material. Experimentally, such setups have been used to demonstrate significant improvements in light emission from optical sources placed in the spacer layer.

In the benchmark, there was a general agreement on the results from most of the methods. The ZIB algorithms were shown to belong to the best performing methods regarding accuracy and convergence. With such qualified numerical methods, ZIB can continue contributing to technological enhancement, in this case especially in the field of optical quantum technologies.



Collaborations

Prof. Stephan Reitzenstein,
Technische Universität Berlin

Prof. Oliver Benson,
Humboldt-Universität zu Berlin

Dr. Martin Hammerschmidt,
Dr. Philipp Schneider, Dr. Lin Zschiedrich,
JCMwave GmbH

Prof. Christiane Becker,
Helmholtz-Zentrum Berlin für
Materialien und Energie

Prof. Philippe Lalanne,
University of Bordeaux



Stefan Zachow | zachow@zib.de | +49-30-84185-275

HOW MATHEMATICAL INTELLIGENCE CAN HELP WITH YOUR KNEE PAIN

Knee osteoarthritis (OA) is a chronic, degenerative joint disease affecting an increasing number of the human population, in particular in the older generation.






Due to the rising average life expectancy and the desire for a prolonged active lifestyle, it is becoming more and more important to understand the mechanisms of OA and to develop individual therapeutic concepts for each patient.

MORE IS BETTER: IMPROVING DIAGNOSTICS WITH MULTI-OMICS

Researchers at ZIB are working on methods to make the lives of potential and actual OA patients better: better diagnostics based on multi-omics technologies, more precise progression assessment and therapy planning using artificial-intelligence technology, and – if nothing else works – pain management with less side-effects and optimized prostheses.

Prediction of an individual's medical status is necessary for clinical decision-making, such as diagnosing osteoarthritis. Often, disease prediction is based on traditional clinical variables such as blood values, general patient condition, and disease history, where the available data of a patient is compared against population statistics. However, in many cases, these clinical parameters are not enough to accurately diagnose a patient's disease since similar data can be the result of completely different diseases on a molecular level. Therefore, not only these "traditional" clinical values should be used for diagnosis but should be complemented by new data types, e.g. from the molecular level. Thanks to today's so-called omics technologies, we can obtain a patient's molecular profile in great detail on multiple levels, e.g. using billions of DNA bases or analyzing the concentration of hundreds of thousands of proteins and metabolites.




Researchers at ZIB aim to build new integrative models for disease diagnostics based on artificial intelligence. These methods allow the combination of the established traditional data with the unprecedented wealth of data coming from different omics measurements, e.g. genomics, proteomics, or metabolomics. Unfortunately, this has been found to be very challenging due to the high complexity and volume of omics data and – more importantly – due to the comparatively small number of training samples that is available for developing these models. This problem is often addressed as the curse of dimensionality and can cause a lack of statistical significance and over-fitting of the learned models.

Fortunately, there seems to be a way out of this dilemma: recent work shows that a small but smartly selected part of the combined data is more informative compared to analyzing the full individual large datasets. This means that the problem has to be reformulated as a preselection problem of an as-small-as-possible set of important features based on the full available data-sets. This set should capture the footprint of the disease phenotype, e.g. the stage of a developing OA.

The new approach developed by ZIB researchers proposes a new strategy on how to identify this feature set. The key insight is that most biological species such as genes, proteins, or metabolites are connected in a particular way. Examples are activation or inhibition of each other or being part of a signal-transduction process, where a signal is transmitted through a series of molecular events. The knowledge about who influences whom is available through established biological networks. The main idea of the new approach is to first combine all available networks across omics layers (e.g. genomics networks with proteomics networks). Then, artificial-intelligence algorithms are used to identify a relevant part of the integrated network that can be used to define the search space for the desired set of important features.

In recent studies, we could show that using the resulting phenotype-oriented network-based feature-selection framework, we could obtain robust feature sets that enable new diagnostic methods. We could also show that the diagnostic performance improves when multiple omics data types are used.

IT'S ALL IN THE IMAGES ...





- Medical imaging has become the method of choice for diagnosing a wide variety of diseases. In the case of osteoarthritis (OA), X-ray images are usually acquired for an initial assessment of a disease state. Magnetic resonance imaging (MRI) is used for a more differentiated diagnosis involving all soft tissue structures such as meniscus, cartilage, and ligaments. If knee surgery is indicated, a computed tomography (CT) scan is often made for the planning of an artificial joint replacement, since CT images depict bony structures in very high detail. For the assessment of medical images, computer algorithms are increasingly applied to assist in analyzing and visualizing the different types of medical image data and to help medical doctors in their diagnosis and therapy planning.
- At ZIB, we are working on algorithms for medical image computing and 3-D anatomy reconstruction. In the case of osteoarthritis, we are developing algorithms for automatic classification of the OA state from medical image data based on machine learning [1,2]. Anatomical structures are automatically segmented and three-dimensionally reconstructed with the help of statistical shape knowledge and convolutional neural networks [3,4,5]. Our current algorithm for 3-D reconstruction of knee structures from MR images is ranked first in the international SKI10* contest on segmentation of knee images. The results produced by our algorithms are comparable to those by human experts.

· * www.ski10.org/results.php

... FOR DIAGNOSIS AND THERAPY PLANNING

3




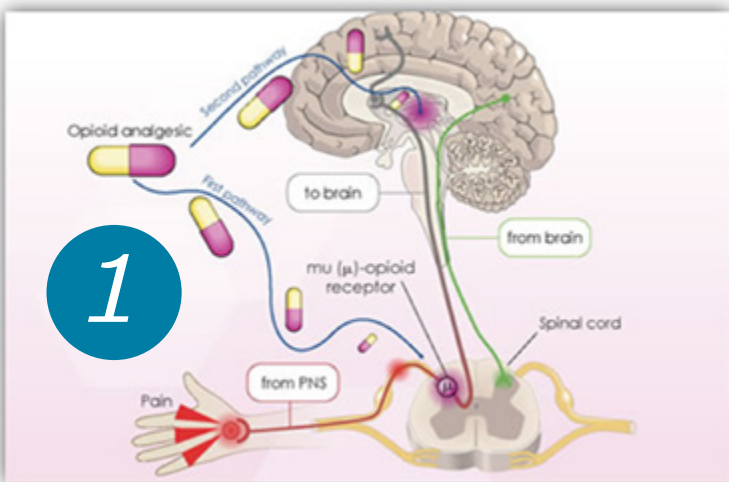
- 1 MR images of a knee.
- 2 3-D segmentation of knee structures.
- 3 Patient-specific 3-D planning of knee spacer therapy.

Researchers at ZIB team up with clinical practitioners and domain experts to develop innovative methods for diagnosis, therapy planning, and computer-assisted intervention with the aim to support medical doctors in their decision-making processes.

Most of our activities are embedded in larger research initiatives, for example in a BMBF-funded research network for musculoskeletal diseases. Here, we are investigating new OA scores that might improve the way high-risk patients are identified and diagnosed very early on – even before they encounter the first signs of arthritis themselves. That way, early individualized actions or a change in lifestyle can be advised to the patients (e.g. to increase activity or to reduce weight) to prolong a pain-free life.

In other projects, we develop methods that can avoid (or at least delay) joint replacements by optimally planning individualized biomimetic spacer therapies that can slow down the progression of osteoarthritis [6]. However, if a joint replacement becomes unavoidable, our methods can help surgeons by computing optimal replacement guidelines to position cutting guides and prostheses in an anatomically plausible manner [7,8].

OPIOIDS: MANAGING THE




OVERCOMING THE OPIOID EPIDEMICS

Painful osteoarthritis is often connected to inflammation. Also many knee surgeries end up in a treatment of a severe inflammatory pain. Strong pain killers (stronger than morphine) are often used with potentially lethal side effects. ZIB has developed a method that allows for simulating the inflammatory environment and, thus, enabling for a rational drug design.

The treatment of chronic pain is still a largely unexplored area. However, knee pain is often associated with acute situations (knee surgery) or inflammatory effects and can therefore be treated with acute pain therapy. Often the pain is so great that strong analgesics are administered (Fig. 1). One opioid used in clinical practice is fentanyl. It is a synthetic painkiller that was developed in the sixties and is about 100-times more effective than morphine. Although the current laws only allow the approval of such drugs, which avoid the most dangerous side effects, unfortunately fentanyl has the same fatal side effects as morphine. And there are many deaths annually due to the use of such opioids. Fentanyl plays

a role in up to 52 percent of drug-related opioid deaths [1]. In a joint publication in *Science*, researchers of ZIB and of Berlin's university hospital Charité proposed a new approach to avoiding these deaths [2]. The idea is based on changing the molecular structure of fentanyl in a way such that the new drug is active only in inflamed tissue while it stays inactive in "normal", healthy tissue, thus avoiding most of the undesired side effects. The pH value of inflamed tissue is significantly smaller than that of healthy tissue. The key idea of the proposed approach was to change the molecular structure of fentanyl such that the new drug would bind to the μ -opioid receptor (i.e. be "active") only for small pH values.

PAIN



3

NEW COMPOUNDS FOR PAIN TREATMENT

This necessary change in the molecular structure is done by replacing one hydrogen atom of fentanyl by a fluorine atom. Based on molecular dynamics computations, ZIB proposed several possibilities to place fluorine atoms in fentanyl such that the ability of fentanyl to bind an additional proton is just weakened in the correct intensity (which is measured as its pKa value). Together with Charité we published the corresponding testings of the new compounds [3]. It turned out that these new compounds are really able to reduce side effects significantly in animal testings.

MATHEMATICAL MODELLING

The pKa value prediction is only one part of the computations that has been done at ZIB. The main part corresponds to detailed simulations that are needed in order to estimate the ability of the new compounds to bind to the μ -opioid receptor (Fig. 3) for different pH regimes. This has been done based on molecular dynamics (MD) simulations of the μ -opioid receptor [4] as well as on careful investigations of proton transfer mechanisms [5]. The new mathematical techniques developed for these simulations resulted from ZIB's continued effort in conformation dynamics. In addition, these techniques allowed application to other drug-development strategies as well [6], and finally resulted in a completely new mathematical framework for performing classical MD simulations more efficiently [7], based on joint work of applied mathematicians and theoretical chemists in the framework of the Collaborative Research Center CRC 1114.

- 1 Mode of action of opioids in pain therapy.
- 2 Idea of fentanyl derivatives.
- 3 Simulation of the μ -opioid receptor in its cell membrane.

STEP BY STEP

In severe cases of osteoarthritis, the affected joints must be replaced by implants to allow pain-free motions. One of the factors limiting the life time of endoprostheses is inflammation due to implant wear. Therefore, implant designs must pass a wear test for market admittance.

1

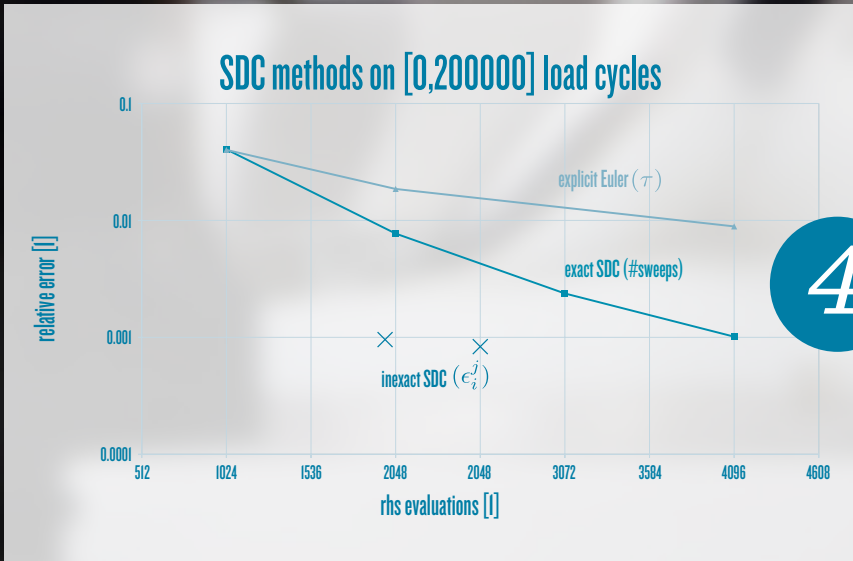



HARDWARE WEAR TESTS

For wear testing, implants undergo one to five million load cycles in testing machines (Fig. 1), which take one to three months at a cost of around 30,000 €. This imposes a significant burden on the implant design process, that would benefit from fast and frequent feedback on the quality of the design. Numerical simulation of wear testing promises to be a faster and cheaper alternative, reducing the number of required hardware tests dramatically.

3

$$\dot{\overline{\partial\Omega}}(t) \approx \frac{1}{\tau} \int_{s=t}^{t+\tau} w(u(\overline{\partial\Omega})(t), s-t) ds$$



4

MULTISCALE SIMULATION

VIRTUAL WEAR TESTS

- 1 Knee implants undergoing wear testing.
- 2 A finite element grid of a knee implant's two main parts: the femoral component (top) usually consisting of cobalt-chrome alloy and the tibial component made of polyethylene.
- 3 Averaging over load cycles.
- 4 Higher efficiency of inexact spectral deferred correction methods exploiting coarser load cycle simulations during early iterations, compared to fixed-accuracy Runge-Kutta schemes.

Modeling the process of wear involves the solution of dynamic contact problems on a geometry that changes due to wear. Using a sufficiently fine finite element grid (Fig. 2), the displacements and contact forces due to prescribed loads can be computed by solving a constrained nonlinear optimization problem [1]. For reasonable accuracy, a load cycle has to be divided into more than 100 time steps. Neglecting inertia, which is small compared to the applied loads, leads to a quasi-static approximation and allows to solve stationary contact problems in parallel. Finally, the loss of material and the resulting shape change can be computed from contact pressure, implant motion, and Archard's wear law (Fig. 4).

Tracking each load cycle and the resulting shape change is prohibitively expensive. But as the shape change over a single cycle is negligible, one can exploit the separation of time scales between wear and kinematics by integrating an averaged shape evolution (Fig. 3), where the right hand side is computed by simulating a single load cycle. Integrating the averaged evolution with an explicit Euler method is known as "cycle jumping". Higher order time steppers such as Runge-Kutta methods can increase the efficiency significantly. Inexact spectral deferred correction methods developed at ZIB [2] can additionally exploit the accuracy-effort trade-off available in evaluating the right hand sides by simulating the load cycle on coarser space and time grids in early iterations, leading to higher efficiency.

DIGITAL CULTURAL ASSETS – DIGIS OR THE BERLIN MARVEL



SIDENTZ STADT BERLIN



1 Stiftung Stadtmuseum Berlin. Plan von der Konigl. Residentzstadt Berlin. 1723.

1

A “cultural memory,” as asserted by Aleida and Jan Assmann, is the “result of incessant cultural work.” This cultural memory needs “dialogue and lively confrontation with the contemporary world.”[3]

THE BERLIN MODEL

Today's society, however, is largely digital. Digital change is now present in all areas of social life and also challenges culture and cultural heritage institutions. The digitization of cultural assets offers enormous opportunities in this respect. If digital artifacts are no longer bound by space and time, they are in principle available to all. A form of participation and creative confrontation – appropriation and reinterpretation – becomes possible in dimensions that were unthinkable in the past. Interested citizens, artists, protagonists from the (creative) economy, scientists, and researchers are beginning to rediscover cultural heritage digitally, place cultural data in its own context, and demand extended (re)use possibilities.[5]

The cultural heritage institutions as traditional guardians of our memory are confronted with new challenges: new competencies, experiments, and new perspectives are required. As Frankfurt curator Franziska Mucha states: "We can't just copy and paste analog museum ideas into the digital realm one-to-one – we need moderators, adjustments, and reiterative project development." [1] Ideally, concrete projects would be the best way to test, evaluate, improve, and implement ideas, and, for this to be sustainable, a long-term technical, legal, and organizational infrastructure would be a wonderful complement. And this is exactly where the marvel of Berlin enters the picture.

Since 2012, Berlin cultural heritage institutions have been called upon to make their collections available in a digital format and are therefore financially supported by the Senate Administration for Culture and Europe through a competitive funding program for the digitization of cultural heritage objects. At the same time, a central coordination and service center (digiS), linked to the Zuse Institute Berlin (ZIB), was conceived and implemented to accompany the funding program. Today digiS operates under the name of Forschungs- und Kompetenzzentrum Digitalisierung Berlin (Research and Competence Center Digitalization Berlin). In 2018, after five successful years as a project, the funding program was made permanent and digiS was institutionally anchored at ZIB, in itself a small cultural-political miracle.

The promotion from interim project to a permanent institution ensures continuity in the objectives of the funding program and convincingly demonstrates the importance that Berlin's state policy attaches to the topic of "digitizing cultural heritage." digiS continues to discharge its task of networking, advise the funded institutions, and provide technical expertise, in particular for the long-term availability of the data generated under the auspices of the funding program. With its technical infrastructure ZIB, as a nonuniversity research institute for applied mathematics and computer science, provides an excellent basis for this undertaking.



“digiS supports institutions in making Berlin’s cultural heritage digitally and permanently available, publicly accessible, and reusable for all. This includes the development of services as well as corresponding research activities.”^[4]



- 2 *Allierten Museum Berlin. Open day at the Tempelhof air base.*
- 3 *Stiftung Berliner Mauer / Foto: Paul Kremer. Viewing platform and visitors at Potsdamer Platz at the S-Bahn entrance.*

THE DIGIS PORTFOLIO

digiS accompanies the digitization projects through the entire workflow starting with legal advice, then indexing and the actual digitization, and finally the preparation and enrichment of the data, its presentation in the German Digital Library (DDB) and on the institutions' own websites, and the archiving and/or ensuring the permanent availability of the project data.

Thus from the very start of a digitization project, issues regarding data modeling have to be clarified in order to be able to translate and apply the concepts, standards, and standardized vocabularies specific to a field to the data collection in question. The digiS project partners benefit from an extensive workshop program [2], the offer of

an individual one-on-one consultation, and continuous support from digiS over the entire funding period and often beyond. Only by standardizing objects and describing and enriching them with publicly accessible vocabularies can the data, as a raw material, be refined and made available for use by third parties. After all, development issues must also be taken into account with regard to the long-term availability of the data right at the start of the digitization process. What is to be preserved for eternity in the first place? How should digital objects be structured and described so that both their formats remain readable and their context understandable for future users in 10, 30, 50, or 100 years? There are no simple, universal answers to these

questions. However, preserving culture must also apply to digital cultural data and therefore be rethought conceptually (Fig 1, 2).

In addition to further training, digiS provides support in terms of tangible technical services for the structuring and enrichment of data and the creation of data mappings for long-term archiving and presentation [6]. digiS also broadens the ranks of aggregators by supplying the Deutsche Digitale Bibliothek through the informal role of “dark aggregator,” i.e. an aggregator without official status. In close cooperation with the DDB's various specialist departments [7], digiS assists its project partners in the transfer of project data to the national cultural heritage portal.



4



5

DATA FOR ETERNITY (IN GERMAN: EWIGKEIT)

In close cooperation with the KOBV (Cooperative Library Network Berlin-Brandenburg) at ZIB, digiS has developed services to ensure the long-term availability of cultural data. This interdisciplinary cooperation at ZIB is expressed in the Digital Preservation joint working group. The open-source-based long-term availability system EWIG [8] is just as suitable for archiving digital artifacts from the field of cultural heritage as it is for archiving the products of scientific research. Core components of the architecture, which is based on the OAIS reference model, are the Archivematica archiving system, the iRODS storage management system, and the Fedora

repository. Due to a framework agreement concluded between the State of Berlin and ZIB in the course of the continuation of the funding program, all project data will be archived by digiS for ten years – without generating further costs for the institutions providing the data. For the cultural heritage institutions under the direct responsibility of the State of Berlin, the framework agreement guarantees the permanent archiving of the digitized objects. In doing so the State of Berlin once again assumes a pioneering role in the area of securing the long-term availability of heterogeneous data sets (Fig 3, 4).

- 4 *Archiv des Lette-Vereins Stiftung des öffentlichen Rechts. Der Lette Verein. 1904.*
- 5 *Stiftung Stadtmuseum Berlin. Heinrich Zille: Die Untermieterin. um 1905.*
- 6 *Stiftung Stadtmuseum Berlin. Heinrich Zille: Nauener Platz. 1911.*



6

WHAT IS CULTURAL HERITAGE AND WHO OWNS IT?

The term “cultural heritage” is a relatively young concept in the German language. After two catastrophic world wars, it first appeared in the 1954 Hague Convention for the Protection of Cultural Property in the Event of Armed Conflict.[9] It defines cultural property as “movable or immovable property of great importance to the cultural heritage of every people [...]” These include movable cultural assets such as paintings, sculptures, archaeological finds, books, manuscripts, and archival documents, as well as buildings and monuments. In this context, UNESCO speaks of human heritage and at the same time addresses the question of cultural heritage ownership. The aim of protecting cultural property is to “preserve the cultural heritage in order to be able to pass it on to future generations unscathed. Since cultural property is always a testimony to human development, its protection always benefits the general public.”[10]

Cultural heritage, understood as a collection of man-made and purposefully designed artifacts, which are in addition to natural artifacts, spans a broad spectrum in the context of digiS, ranging from dinosaur bones and dance films to Zille drawings (Fig. 5, 6). The diversity of the institutions with which digiS cooperates is equally impressive. The project partners include major Berlin state institutions – the Zentral- und Landesbibliothek Berlin,

the Landesarchiv Berlin, the Stiftung Stadtmuseum Berlin, Berlinische Galerie, the Deutsches Technikmuseum Berlin – but also institutions that are simply located in the Berlin boroughs (such as the FHXB Museum or the archives and collections of the Tempelhof-Schöneberg museums). Archives of movements, such as the Archive of Youth Cultures, are just as much a part of this, as are archives and collections from the independent scene – such as the Mime Centrum Berlin. Collections from research museums such as the Museum für Naturkunde and increasingly also from universities and other tertiary academic institutions are digitized via the funding program, the latter also because they cannot be directly integrated into a specific research project, but nevertheless play an outstanding role in their respective discipline.[11]

Writings and images of revolutions (1848/1918) have been digitized (Fig. 7), Berlin (hi)stories – from a sociohistorical, architectural, and cultural perspective (Fig. 8) – have been written, and works and networks of Berlin artists – from Hannah Höch to Egmont Schäfer – have been prepared for digital consumption in the present. Berlin’s theater and dance culture from the beginning and up until the end of the 20th century has been given a digital stage, and natural history artifacts – animal calls and dinosaur bones from the bone depository of





7

7 *Stiftung Stadtmuseum Berlin. Car with machine guns of the Workers and Soldiers Council at the Brandenburg Gate. 1918.*

the Museum für Naturkunde – have been made available in digital form.

In some cases, digital cultural data directly serves as a basis for further scientific research. The project of the Institute of Theater Studies at Freie Universität Berlin is a good example of this collaboration between culture and science. In a pioneering work, the complete director's book by Max Reinhardt for "Dantons Tod" (1916) was digitized and this valuable source material in terms of the history of theater has been published as a digital edition. "The digitization and editorial project offers the opportunity to illustrate Max Reinhardt's development of a new theatrical aesthetic of an excellently suited source with digital methods and thus to make it accessible for research and teaching in a new way." [12] (Fig. 9)

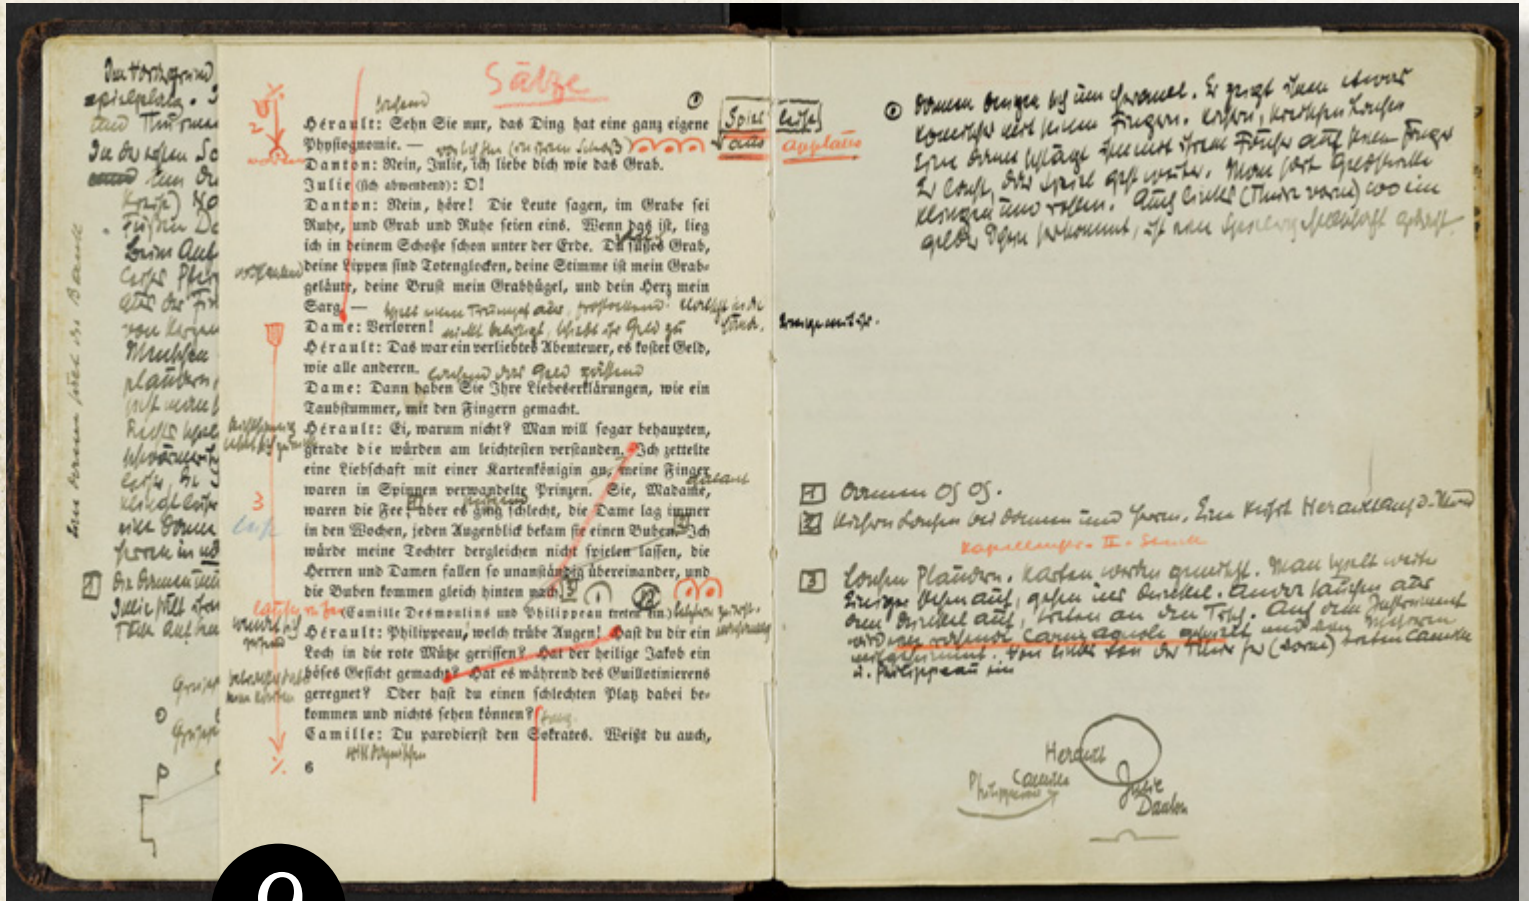
Each year an independent jury comprising five experts is responsible for selecting digitization projects and recommending them to the Berlin Senator for Culture for funding. Over the years, this jury has selected an assortment of projects that have shaped the special character of digiS in its diversity. Over the past six years, a total of 71 projects have been carried out in 32 Berlin institutions and almost one million primary files have been generated. In 2019 a further 13 projects will follow. [13]

CODING DA VINCI

But how is our cultural heritage negotiated digitally and which new forums and forms of appropriation are available? One field of experimentation conceived by digiS with partners from the open data movement and the cultural sector is the cultural data hackathon Coding da Vinci.[14] Since 2014, this event has united people from different fields across the nation: employees from cultural institutions of all branches, and interested programmers, designers, artists, makers, gamers, and culture and cultural data enthusiasts. Coding da Vinci is characterized by its specific format: during a development period of several weeks, interdisciplinary teams create prototypes of their creative ideas on the basis of openly licensed cultural data. More than 100 cultural and scientific institutions from Germany (and in some cases from Europe) have already participated in the cultural hackathon in the past as data providers. Just as many projects have been developed to showcase the enormous potential of our digital cultural heritage.[15] In every single project there has been highly individual confrontation with our cultural memory, our values, and our past.

The cultural data hackathon is an ideal experimental platform for the expansion of the analog museum space through digital means. Here, forward-looking approaches to a fundamentally changed interaction between cultural institutions and their (digital) users are emerging.






9

- 8 FHXB Friedrichshain-Kreuzberg Museum. Etching by the artist Erhard Gross. Title: Kreuzberg shop scenes 1950, 1980, 1985. Ed. 4/30, 1986.
- 9 Institut für Theaterwissenschaft der FU Berlin, THistorical theatre collections: Director's book "Dantons Tod". 1916.

SUMMARY

The institutionalization and consolidation of digiS and the Berlin funding program can be celebrated as a great success. It is a success for the general public: for culture, science, and society, which has produced this cultural heritage and which is constantly renegotiating the canon. With digitization, our heritage is kept alive and potentially shared to an extent that was not possible before. [16] To quote Aleida and Jan Assmann again: "Society needs a memory, just as an individual needs one: to know who we are and what we can expect, to orient ourselves, and to develop." [17]




$$\ln c^T x$$

$$x(\delta^-(w)) \geq 1 \quad \forall w \in V, (w, T) \neq \emptyset, r \notin w$$

$$x(\delta^+(v)) \geq x(\delta(v)) \quad \forall v \in V \setminus T$$

$$x(e) \in \{0, 1\}^n \quad \forall e \in E$$

$$\xrightarrow{LP} x(e) \in [0, 1]^n$$



S

OPTIMIZATION MEETS UPERCOMPUTING

Tackling Hard
Combinatorial Structures
By Parallel Algorithms

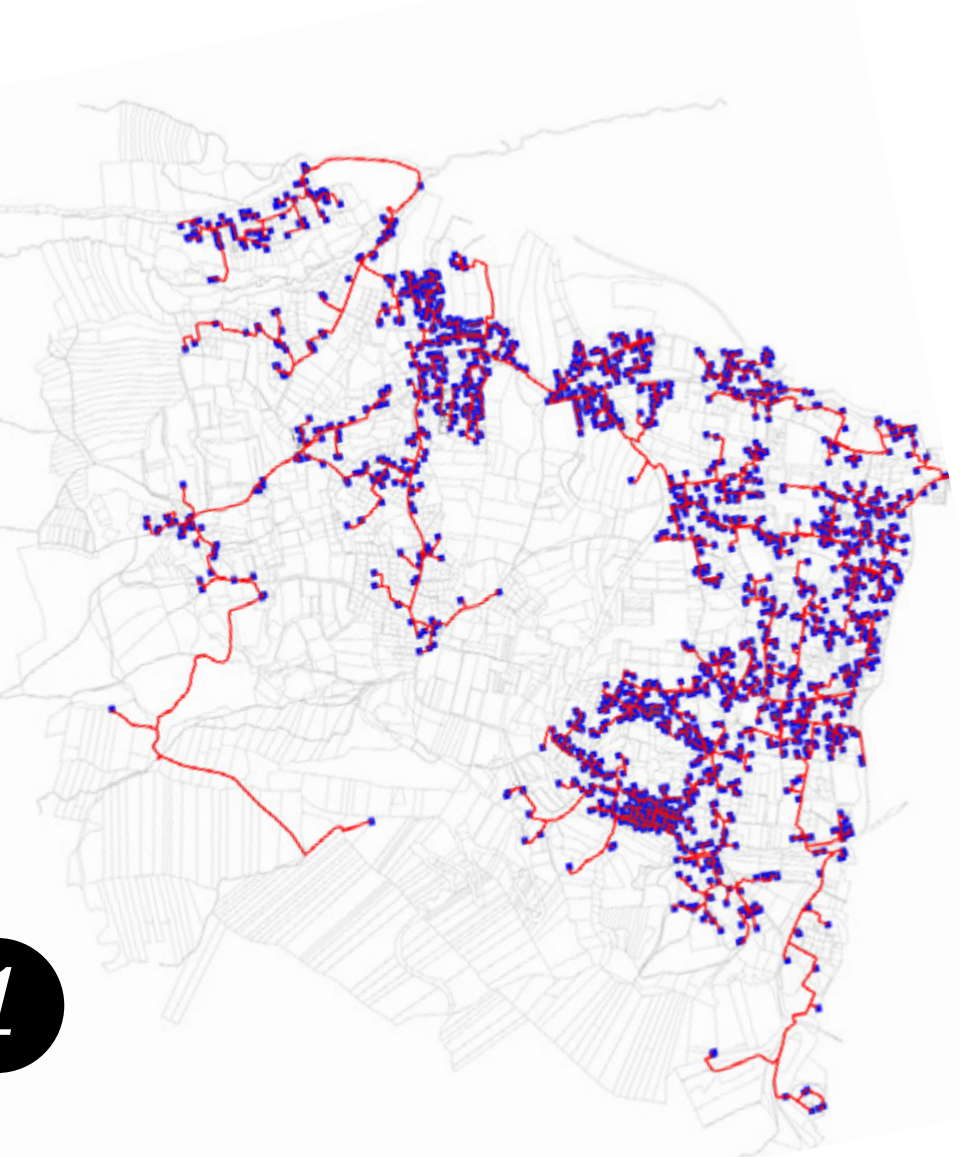
Since the turn of the century, the steady increase of single-core computing performance has come to a sobering halt. To overcome this roadblock, researchers at ZIB develop efficient algorithms for challenging optimization problems and work on their transformation to exploit the growing availability of parallel hardware.

TACKLING HARD COMBINATORIAL STRUCTURES THROUGH PARALLEL ALGORITHMS

Since the turn of the century, the steady increase in single-core computing performance has come to a sobering halt. To overcome this roadblock, researchers at ZIB develop efficient algorithms for challenging optimization problems and work on their transformation to exploit the growing availability of parallel hardware.

In a technologically driven society we increasingly encounter seemingly technical questions that require mathematical answers. Consider, for example, the construction of high-bandwidth telecommunication networks. Operators are faced with the task of connecting customers and deciding where to route new cables in the most cost-efficient way.

This scenario is also faced in many other applications. It leads to a fundamental optimization problem, the Steiner tree problem in graphs, which is named after 19th century Swiss mathematician Jakob Steiner.



STEINER TREES: A CLASSIC GRAPH PROBLEM REVISITED

From a mathematical viewpoint, such problems can be abstractly modeled by using a graph, a mathematical representation of what one might consider a “network”. A graph consists of a set of so-called vertices (corresponding to locations) that are connected or can be connected by so-called edges (cables). One can now associate a non-negative weight (distance or cost) with each edge. A subset of the vertices is designated as terminals (customers). A tree – a connected subgraph without any cycles – is called a Steiner tree if it connects all terminals. In applications, one is typically interested in computing a Steiner tree of minimum length or cost. This seemingly simple problem called the Steiner tree problem in graphs, or SPG for short, has proven to be notoriously hard to solve, both in theory and in practice.


For a theoretical mathematician the solution might seem astonishingly trivial: just check each of the finite number of Steiner trees and select one of minimum length. For practical purposes, however, such an approach requires a good deal of patience. Even for small problems with a few thousand vertices, this plain enumeration might take billions of years even on the fastest supercomputers that are currently available.

However, years of mathematical research have led to more targeted and more powerful algorithms that allow the solution of many real-world SPG instances in a matter of seconds on a regular desktop computer, in defiance of the exponential worst-case complexity. One of the fastest Steiner tree solvers available worldwide today is the SCIP-Jack solver developed at ZIB [1].


2




- 1 *A real-world telecommunication network from a city in Austria with an optimal Steiner tree as a solution marked in red.*
- 2 *From 1834 until his death, Swiss mathematician Jakob Steiner held a chair of geometry in Berlin.*
- 3 *A small example for an abstract graph representation of a Steiner tree problem. In the right picture, the solid edges indicate an optimal solution.*



3





HANDLING REAL-WORLD DIVERSITY

While there are practical applications that lead to a pure SPG formulation, it is far more common to encounter applications that require slight variations. As an example, consider the telecommunication network problem described above. For a company it may not pay off to connect every customer. So instead of aiming to include all customers, one might associate a non-negative prize with each customer and search the corresponding graph for a tree that maximizes the sum of its prizes minus its length. Indeed, this so-called prize-collecting Steiner tree problem has been used for planning fiber-optic networks in Austrian cities [2]. The various fields where one encounters problems that are closely related to the SPG range from computational biology and cancer research to computer chip design, computer vision, and even the deployment of drones. SCIP-Jack currently provides specialized support for the solution of thirteen different SPG variants.

In order to handle the large variety and complexity of Steiner tree problems, ZIB's SCIP-Jack solver incorporates a diverse set of algorithmic components, which closely interact. The overall solution algorithm follows the branch-and-bound paradigm, an advanced divide-and-conquer approach that explores the space of solutions by dividing it into increasingly smaller regions. Over- and underestimates of the objective function value are used to identify and exclude entire regions that do not contain optimal solutions as early as possible. Although the branch-and-bound paradigm is guaranteed to converge to a global optimal solution in finite time, this approach alone suffers from the same exponential worst-case behavior as plain enumeration. Hence, SCIP-Jack implements many supplementary techniques including graph transformations to a beneficial standard form, heuristics for finding primal solutions and dual bounds on the best possible objective value, domain propagation, cutting planes, and branching rules.

- 4 *Sophisticated preprocessing reductions are crucial to the performance of state-of-the-art solvers like SCIP-Jack. In this example, SCIP-Jack was able to reduce an input graph with initially 5,000 edges to only 17 edges.*
- 5 *The main SCIP-Jack developer Daniel Rehfeldt at the award ceremony of the PACE 2018 challenge.*
- 6 *In 2018, SCIP had registered 14,000 downloads from users worldwide.*

5



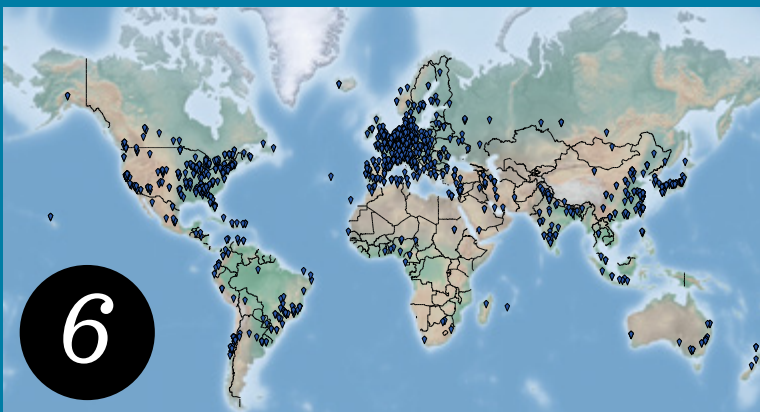
FIGHTING EXPONENTIAL SEARCH

The single most important component is perhaps the reduction techniques applied to the formulation prior to the branch-and-bound search. During this initial preprocessing phase, a plethora of mathematical arguments are applied in order to reduce the solution space of a given problem in such a way that each optimal Steiner tree in the reduced space can be equivalently transformed back to an optimal Steiner tree for the original problem. These techniques often lead to a dramatic reduction in problem size as shown in Figure 4. While the recent development of SCIP-Jack at ZIB has certainly spurred on many algorithmic

innovations, it should be noted that it is also based on decades of prior research. The classic Steiner tree problem in graphs has been the subject of hundreds of research articles, and there are thousands of articles devoted to its numerous variations.

Versatility is only one of SCIP-Jack's strong points. Over the years, it has also become the fastest solver worldwide for most of the thirteen problem classes it can handle. Just recently, the strong performance of SCIP-Jack was, for instance, demonstrated at the international PACE Challenge 2018 [3]. This competition evaluated solvers for a subclass of the

SPG, so-called fixed-parameter tractable SPGs, for which more efficient algorithms than for the general SPG are known. For PACE 2018, 33 teams of scientists from more than 16 countries submitted their software. The winning criterion was the number of benchmark problems that could be solved within 30 minutes. Even though SCIP-Jack does not contain any specialized algorithms for fixed-parameter tractable SPGs, it achieved top ratings in all three categories at PACE 2018: third place in track A, first place in track B, and second place in track C.



The SCIP-Jack Steiner tree solver is implemented as an extension of the more general SCIP optimization framework. With over 14,000 downloads per year, SCIP is one of the most used research frameworks for branch-cut-and-price algorithms worldwide – also because SCIP is currently the fastest mixed-integer linear and nonlinear programming solver that is openly accessible in source code. With the accompanying UG framework (see page 56/57) users can additionally exploit the computing power of parallel hardware architectures. The SCIP Optimization Suite is freely available for academic and nonprofit research at <http://scip.zib.de/>.



TREES OF TREES: PARALLELIZING SEARCH


Most of the recent advances in solving Steiner tree problems have focused on improving the performance of sequentially executed methods. However, in the age of increasingly available data sources, the growing complexity and dimension of real-world instances pose new challenges to sequential solvers. Even after executing today's sophisticated preprocessing techniques, the reduced problems may require an extremely expensive branch-and-bound search to compute a provably optimal Steiner tree solution. Motivated by this bottleneck, researchers at ZIB have developed one of the world's leading software frameworks that accelerates branch-and-bound searches by distributing the search effort using massively parallel hardware.

BRANCH-AND-BOUND ANOMALIES MULTIPLIED

At first sight, parallelizing the search effort in combinatorial optimization may seem simple. Naively, the 2^n possible assignments of Steiner trees with n edges can be explored m times as fast by simultaneously using m computers instead of a single one. However, state-of-the-art branch-and-bound solvers such as SCIP-Jack heavily rely on the fact that large parts of the solution space can be identified as suboptimal during the course of the algorithm. Premature enumeration effort invested in such regions makes parallelization ineffective.

More specifically, the divide-and-conquer approach organizes the exploration of the solution space hierarchically in form of a search tree. (Please be aware that these trees are unrelated to the Steiner trees in the previous article.) The top node represents the entire

solution space. Recursively, each node is subdivided into two or more disjoint regions one level below. The bottom level represents a single solution, but ideally nodes are identified as suboptimal at an early level and pruned without further subdivision. However, the a priori shape of a search tree is undetermined and its size can be strongly affected by the order in which nodes are processed. Already in a sequentially executed search, this dependency on the order leads to well-known anomalies and highly variable solution times. In a parallel search, these anomalies loom even larger. The challenge of creating an effective parallelization scheme is to dynamically determine the parts of the search tree where enumeration is indispensable and distribute this work.



7 Trees of trees: the dynamic load balancing of the UG framework distributes a complete search tree to parallelly operating subsolvers, distinguished by color.

8 UG has been used on many parallel hardware architectures worldwide. Here: The TITAN supercomputer at Oak Ridge National Laboratory, USA.



8

UG: THE UBIQUITY GENERATOR FRAMEWORK


Since many branch-and-bound solvers face these difficulties, researchers at ZIB have decided to develop a framework for parallelizing existing sequential branch-and-bound software in a generic fashion. The UG framework provides a collection of C++ classes, which define interfaces that can be customized for any such sequential base solver. They allow the translation of subproblems and solutions into a solver-independent form. Implementations of ramp-up, dynamic load balancing, and checkpointing and restarting mechanisms are available as a generic functionality. The latter is particularly important for obtaining a robust piece of software since the probability of hardware failure increases exponentially with the number of computers involved.

One flagship of the UG project has been ParaSCIP [4], a massively parallel version of the academic MIP solver SCIP. Since 2009, ParaSCIP has regularly succeeded in solving previously unsolved benchmark instances. Although UG relies on centralized load balancing with a single load coordinator, it has demonstrated the capability to utilize some of the largest supercomputers available today. The biggest computation conducted with ParaSCIP so far has harnessed 80,000 cores on the TITAN supercomputer at Oak Ridge National Laboratory. However, since UG provides interfaces for message-passing protocols in shared-memory environments, such as C++11 threads, it can be used at the same time to accelerate solvers on standard desktop computers.

BOOSTING SCIP-JACK: TACKLING OPEN STEINER TREE PROBLEMS

The versatility of UG has also been proven in the context of Steiner trees. With only 200 lines of glue code, the SCIP-Jack sequential solver could be transformed into a parallel Steiner tree solver that can exploit multicore desktop machines to large supercomputers. In the 11th DIMACS Implementation Challenge [5] held in 2014, SCIP-Jack was the only solver that could run on a distributed environment like a supercomputer. This could be harnessed to update 14 best known solutions out of 32 previously unsolved instances and close three open instances.

Certainly, the merely linear increase in computing power provided by supercomputers is only a limited remedy against the exponential worst-case complexity of combinatorial optimization problems. Nevertheless, the capability to momentarily boost the performance of state-of-the-art algorithms by large-scale parallelization can provide highly valuable insights for research. On several occasions, the opportunity to “look ahead in time” regarding the potential solvability of instances has already helped to guide the development of algorithmic innovations for the underlying sequential solvers.



In some applications the computational bottleneck lies not in the exploration of large search trees, but somewhat orthogonally in the solution of the so-called linear programming relaxation at each node of the search tree. One example are large-scale energy systems models that aid the integration of renewable energy sources into the future energy mix. To alleviate this bottleneck, ZIB works on extending PIPS-IPM, an open-source solver for block-structured LPs, within the BEAM-ME project funded by the Federal Ministry for Economic Affairs and Energy (www.beam-me-projekt.de).




Jan Skrzypczak | skrzypczak@zib.de | +49-30-8418-5474
Masoud Gholami Estahbanati | gholami@zib.de | +49-30-8418-5318

FAULT TOLERANCE IN DISTRIBUTED SYSTEMS

From Small to Large Scale



A server room with a glowing blue grid overlay on the floor and a row of monitors in the background. The monitors display various data and the word 'HLRI' is visible on one of them. The overall scene is dimly lit with blue and purple ambient lighting.

Fault tolerance, the ability of a system to survive different types of unexpected failures and to continue operation nevertheless, becomes essential for efficient resource usage and system stability. User applications have become more and more complex and resource demanding in recent years. They need to be executed and distributed among a growing number of machines. Consequently, they need larger systems consisting of more components, which also increases the rate of unexpected failures.

At ZIB, we study fault tolerance approaches on different scales, ranging from byte-granular replication to checkpoints and restarts in exascale computing.

BUILDING FAULT-TOLERANT SERVICES USING CONSENSUS

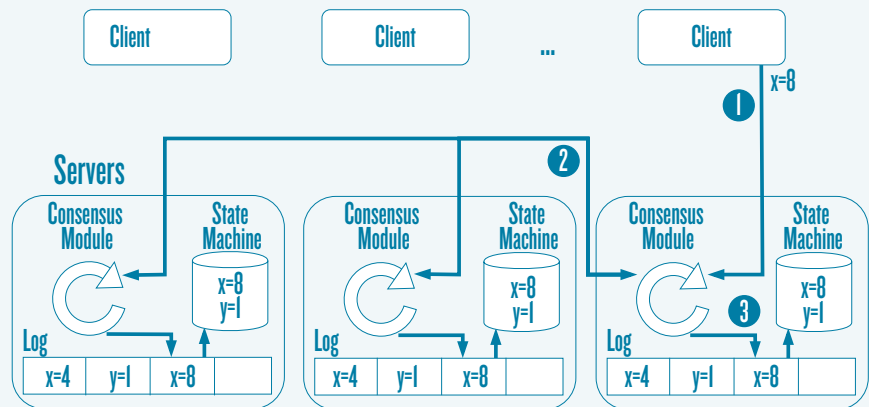
Today's society is largely dependent on network-based services. E-mail, instant messaging, collaboration tools, and data storage are just some examples of services that are used by end users and corporations alike on a daily basis. Many of these services are mission critical systems. Therefore, it is crucial to guarantee the service's availability even if one or more servers fail.

Adding redundancy to the system's architecture helps to achieve fault tolerance. One common approach, called state machine replication, is to replicate the state of the server across multiple machines. The state of all replicas must be kept synchronized at all times, which ensures that all replicas give the same reply to a client's request. This way, each replica can seamlessly take over the responsibilities of a crashed replica.

REPLICATED STATE MACHINE ARCHITECTURE

Many replicated state machine (RSM) implementations share the same architecture, as depicted in Figure 1. The RSM consists of a number of identical servers, called replicas. Clients can interact with the RSM by sending a read/write command to one of the replicas (step 1). However, there can be many clients that concurrently send requests to different replicas. As replicas must not diverge, concurrent commands must be ordered. This is done by employing a consensus

protocol like Paxos [1] (step 2). Each instance of Paxos can learn a single command by a majority decision. By chaining multiple Paxos instances together, a sequence of commands can be learned. The learned commands are entered into a command log and then applied to the state of the replica (step 3). Any client that submits a read command hereafter will observe the new state. If a replica fails, it can restore its state by replaying the commands entered into its log.



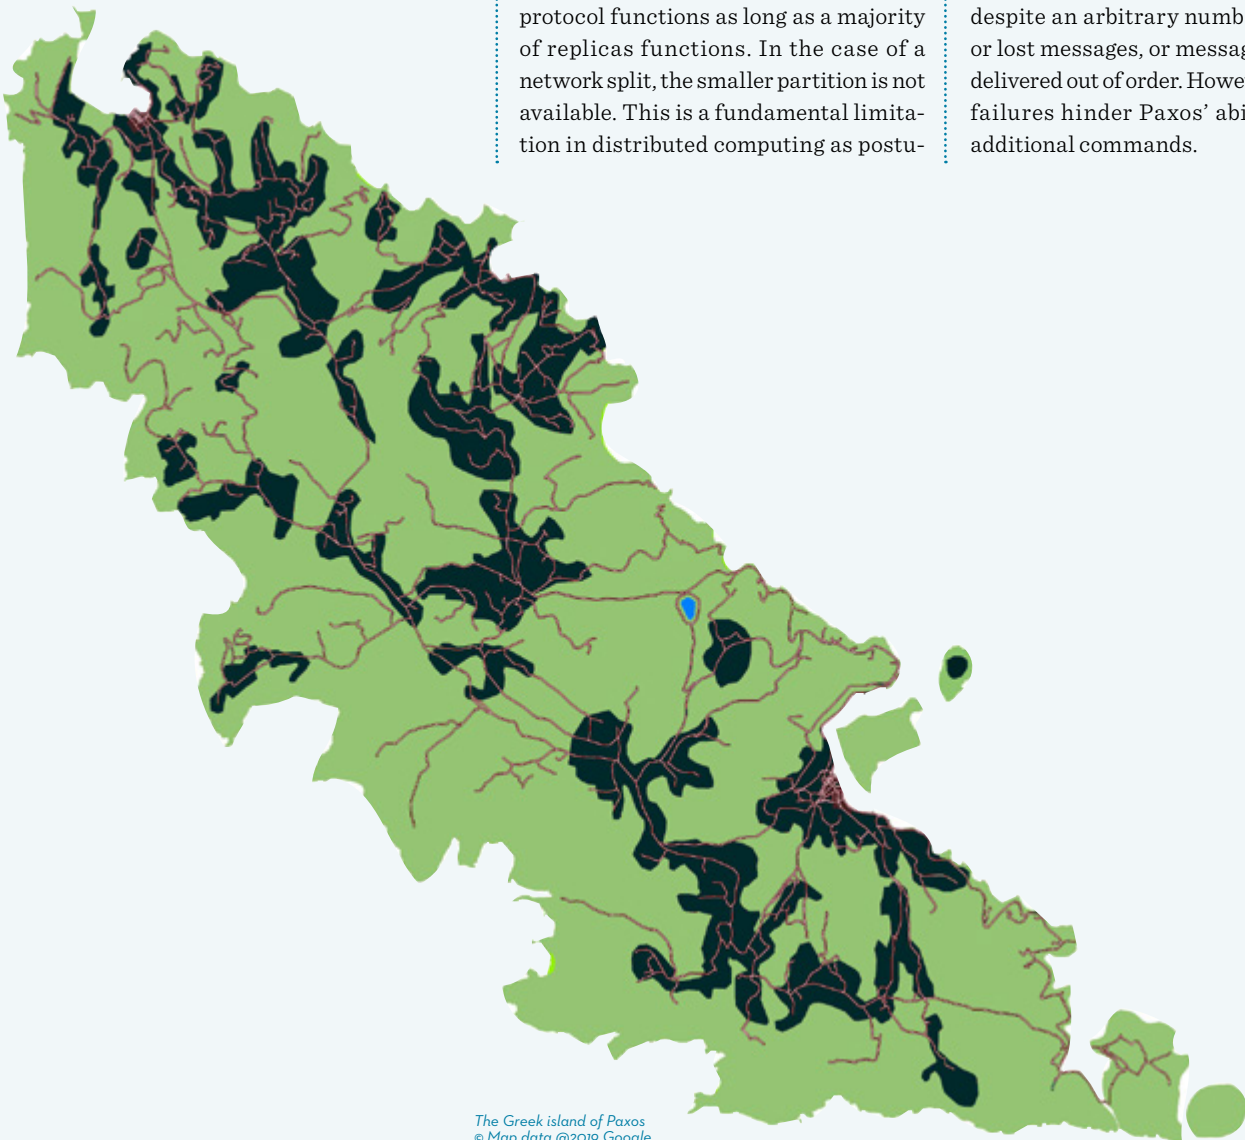
1 High-level architecture of a replicated state machine.

FAILURE TYPES

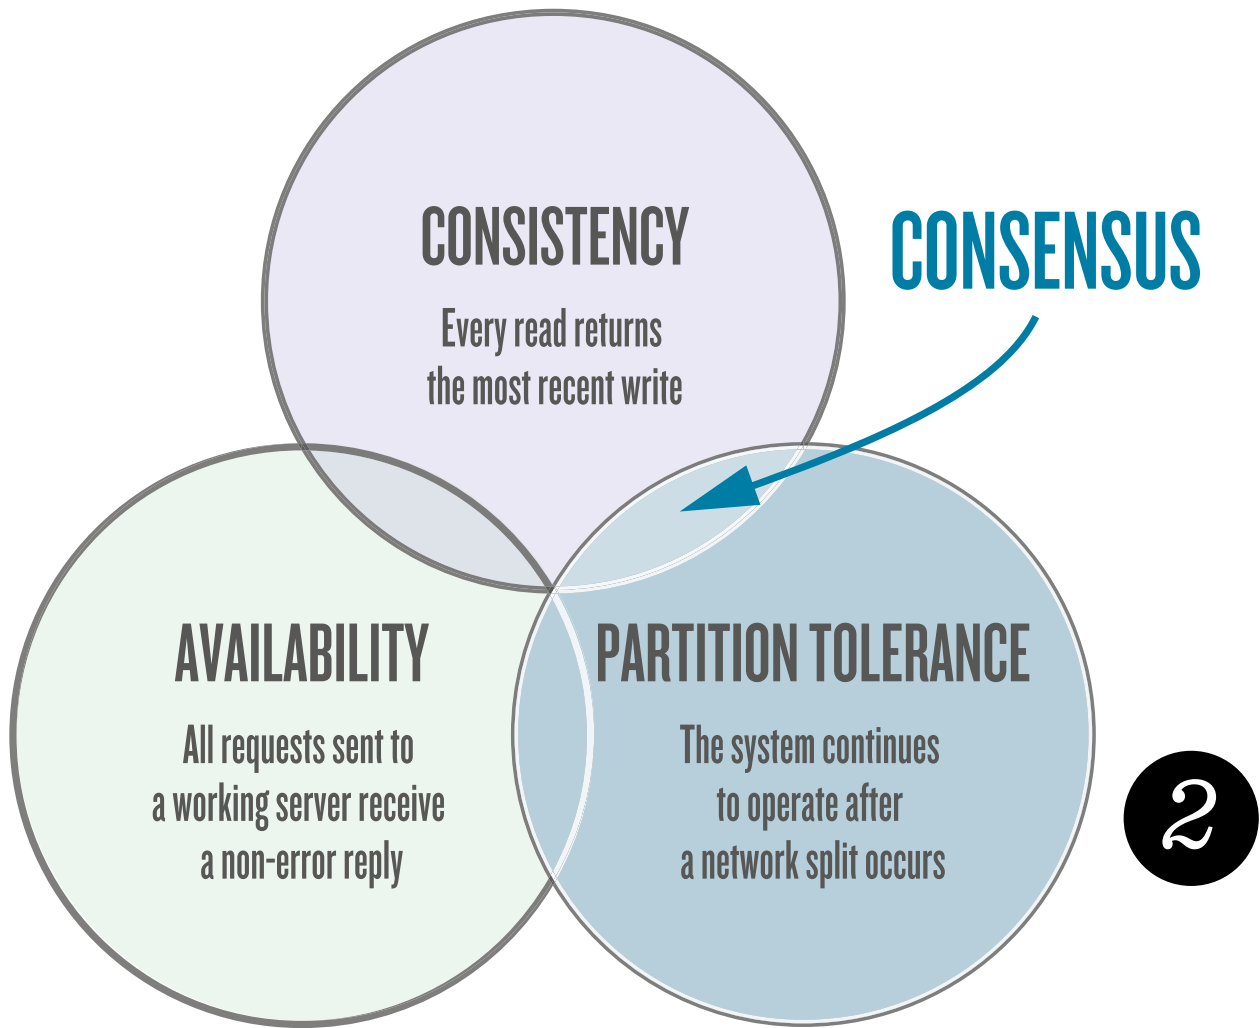
The types and number of failures that can be tolerated by the RSM approach strongly depend on the consensus protocol in use. Here at ZIB, we focus on the development of Paxos-derived protocols. As Paxos relies on majority decisions, the protocol functions as long as a majority of replicas functions. In the case of a network split, the smaller partition is not available. This is a fundamental limitation in distributed computing as postu-

lated by the CAP theorem [2]: no system can be both consistent and available when faced with network partitioning.

In addition to replica failures, Paxos is also resilient against communication failures. Replica states will never diverge despite an arbitrary number of delayed or lost messages, or messages that were delivered out of order. However, too many failures hinder Paxos' ability to learn additional commands.



The Greek island of Paxos
© Map data @2019 Google



CHALLENGES IN USING CONSENSUS

Two main challenges must be addressed when attempting to implement the current state-of-the-art approaches to state machine replication. First, the management of the command log requires additional maintenance. As new commands are always appended, it grows larger and larger over time. To restrict the log size, snapshot and other log truncation mechanisms must be implemented and it must be insured that these mechanisms are kept consistent with the state machine. Additionally, this incurs extra I/O costs on the persistent medium.

The second problem is that no consensus protocol exists that always terminates in a finite number of steps in a fully asynchronous system, i.e. a system

where communication delays are potentially unbounded. This limitation is also known as the FLP result [3]. For Paxos, this means that concurrent requests might never be ordered successfully as they can block each other. Of course, this is not acceptable for most practical systems if they have strict latency requirements. The usual approach to this problem involves the election of a leader process which acts as the coordinator of concurrent requests. However, such a leader often becomes a bottleneck of the system as all requests must be handled by it. In addition, it is a single point of failure – if the leader fails, the system becomes unavailable until a new leader is elected (see Figure 3).

- 2 *Classification of consensus in the CAP theorem.*
- 3 *Typical impact of leader failure on system throughput.*

FINE-GRAIN CONSENSUS WITHOUT LOGS

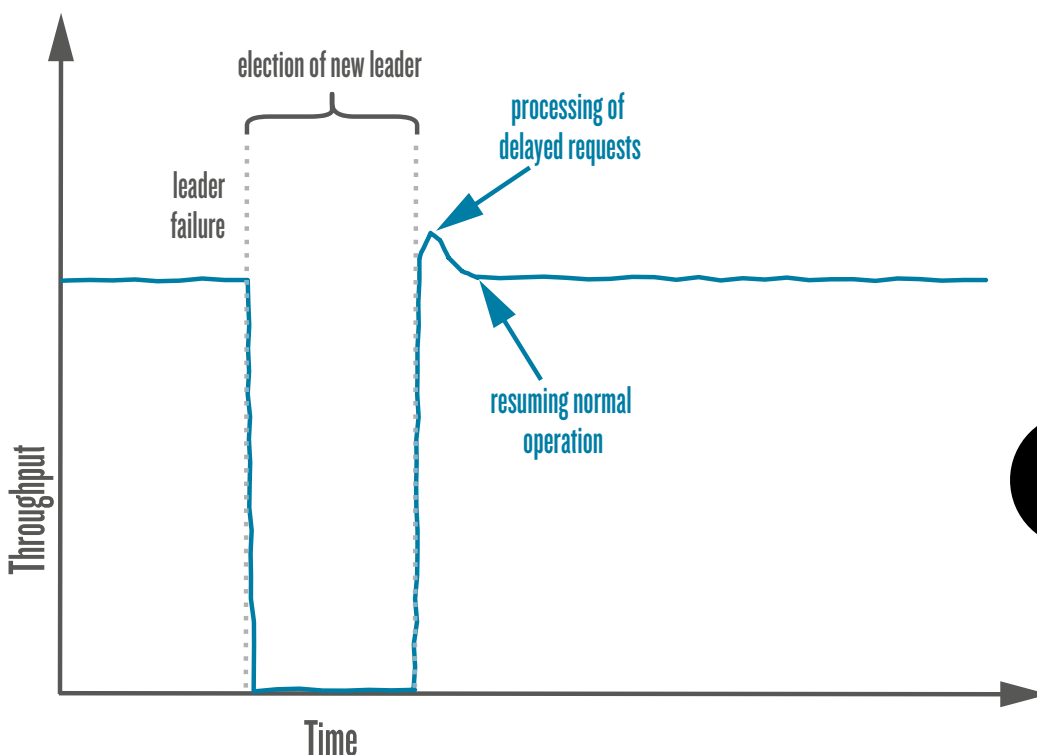
In our work, which is part of the DFG Priority Program 2037 (Scalable Data Management for Future Hardware), we aim to address the aforementioned challenges by designing a consensus approach that can be used on a more fine-grain level than existing protocols. Approaches maintaining a command log can only handle a handful of parallel consensus instances due to the complexity. Therefore, our work focuses on designing Paxos-based protocols that can make and implement a sequence of consensus decisions without the need of employing a command log. Instead of using consensus to learn a sequence of commands, our approach involves learning a sequence of state machine states. Thus, the command log is no longer needed.

This approach is suitable for replicating state machines that handle small data, for example a single key-value pair in a key-value store. Since our consensus protocol does not require log management, it is easy to implement each key-value pair in the store as its own replicated state machine. As an additional advantage, this design distributes the load across many consensus protocol instances. Depending on the access distribution of the keys, multiple leaders can be used, who all manage a subset of the key space, or the need for a leader can be alleviated entirely. In both cases, a singular process failure no longer leads to a period of unavailability of the whole system.



SCALARIS

To check the correctness and performance of our protocols, we implement them in the context of Scalaris. Scalaris is a transactional, distributed key-value store that is actively being developed at ZIB. It was the first NoSQL database that supported ACID-compliant multi-key transactions. In addition to our work on consensus, Scalaris is used to implement and evaluate transaction protocols, distributed hash tables, deterministic protocol testing, and more.

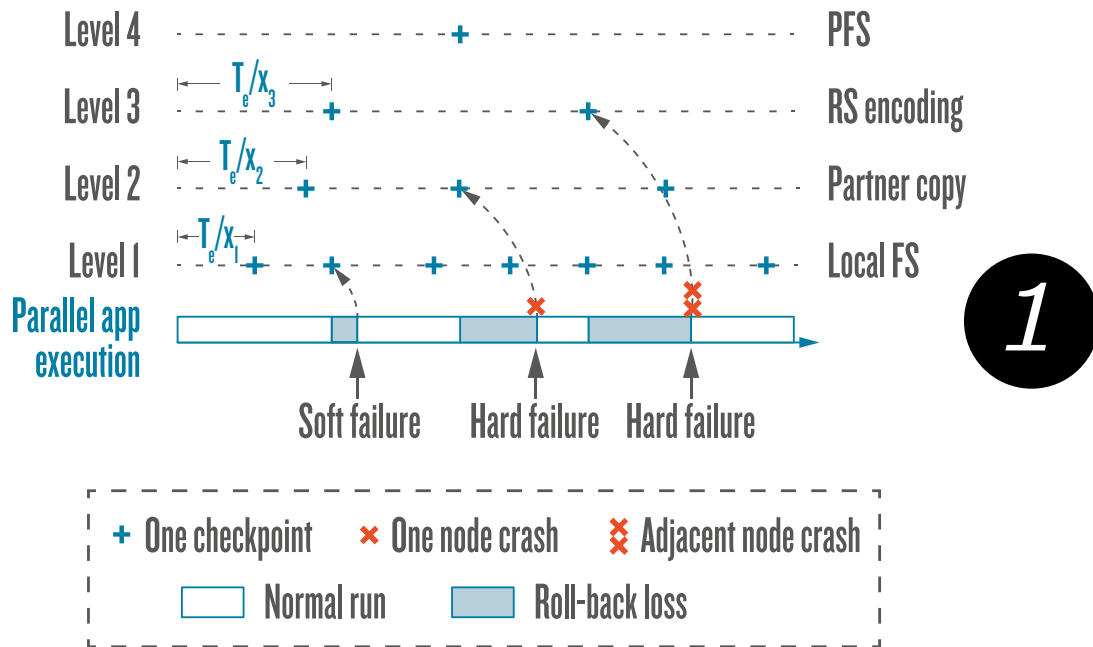


FAULT TOLERANCE: CHECKPOINT/RESTART MECHANISM

High Performance Computing (HPC) domain applications use a vast number of nodes for a relatively long time [4], and will continue to do so especially in the upcoming era of exascale computing. This increases the probability of encountering different types of failures frequently during an application's lifetime [5]. Many analyses are carried

out on HPC workload and failure logs to better inspect the rate, origin, and impact of common faults occurring in such systems. They indicate a pattern in which the number of failures on HPC systems has increased over time due to the tendency of executing larger scientific jobs on larger HPC systems.

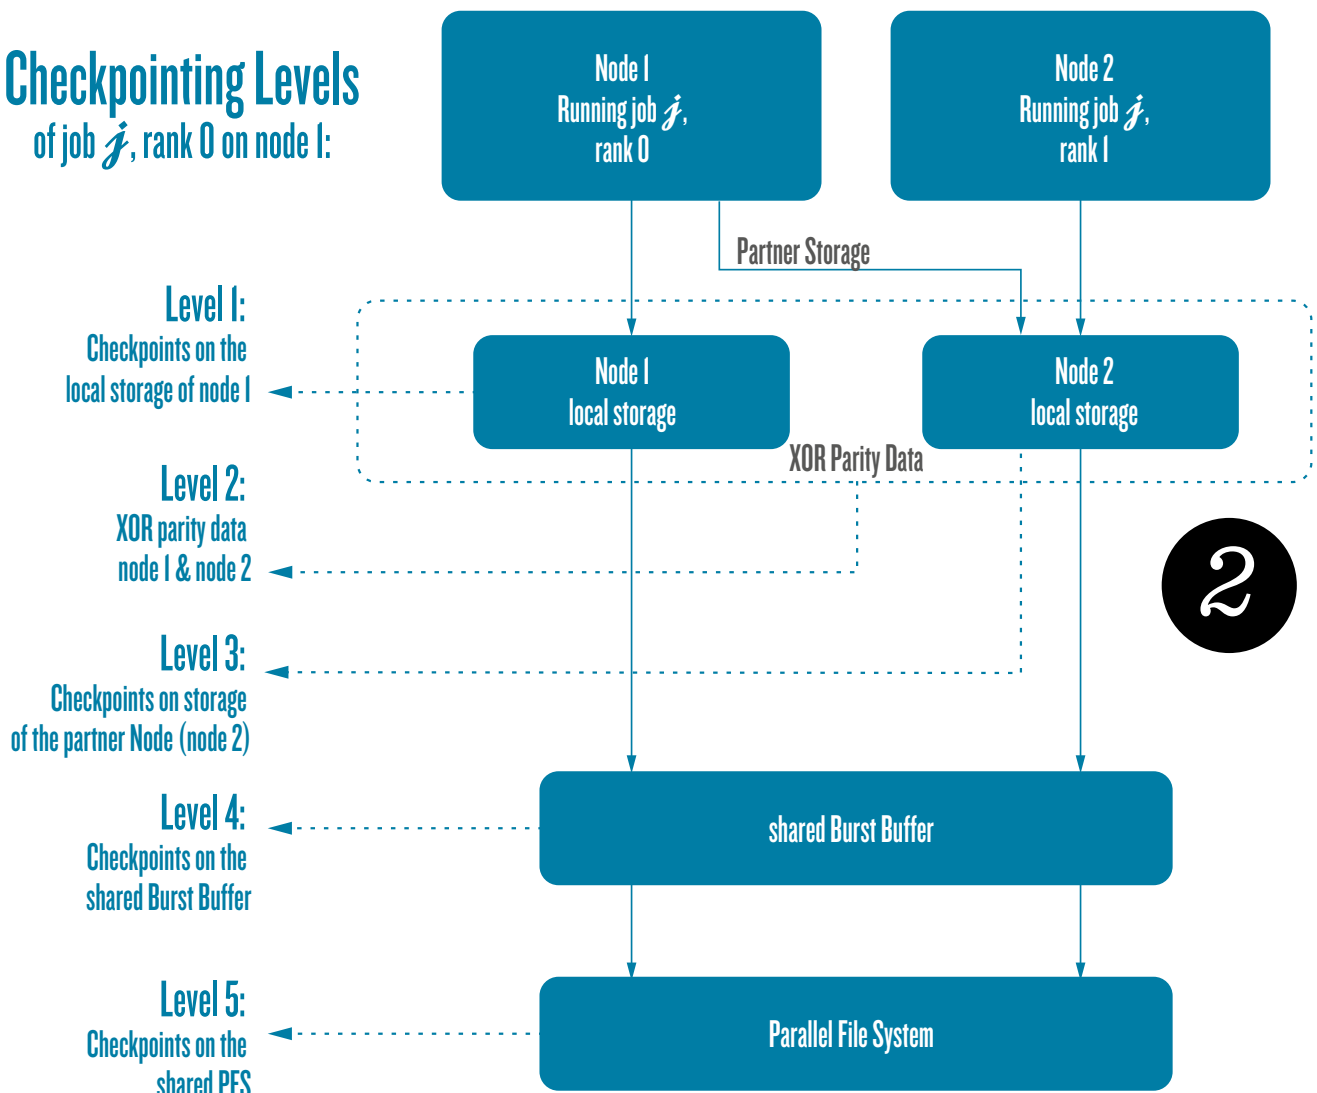
To overcome this problem, periodical checkpoints are introduced [6, 7] as an effective mechanism to survive failures during execution of applications – the job periodically stores its checkpoint in a Parallel File System (PFS) and restarts from the most recent checkpoint in case of a failure.



1

1 Different checkpoint levels in FTI [8].

Checkpointing Levels of job j , rank 0 on node 1:



MULTILEVEL CHECKPOINT/RESTART

For large jobs that have long lifetimes and large checkpoints, the classic checkpoint/restart (C/R) approach suffers from a significant decrease in performance as discussed in [8]. Therefore, we introduce – as part of the FFMK project in the DFG Priority Program 1648 (SPPEXA) – a multilevel checkpoint/restart model as a solution which considers various types of failures, and uses different types of storage devices that are able to survive specific levels of failures. On the one hand, transient levels encounter failures more frequently, and,

on the other hand, writing checkpoints on stable levels is more expensive [8].

Software libraries supporting multilevel checkpoint/restart models like SCR [9] and FTI [10], can deploy checkpoints at different levels. For example in the Fault Tolerance Interface (FTI), possible levels are local storage, RS encoding, partner nodes, PFS, and possibly other shared storage devices as shown in Figure 1. The frequency of the checkpoints for each level can be defined manually or computed by an algorithm.

COMBINING SYSTEM LEVEL WITH USER-DEFINED LEVEL

We combine system level with user-defined level checkpoints to benefit from the combined knowledge, which is available to both the system and the application [11]. The application can create checkpoints periodically when its internal state is small. The system, however, decides when to accept or reject the checkpoint requests. To make such decisions the system deploys its knowledge about the size of the job (number of nodes, processes, checkpoint sizes, etc.) and estimates the failure rate of nodes, storage devices, network components, different storage levels of checkpoints, and other jobs using the system at the same time.

According to the estimations of failure rates and applying the information about the job, the system decides when it is optimal to accept or reject the checkpoint offers from the job.

DIFFERENT FAILURE TYPES

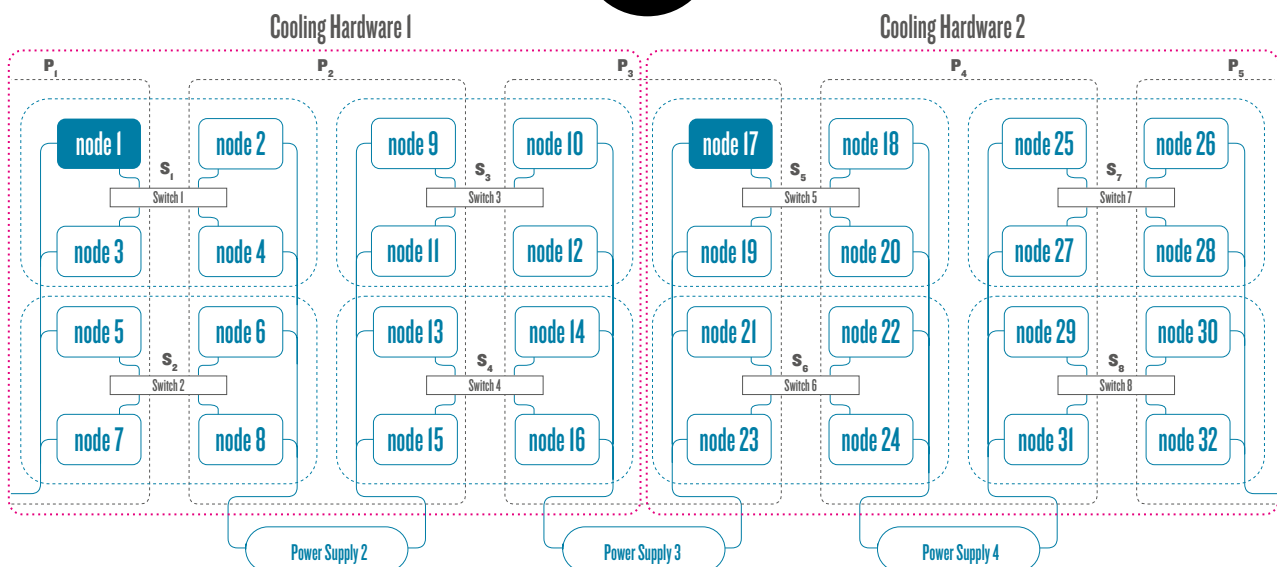
To support our approach, we developed a comprehensive model of an HPC job running on nodes and explored a wide range of failure types occurring in jobs, nodes, network switches, power supplies, and different types of correlated failures in general.

The model is later deployed to calculate the failure rates of all available C/R levels. This is the rate of errors that makes a specific C/R level inaccessible or useless for restart and therefore means that more stable levels should be engaged.

Additionally, by performing different simulations of a given job running on a cluster and deploying multilevel check-

point/restart, we were able to recognize further cases of multilevel failures and their impact on the job's execution. For instance, we explored the effect a chain of failures occurring in clusters has on the availability of different checkpoint levels. Another case was to investigate simultaneous correlated/uncorrelated faults and to introduce a mechanism to overcome their destructive side effect on the executions.


The discovered cases of failures (in simulations) are mathematically modeled and their occurrence rate is also considered in the computations of the desired failure rates of levels.




OPTIMUM CHECKPOINT INTERVALS OF LEVELS

Based on the calculated failure rates and also corresponding C/R cost of each level, the optimum checkpoint interval for each checkpoint level is defined, which minimizes the overall C/R overhead. The system then accepts or rejects the checkpoint requests based on the computed optimum checkpoint intervals. Finally, we select the optimum set of checkpoint levels that minimize the expected run time of a given job.

Additionally, because of higher checkpoint costs on stable levels, checkpoints are written asynchronously on the storage devices. This is done by writing the checkpoint on the local storage and letting the job continue with its computation early. Then, the checkpoint can be transferred to more stable levels by the library in the background.



2 A general cluster model containing different correlated groups of nodes: network, power supply, cooling hardware, etc.






Florian Wende | wende@zib.de | +49-30-8418-5323

THE DEEP PARALLELISM CHALLENGE

How to Use
Multiple Levels
of Parallelism
in Your Code



Architectural changes of modern CPUs urge software developers to approach multilevel parallelism in their programs in order to benefit most from the available compute capabilities. Adapting algorithms and codes are complementary tasks to cope with, which require a broad understanding of the CPU architecture and how to code for it effectively.

MODERN PARALLEL COMPUTER PLATFORMS

State-of-the-art high-performance computer (HPC) systems expose a deep hierarchy of parallelism to algorithm designers and software developers. Typical HPC systems are built from a large number of computer nodes together with a high-speed interconnect. A single node usually comprises two or more CPUs, each of which with tens of cores and per-core SIMD (single instruction, multiple data) units for data-parallel processing.

Figure 1 compares some of the key properties of the next-generation HLRN-IV system at ZIB against the previous generation. Even with a reduced node count, the new system provides a significantly increased compute performance because of a larger per-CPU core count and wider data parallel execution paths. Mastering these multiple levels of parallelism is today the only way to achieve a high level of sustained application performance. The increased intra-node parallelism means that software developers have to complement process-level with thread-level parallelism in order to benefit (more) from shared-memory communication and to avoid potential data duplication.

Additionally, threads need to execute mostly data-parallel instructions to fully exploit the available compute resources.

Beside orchestrating the multilevel parallel program execution so that it is evenly balanced across all execution units, adapting data layouts for improved data access is similarly relevant to programming for multilevel parallelism:

- data races must be resolved,
- data structures need to be set up such that (i) element accesses result in contiguous loads and stores from/to the main memory to improve or enable SIMD vectorization, and (ii) non-uniform memory access (NUMA) on multi-CPU nodes is approached, and
- additional levels in the memory hierarchy – high-bandwidth on-chip memory may require adaptations to the program code in order to be beneficial.

The most challenging part is maybe putting all this together at a reasonable cost and ideally such that code (and maybe performance) portability aspects are satisfied: different platforms (many-core CPUs, GPUs, vector processors) may require different optimization strategies and do not necessarily follow the same programming model.

¹ Key properties of the next-generation HLRN-IV complex at ZIB and the current HLRN-III complex (the HLRN-IV complex will be installed at the end of 2019).

HLRN-III

Complex at ZIB (2012-2019)

HLRN-IV

Complex at ZIB (from 2019)

1

1,872




>1,140 NODES

10



12 COMPUTE CABINETS

3,744



>2,280 CPUs

44,928



>109,000 CORES

DRAGON FLY



FAT TREE

INTERCONNECT TOPOLOGY


1.4 PFLOP/S



5× INCREASE IN

PEAK PERFORMANCE

120 TB



240 TB

DISTRIBUTED MAIN MEMORY

4.2 PB



8.1 PB

ONLINE STORAGE CAPACITY

THREAD-LEVEL PARALLELISM THROUGH OPENMP

Process- and thread-level parallelism are both means to effectively duplicate the executing context of a program (section) for running it on multiple CPU cores.

While processes usually have separate address spaces and coordinate their work through communication channels, threads are independent execution states within a process and share the process' address space. Beside the fact that threads are scheduled in the user space without kernel intervention, multi-threading allows for faster communication and data exchange through shared memory with (potentially) less data duplication compared to multipro-

cessing. Thinking of multi-node program executions, however, the combination of the two is obvious with at least one process per compute node.

Faced with programming for hybrid process- and thread-level parallelism, the Open Multi-Processing (OpenMP) standard enables portable shared-memory programming through an API and compiler directives for the C/C++ and Fortran programming language (see Listing 1), together with a set of environment variables to control the OpenMP runtime.

Over the last years, OpenMP has been extended to not only support a fixed

static/dynamic assignment of work to threads, but an adaptive one as well via tasking and a work-stealing scheduler. As most OpenMP functionality is accessible through compiler directives (**#pragma omp <directive [clauses]>**), code adaptations are manageable, at least in the first instance. However, the transition towards a thread-parallel code using shared memory usually involves some data layout changes to improve the per-thread data locality and to avoid data races. The latter can come along with additional thread synchronization and locking of shared data structures, and algorithmic changes, potentially.

Listing 1: Using OpenMP compiler directives to enable multi-threading in a C/C++ program.

```
// process context: 1 thread running
//
#pragma omp parallel
{
    int threadId=omp_get_thread_num();

    kernelA(data[threadId]);

    #pragma omp for schedule(dynamic)
    for(int i=0; i<n; ++i) kernelB(someOtherData[i]);
    #pragma omp task
    kernelC(againSomeOtherData[threadId]);
}
//
// process context: 1 thread running
```

THE MPI – OPENMP INTEROPERABILITY CHALLENGE

The Message Passing Interface (MPI) provides an abstraction of process-level parallelism beyond the node border. Various MPI implementations have been established in the HPC community for decades already, and are used to execute parallel programs across compute clusters or supercomputers. Even though the MPI standard has evolved over time and now offers support for shared-memory programming (using API calls), adapting codes for hybrid MPI+OpenMP parallelism seems more natural, as multi-threading is added as an optional level of parallelism that can be enabled or disabled through the environment. The process and thread configuration can therefore be configured in a flexible way. The benefits of executing a program with N/T processes and T threads per process

instead of N processes with no threads include

- a (potentially) reduced memory footprint of the execution due to less data duplication,
- a smaller number of processes loading the network and (potentially) larger message sizes per process that render the communication from latency- to network-bandwidth-dominated,
- and the use of multi-threaded library calls that might improve performance over their sequential equivalents.

In any case, the actual process-thread configuration is a tuning parameter that depends on the program itself and the system it is executed on. For applications with nontrivial communication characteristics one might even see a reduction

in performance if T becomes too large and N/T (the number of processes) becomes too small as a consequence, and if communication is triggered by processes only.

Even though MPI allows for multi-threaded communication, all major implementations suffer internal synchronization overheads if multiple (OpenMP) threads access a shared communicator. A recent development to compensate for this issue is the multiple-endpoint MPI feature, which basically assigns each thread a duplicate of the process' communicator (see Listing 2). This way, threads can communicate through their own communicator without synchronization.

Listing 2: User-level implementation of the multiple-endpoint MPI feature with OpenMP. Threads can communicate through their own communicators for improved MPI+OpenMP interoperability.

```
MPI_Comm endpoint[omp_get_max_threads()];
..
#pragma omp parallel
{
    int threadId=omp_get_thread_num();
    MPI_Comm_dup(MPI_COMM_WORLD, &endpoint[threadId]);
    ..
    #pragma omp for
    for(int i = 0; i<n; ++i)
        MPI_Allreduce(endpoint[threadId],data[threadId],...);
    ..
}
```

DATA-LEVEL PARALLELISM THROUGH SIMD VECTORIZATION

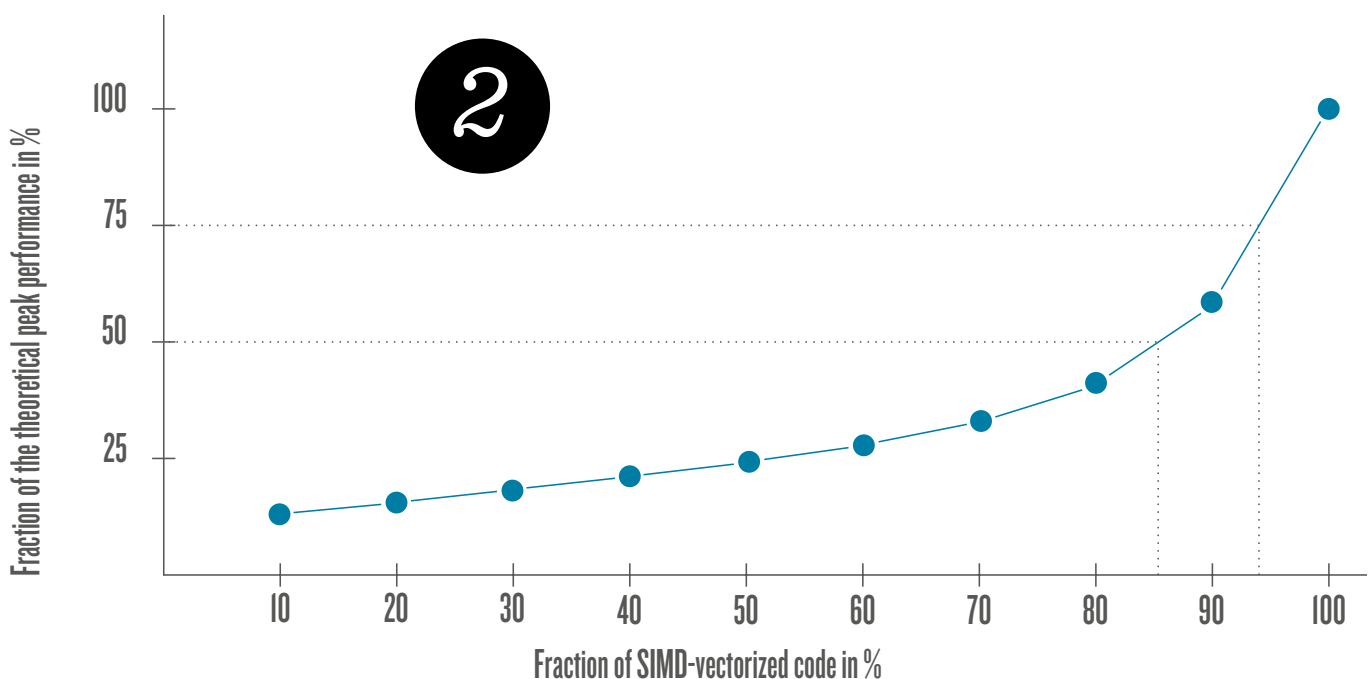
Data-level parallelism refers to executing the same code on all data elements or subsets of the data either consecutively in a pipeline (vector processors), concurrently (array processors), or in a mixed fashion. Modern x86 CPUs follow the latter by implementing the single instruction, multiple data (SIMD) execution model for short vectors with the execution happening in the process or thread context.

With up to eight 64-bit execution streams per SIMD unit per x86 CPU core, data-level parallelism can contribute significantly to the program performance if the number of arithmetic operations per byte loaded is sufficiently large and

the computation is compute-bound. The latter, however, is not necessarily satisfied and not all sections of a program are necessarily data-parallel: any kind of communication or access to shared resources add to the sequential part of the execution. In reality, only a fraction of the theoretical peak performance can therefore be achieved as shown in Figure 2 for a certain portion of the program code being (SIMD-)vectorized.

While processes (threads) act independently, SIMD-execution streams (or SIMD lanes) are organized into groups of size VL (vector length or SIMD width) according to the SIMD-register size in bytes (currently up to 64 bytes) and the byte size of the data being processed

(between 1 and 8 bytes), meaning that $VL=8 \cdot 64$. Unlike processes or threads, SIMD lanes in the same group share one program state. Data-parallel execution is therefore most efficient without any control flow divergences. The latter, however, cannot be avoided in any case and is resolved by executing the different branch paths one after the other and blending all intermediate results together for the final result (see Figure 3). Depending on the degree of divergence, SIMD execution might even lower the program performance because of SIMD-related overheads and potentially lower clock frequencies in the case of processing SIMD-instruction streams.



Optimizing program codes for SIMD execution, however, is an important investment into the future of a code as many upcoming architectures will draw on data-level parallelism to reach their theoretical peak performance.

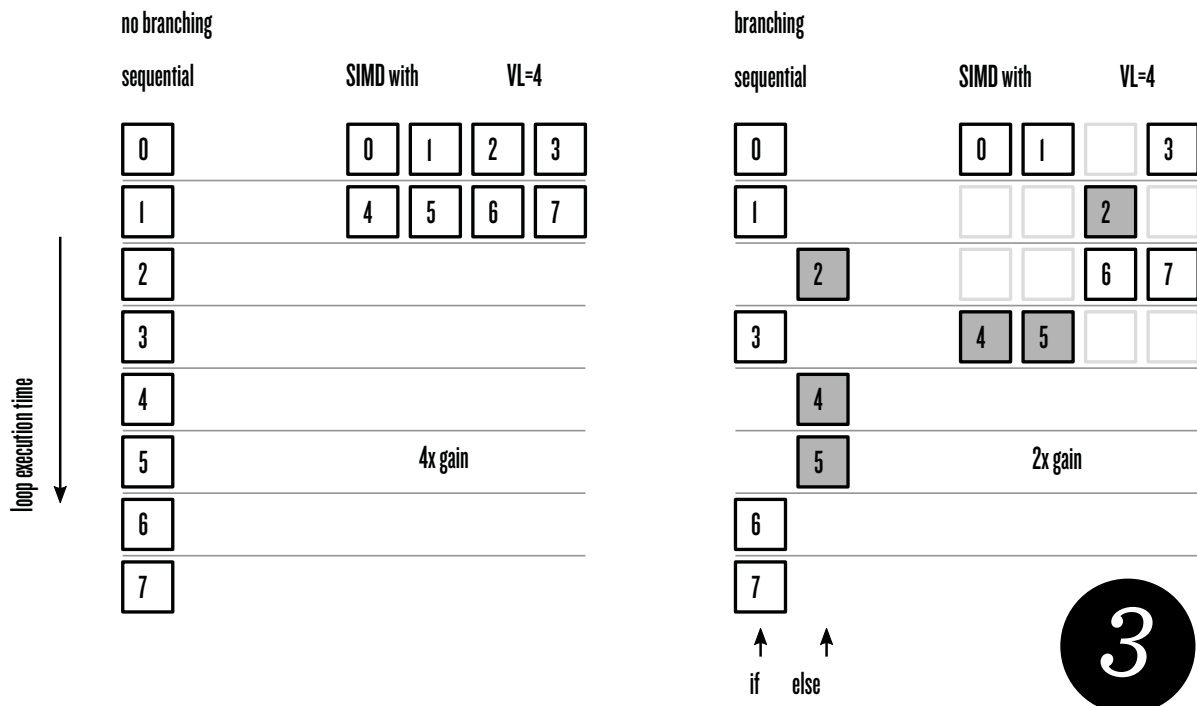
With the OpenMP 4 standard, not only multi-threading can be targeted through compiler directives, but also loop (`#pragma omp simd`) and function (`#pragma omp declare simd`) vector-

ization. Using any of the latest major compilers supporting OpenMP (such as Intel, GNU, Clang/LLVM), even complex data-parallel codes can be vectorized in a portable way.


Everything beyond that requires advanced SIMD-programming techniques to explicitly instruct the compiler on which vector data types and operations to use.

2 Amdahl's law applied to SIMD parallelism with eight data-parallel execution streams.

3 Sequential and SIMD-like loop execution with and without control flow divergence (branching in this case). Boxes with ascending numbers correspond to successive loop iterations. Different colors represent different instructions. Although boxes of the same color can all be executed at once, they are processed in groups of size VL=4 in the illustration.



PUTTING IT ALL TOGETHER: A REAL-WORLD EXAMPLE



In a joint collaborative effort with the University of Vienna, Intel Corp., and NERSC, we applied all optimization strategies as described in the previous sections to the Vienna Ab initio Simulation Package – VASP (more details can be found in [1]). In particular, we

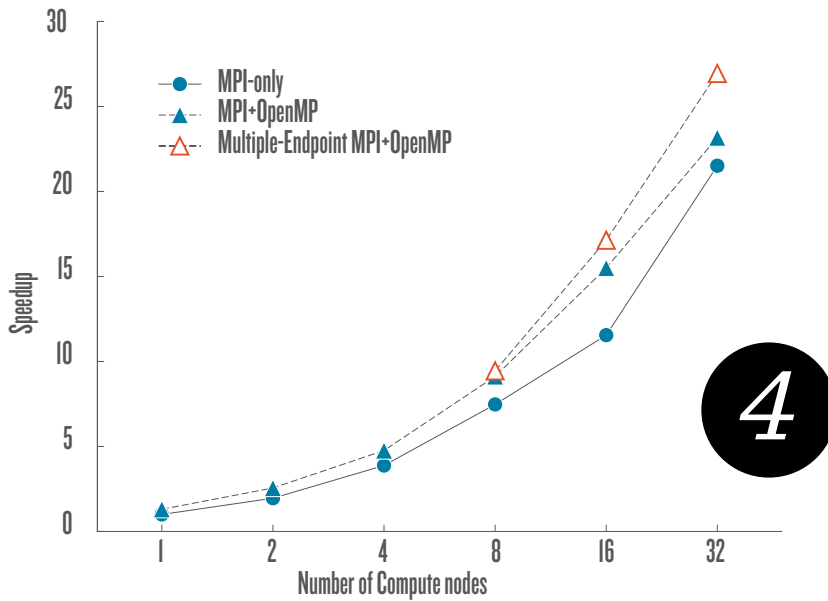
- made the transition from a pure MPI to a hybrid MPI+OpenMP code base at different levels by
 - using OpenMP compiler directives for loop parallelization,
 - adapting data structures for multi-threaded access, and
 - taking advantage of multi-threaded libraries (BLAS and FFT) where possible.
- implemented the multiple-endpoint MPI feature according to Listing 2 as a proof of concept to assess its potential.
- enabled SIMD vectorization using OpenMP 4 for relevant data-parallel loops and within VASP's GGA routine. We demonstrate how to vectorize along the calling tree for different GGA functionals using a high-level vector approach which originates from manual vector programming, but draws on OpenMP 4 loop vectorization for portability reasons.

The Vienna Ab initio Simulation Package – VASP is a state-of-the-art electronic structure code that is widely used for the atomic scale modeling of material properties of condensed matter in the fields of computational materials science, applied materials research, and development of novel materials.

VASP [2][3] supports a broad range of methods to approximate electron-electron interactions, ranging from density functional theory (DFT), the Hartree-Fock (HF) method, and hybrid HF/DFT functionals, to many-body perturbative approaches based on the random phase approximation (RPA) such as GW and ACFDT [4][5][6]. Electronic interactions are described at the all-electron level by means of the projector augmented wave method (PAW)[7][8].

VASP is one of the most used codes on the HLRN system at ZIB and consumes a significant portion of the available compute resources. Therefore, any performance optimizations of VASP translate directly into advantages for many scientists using the HLRN resources for their scientific projects.

Strong Scaling of the Development Version of VASP
(2x20 cores per compute node, 10 OpenMP threads per MPI process)



The results of a strong scaling study with our development version of VASP on the Intel Skylake architecture are shown in Figure 4. The hybrid MPI+OpenMP version shows significantly improved scalability over pure MPI, which for the most part originates from less time being spent in MPI communication. For higher node counts, the performance potential of multiple-endpoint MPI in the multi-threaded context can be clearly demonstrated.

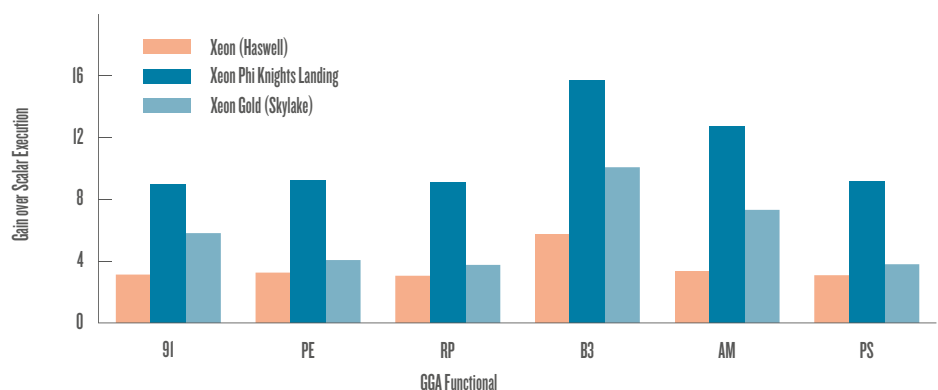
The SIMD vectorization of complex loops, for example those with deep calling trees and non-trivial control flow inside the loop body, as is the case for VASP's GGA routine, poses a huge challenge to the programmer. To deal with this, we developed a high-level vector approach which enables an effective OpenMP 4-based SIMD vectorization of these kinds of loops. Figure 5 shows the performance

gains of our SIMD version of the GGA routine over the original scalar version for different GGA density functionals and three processor platforms. For an execution with 64-bit (double precision) words, we expect a factor 4 gain (with one 256-bit SIMD unit per CPU core) on the Haswell platform and a gain up to a factor of 16 on Skylake and Xeon Phi Knights Landing (both with two 512-bit SIMD units per CPU core) respectively. While the Haswell and the Xeon Phi platform meet the expectations, the execution on Skylake results in a performance gain of factor 8 and 16 only, which to some extent goes back to an instruction mix with 256-bit and 512-bit SIMD vectors because of performance reasons on this platform. The performance potential of SIMD data-parallel processing, however, is clearly demonstrated even if complex code structures are considered.

- 4 *Scaling of the development version of VASP on Intel Skylake architecture. Test case is a unit cell with 256 Si atoms and the HSE density functional is used for the computation.*
- 5 *Performance gains of VASP's GGA routine for different GGA density functionals and processor platforms over the sequential execution (a recent Intel compiler version is used for code generation). at once, they are processed in groups of size VL=4 in the illustration.*

5

SIMD-Vectorization of VASP's GGA Routine with the Intel Compiler for different GGA Functionals and Platforms



ZIB PUBLICATIONS

PEER-REVIEWED (156)

Muluneh G. Abebe, Aimi Abass, Guillaume Gomard, Lin Zschiedrich, Uli Lemmer, Bryce S. Richards, Carsten Rockstuhl, Ulrich W. Paetzold (2018). **Rigorous Wave-Optical Treatment of Photon Recycling in Thermodynamics of Photovoltaics: Perovskite Thin-Film Solar Cells.** *Phys. Rev. B*, 98:075141. <https://doi.org/10.1103/PhysRevB.98.075141>

Mohsen Akbari-Shandiz, Paul Boulos, Stefan Sævarsson, Heiko Ramm, Chun Kit Fu, Stephen Miller, Stefan Zachow, Carolyn Anglin (2018). **Changes in Knee Shape and Geometry Resulting from Total Knee Arthroplasty.** *Journal of Engineering in Medicine*, 232(1):67-79. <https://doi.org/10.1177/0954411917743274>

M. Alchikh, T. Conrad, Ch. Hoppe, X. Ma, E. Broberg, P. Penttinen, J. Reiche, B. Biere, B. Schweiger, B. Rath (2018). **Are We Missing Respiratory Viral Infections in Infants and Children? Comparison of a Hospital-Based Quality Management System with Standard of Care.** *Clinical Microbiology and Infection*, 1-1. <https://doi.org/10.1016/j.cmi.2018.05.023> (accepted for publication 2018-06-12)

Felix Ambellan, Alexander Tack, Moritz Ehlke, Stefan Zachow (2019). **Automated Segmentation of Knee Bone and Cartilage Combining Statistical Shape Knowledge and Convolutional Neural Networks: Data from the Osteoarthritis Initiative.** *Medical Image Analysis*, 52(2):109-118. <https://doi.org/10.1016/j.media.2018.11.009> (epub ahead of print 2018-11-17)

Felix Ambellan, Alexander Tack, Moritz Ehlke, Stefan Zachow (2018). **Automated Segmentation of Knee Bone and Cartilage combining Statistical Shape Knowledge and Convolutional Neural Networks: Data from the Osteoarthritis Initiative.** In *Medical Imaging with Deep Learning*. Eds.: Shi, Yinghuan, Suk, Heung-Il, Liu, Minigxia (Eds.)

Lovis Anderson, Benjamin Hiller (2018). **A Sweep-Plane Algorithm for the Computation of the Volume of a Union of Polytopes.** *OR 2018 Proceedings*. (accepted for publication 2018-11-16)

Daniel Anderson, Gregor Hendel, Pierre Le Bodic, Jan Merlin Viernickel (2018). **Clairvoyant Restarts in Branch-and-Bound Search Using Online Tree-Size Estimation.** In *AAAI-19: Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*. (accepted for publication 2018-10-31)

Carlo Barth, Christiane Becker (2018). **Machine Learning Classification for Field Distributions of Photonic Modes.** *Commun. Phys.*, 1:58. <https://doi.org/10.1038/s42005-018-0060-1>

Isabel Beckenbach, Ralf Borndörfer (2018). **Hall's and König's Theorem in Graphs and Hypergraphs.** *Discrete Mathematics*, 341(10):2753-2761. <https://doi.org/10.1016/j.disc.2018.06.013>

Christiane Becker, Sven Burger, Carlo Barth, Phillip Manley, Klaus Jäger, David Eisenhauer, Grit Köppel, Pavel Chabera, Junsheng Chen, Kaibo Zheng, Tõnu Pullerits (2018). **Nanophotonic-Enhanced Two-Photon-Excited Photoluminescence of Perovskite Quantum Dots.** *ACS Photonics*, 5:4668. <https://doi.org/10.1021/acsp Photonics.8b01199>

Felix Binkowski, Lin Zschiedrich, Sven Burger (2018). **Analysis of Light-Matter Interaction of Optical Sources with Dispersive Nanoresonators via Contour Integration.** In *XXVI International Workshop on Optical Wave & Waveguide Theory and Numerical Modelling, Proceedings*, 40.

Andreas Bitttracher, Péter Koltai, Stefan Klus, Ralf Banisch, Michael Dellnitz, Christof Schütte (2018). **Transition Manifolds of Complex Metastable Systems: Theory and Data-Driven Computation of Effective Dynamics.** *Journal of Nonlinear Science*, 28(2):471-512. <https://doi.org/10.1007/s00332-017-9415-0>

Ralf Borndörfer, Stephan Schwartz, Gerald Bartz (2018). **The Graph Segmentation Problem.** In *INOC 2017 - 8th International Network Optimization Conference, Volume 64 of Electronic Notes in Discrete Mathematics*, Pages 35-44.

Ralf Borndörfer, Matthias Breuer, Boris Grimm, Markus Reuther, Stanley Schade, Thomas Schlechte (2018). **Timetable Sparsification by Rolling Stock Rotation Optimization.** In *Operations Research 2017*, 723-728. https://doi.org/10.1007/978-3-319-89920-6_96

Ralf Borndörfer, Marika Karbstein, Christian Liebchen, Niels Lindner (2018). **A Simple Way to Compute the Number of Vehicles that Are Required to Operate a Periodic Timetable.** In *18th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2018)*, Volume 65 of *Open-Access Series in Informatics (OASIS)*, Pages 16:1-16:15. <https://doi.org/10.4230/OASIS.ATMOS.2018.16>

Thomas Breuer, Michael Bussieck, Karl-Kien Cao, Felix Cebulla, Frederik Fiand, Hans Christian Gils, Ambros Gleixner, Dmitry Khabi, Thorsten Koch, Daniel Rehfeldt, Manuel Wetzel (2018). **Optimizing Large-Scale Linear Energy System Problems with Block Diagonal Structure by Using Parallel Interior-Point Methods.** In *Operations Research Proceedings 2017*, 641-647. https://doi.org/10.1007/978-3-319-89920-6_85

Gabrio Caimi, Frank Fischer, Thomas Schlechte (2018). **Railway Track Allocation**. In *Handbook of Optimization in the Railway Industry*, Volume 268 of *International Series in Operations Research & Management Science*, Pages 141–160, Springer International Publishing. <https://doi.org/https://doi.org/10.1007/978-3-319-72153-8>

Natascia Castagna, Jacques Morel, Luc Testa, Sven Burger (2018). **Modelling of Standard and Specialty Fibre-Based Systems Using Finite Element Methods**. *Proc. SPIE*, 10683:1068336. <https://doi.org/10.1117/12.2307372>

Duote Chen, Phillip Manley, Philipp Tockhorn, David Eisenhauer, Grit Köppel, Martin Hammerschmidt, Sven Burger, Steve Albrecht, Christiane Becker, Klaus Jäger (2018). **Nanophotonic Light Management for Perovskite-Silicon Tandem Solar Cells**. *J. Photonics Energy*, 8:022601. <https://doi.org/10.1117/1.JPE.8.022601>

Giovanni Ciccotti, Mauro Ferrario, Christof Schütte (2018). **Molecular Dynamics vs. Stochastic Processes: Are We Heading Anywhere?**. In *Special Issue: Understanding Molecular Dynamics via Stochastic Processes*, *Entropy*, 20(5). <https://doi.org/10.3390/e20050348>

Liviu Coconu, Brygg Ullmer, Philipp Paar, Jing Liu, Miriam Konkel, Hans-Christian Hege (2018). **A Smartphone-Based Tangible Interaction Approach for Landscape Visualization**. In *PerDis '18 Proceedings of the 7th ACM International Symposium on Pervasive Displays, Munich, Germany, June 6–8, 2018*, No. 23. <https://doi.org/10.1145/3205873.3210707>

Natasa Djurdjevac Conrad, Daniel Fuerstenau, Ana Grabundzija, Luzie Helfmann, Martin Park, Wolfram Schier, Brigitta Schütt, Christof Schütte, Marcus Weber, Niklas Wulkow, Johannes Zonker (2018). **Mathematical Modeling of the Spreading of Innovations in the Ancient World**. *eTopoi. Journal for Ancient Studies*, 7.

Natasa Djurdjevac Conrad, Luzie Helfmann, Johannes Zonker, Stefanie Winkelmann, Christof Schütte (2018). **Human Mobility and Innovation Spreading in Ancient Times: a Stochastic Agent-Based Simulation Approach**. *EPJ Data Science*, 7(1). <https://doi.org/10.1140/epjds/s13688-018-0153-9>

M. Dell'Amico, N.S. Hadjidimitriou, T. Koch, M. Petkovic (2018). **Forecasting Natural Gas Flows in Large Networks**. *Machine Learning, Optimization, and Big Data. MOD 2017, Lecture Notes in Computer Science*(Volume 10710):158–171. https://doi.org/https://doi.org/10.1007/978-3-319-72926-8_14

Peter Deuffhard (2018). **The Grand Four. Affine Invariant Globalizations of Newton's Method**. *Vietnam Journal of Mathematics*, 46(4):761–777. <https://doi.org/10.1007/s10013-018-0301-3>

Manuel Dibak, Mauricio J. del Razo, David De Sancho, Christof Schütte, Frank Noé (2018). **MSM/RD: Coupling Markov State Models of Molecular Kinetics with Reaction-Diffusion Simulations**. *Journal of Chemical Physics*, 148(214107).

Luca Donati, Martin Heida, Bettina G. Keller, Marcus Weber (2018). **Estimation of the Infinitesimal Generator by Square-Root Approximation**. *J. Phys.: Condens. Matter*, 30(42):425201–425201. <https://doi.org/10.1088/1361-648X/aadfc8>

Belmiro P.M. Duarte, Guillaume Sagnol, Weng Kee Wong (2018). **An Algorithm Based on Semidefinite Programming for Finding Minimax Optimal Designs**. *Computational Statistics & Data Analysis*, 119:99–117. <https://doi.org/10.1016/j.csda.2017.09.008>

Rainald M. Ehrig, Markus O. Heller (2018). **On Intrinsic Equivalences of the Finite Helical Axis, the Instantaneous Helical Axis, and the SARA Approach. A Mathematical Survey**. *Journal of Biomechanics*. <https://doi.org/10.1016/j.jbiomech.2018.12.034> (accepted for publication 2018-12-19)

Konstantin Fackeldey, Alexander Sikorski, Marcus Weber (2018). **Spectral Clustering for Non-reversible Markov Chains**. *Computational and Applied Mathematics*, 37(5):6376–6391. <https://doi.org/https://doi.org/10.1007/s40314-018-0697-0>

Lisa Fischer, Sebastian Götschel, Martin Weiser (2018). **Lossy Data Compression Reduces Communication Time in Hybrid Time-Parallel Integrators**. *Comput. Vis. Sci.*, 19(1):19–30. <https://doi.org/10.1007/s00791-018-0293-2>

Fabio Furini, Emiliano Traversi, Pietro Belotti, Antonio Frangioni, Ambros Gleixner, Nick Gould, Leo Liberti, Andrea Lodi, Ruth Misener, Hans Mittelmann, Nikolaos V. Sahinidis, Stefan Vigerske, Angelika Wiegele (2018). **QPLIB: A Library of Quadratic Programming Instances**. *Mathematical Programming Computation*. <https://doi.org/10.1007/s12532-018-0147-4> (epub ahead of print 2018-09-22)

Gerald Gamrath, Christoph Schubert (2018). **Measuring the Impact of Branching Rules for Mixed-Integer Programming**. In *Operations Research Proceedings 2017*, 165–170. https://doi.org/10.1007/978-3-319-89920-6_23

Masoud Gholami, Florian Schintke, Thorsten Schütt (2018). **Checkpoint Scheduling for Shared Usage of Burst-Buffers in Supercomputers**. In *Proceedings of the 47th International Conference on Parallel Processing Companion; SRMPDS 2018; The 14th International Workshop on Scheduling and Resource Management for Parallel and Distributed Systems*, 44:1–44:10. <https://doi.org/10.1145/3229710.3229755>

Brady Gilg, Torsten Klug, Rosemarie Martienssen, Joseph Paat, Thomas Schlechte, Christof Schulz, Senan Seymen, Alexander Tesch (2018). **Conflict-Free Railway Track Assignment at Depots**. *Journal of Rail Transport Planning & Management*. <https://doi.org/10.1016/j.jrtpm.2017.12.004> (epub ahead of print 2018-02-03)

Ambros Gleixner, Stephan Maher, Benjamin Müller, João Pedro Pedroso (2018). **Price-and-Verify: a New Algorithm for Recursive Circle Packing Using Dantzig-Wolfe Decomposition**. *Annals of Operations Research*, 1–29. <https://doi.org/10.1007/s10479-018-3115-5>

Uwe Gotzes, Claudia Gotzes (2018). **The Jeep Problem: a Flexible MIP Formulation**. *Journal of Nepal Mathematical Society*, 1:11–17.

Niels Gregersen, Jakob Rosenkrantz de Lasson, Lars Hagedorn Frandsen, Oleksiy S. Kim, Olav Breinbjerg, Fengwen Wang, Ole Sigmund, Aliaksandra Ivinskaya, Andrei Lavrinenko, Philipp Gutsche, Sven Burger, Teppo Häyrynen, Jesper Mørk (2018). **Benchmarking State-of-the-Art Optical Simulation Methods for Analyzing Large Nanophotonic Structures**. In *XXVI International Workshop on Optical Wave & Waveguide Theory and Numerical Modelling, Proceedings*, 9.

Niels Gregersen, Jakob Rosenkrantz de Lasson, Lars Hagedorn Frandsen, Philipp Gutsche, Sven Burger, Oleksiy S. Kim, Olav Breinbjerg, Aliaksandra Ivinskaya, Fengwen Wang, Ole Sigmund, Teppo Häyrynen, Andrei V. Lavrinenko (2018). **Benchmarking State-of-the-Art Numerical Simulation Techniques for Analyzing Large Photonic Crystal Membrane Line Defect Cavities**. *Proc. SPIE*, 10672:106721C. <https://doi.org/10.1117/12.2304338>

Carl Martin Grewe, Gabriel le Roux, Sven-Kristofer Pilz, Stefan Zachow (2018). **Spotting the Details: the Various Facets of Facial Expressions**. In *IEEE International Conference on Automatic Face and Gesture Recognition*, 286–293. <https://doi.org/10.1109/FG.2018.00049>

Boris Grimm, Ralf Borndörfer, Christof Schulz, Steffen Weider (2018). **The Rolling Stock Rotation Planning Problem under Revenue Considerations**. In *Proceedings of the Rail Transport Demand Management Conference*.

Pooja Gupta, Sarah Peter, Markus Jung, Astrid Lewin, Georg Hemmrich-Stanisak, Andre Franke, Max von Kleist, Christof Schütte, Ralf Einspanier, Soroush Sharbati, Jennifer zur Brügge (2018). **Analysis of Long Non-coding RNA and mRNA Expression in Bovine Macrophages Brings up Novel 2 Aspects of Mycobacterium Avium Subspecies Paratuberculosis Infections**. *Scientific Reports in Nature*. (accepted for publication 2018-12-18)

Philipp Gutsche, Philipp-Immanuel Schneider, Sven Burger, Manuel Nieto-Vesperinas (2018). **Chiral Scatterers Designed by Bayesian Optimization**. *J. Phys.: Conf. Ser.*, 963:012004. <https://doi.org/10.1088/1742-6596/963/1/012004>

Philipp Gutsche, Manuel Nieto-Vesperinas (2018). **Optical Chirality of Time-Harmonic Wavefields for Classification of Scatterers**. *Sci. Rep.*, 8:9416. <https://doi.org/10.1038/s41598-018-27496-w>

Sebastian Götschel, Michael L. Minion (2018). **Parallel-in-Time for Parabolic Optimal Control Problems Using PFASST**. In *Domain Decomposition Methods in Science and Engineering XXIV*, 363–371. https://doi.org/10.1007/978-3-319-93873-8_34

Hassan Al Hajj, Mathieu Lamard, Pierre-Henri Conze, Soumali Roychowdhury, Xiaowei Hu, Gabija Marsalkaite, Odysseas Zisimopoulos, Muneer Ahmad Dedmari, Fenqiang Zhao, Jonas Prellberg, Manish Sahu, Adrian Galdran, Teresa Araujo, Duc My Vo, Chandan Panda, Navdeep Dahiya, Satoshi Kondo, Zhengbing Bian, Jonas Bialopetravicius, Chenghui Qiu, Sabrina Dill, Anirban Mukhopadhyay, Pedro Costa, Guilherme Aresta, Senthil Ramamurthy, Sang-Woong Lee, Aurelio Campilho, Stefan Zachow, Shunren Xia, Sailesh Conjeti, Jogundas Armitis, Pheng-Ann Heng, Arash Vahdat, Beatrice Cochener, Gwenole Quellec (2018). **CATARACTS: Challenge on Automatic Tool Annotation for cataRACT Surgery**. *Medical Image Analysis*, 52:24–41. <https://doi.org/10.1016/j.media.2018.11.008> (accepted for publication 2018-11-15)

Martin Hammerschmidt, Jakub Pabisiak, Ana Cristina Perez-Gea, Nicolas Julian, Sean Hon, Sven Burger (2018). **Efficient Finite-Element-Based Numerical Modelling of Large Sub-wavelength Patterned Optical Structures**. *Proc. SPIE*, 10542:105421G. <https://doi.org/10.1117/12.2292438>

Martin Hammerschmidt, Philipp-Immanuel Schneider, Xavier Garcia Santiago, Lin Zschiedrich, Martin Weiser, Sven Burger (2018). **Solving Inverse Problems Appearing in Design and Metrology of Diffractive Optical Elements by Using Bayesian Optimization**. *Proc. SPIE*, 10694:1069407. <https://doi.org/10.1117/12.2315468>

Gregor Hendel, Matthias Miltenberger, Jakob Witzig (2018). **Adaptive Algorithmic Behavior for Solving Mixed Integer Programs Using Bandit Algorithms**. In *OR 2018: International Conference on Operations Research, Operations Research 2018 Proceedings*. (accepted for publication 2018-11-16)

Felix Hennings (2018). **Benefits and Limitations of Simplified Transient Gas Flow Formulations.** In *Operations Research Proceedings 2017*, Operations Research Proceedings, 231-237. https://doi.org/https://doi.org/10.1007/978-3-319-89920-6_32

Benjamin Hiller, Thorsten Koch, Lars Schewe, Robert Schwarz, Jonas Schweiger (2018). **A System to Evaluate Gas Network Capacities: Concepts and Implementation.** *European Journal of Operational Research*, 270(3):797-808.

Rene Hoffmann, Robert Lemanis, Lena Wulff, Stefan Zachow, Alexander Lukeneder, Christian Klug, Helmut Keupp (2018). **Traumatic Events in the Life of the Deep-Sea Cephalopod Mollusc, the Coleoid Spirula Spirula.** *ScienceDirect: Deep Sea Research Part I - Oceanographic Research*. <https://doi.org/10.1016/j.dsr.2018.10.007> (epub ahead of print 2018-10-25)

Kai Hoppmann, Robert Schwarz (2018). **Finding Maximum Minimum Cost Flows to Evaluate Gas Network Capacities.** In *Operations Research Proceedings 2017*, 339-346. https://doi.org/10.1007/978-3-319-89920-6_46

Theresa Höhne, Peter Schnauber, Sven Rodt, Stephan Reitzenstein, Sven Burger (2018). **Numerical Investigation of Light Emission from Quantum Dots Embedded into On-Chip, Low Index Contrast Optical Waveguides.** *Phys. Status Solidi B*. (accepted for publication 2018-12-17)

Theresa Höhne, Lin Zschiedrich, Nasibeh Haghighi, James A. Lott, Sven Burger (2018). **Validation of Quasi-normal Modes and of Constant-Flux Modes for Computing Fundamental Resonances of VCSELs.** *Proc. SPIE*, 10682:106821U. <https://doi.org/10.1117/12.2307200>

Eugene Jennifer Jin, Ferdi Ridvan Kiral, Mehmet Neset Ozel, Lara Sophie Burchardt, Marc Osterland, Daniel Epstein, Heike Wolfenber, Steffen Prohaska, Peter Robin Hiesinger (2018). **Live Observation of Two Parallel Membrane Degradation Pathways at Axon Terminals.** *Current Biology*, 28(7):1027-1038.e4. <https://doi.org/10.1016/j.cub.2018.02.032>

Leonardo Agudo Jácome, Hans-Christian Hege, Olaf Paetsch, Kai Pöthkow (2018). **Three-Dimensional Reconstruction and Quantification of Dislocation Substructures from Transmission Electron Microscopy Stereo Pairs.** *Ultramicroscopy*, 195:157-170. <https://doi.org/10.1016/j.ultramic.2018.08.015>

Klaus Jäger, Phillip Manley, David Eisenhauer, Philipp Tockhorn, Johannes Sutter, Martin Hammerschmidt, Sven Burger, Steve Albrecht, Christiane Becker (2018). **Light Management with Sinusoidal Nanotextures.** In *Light, Energy and the Environment 2018*, OSA Technical Digest, page OT5C.5. <https://doi.org/10.1364/OSE.2018.OT5C.5>

Klaus Jäger, Phillip Manley, Duote Chen, Philipp Tockhorn, David Eisenhauer, Grit Köppel, Martin Hammerschmidt, Sven Burger, Steve Albrecht, Christiane Becker (2018). **Antireflective Nanotextures for Monolithic Perovskite-Silicon Tandem Solar Cells.** *Proc. SPIE*, 10688:106880A. <https://doi.org/10.1117/12.2306142>

Klaus Jäger, Grit Köppel, Martin Hammerschmidt, Sven Burger, Christiane Becker (2018). **On Accurate Simulations of Thin-Film Solar Cells with a Thick Glass Superstrate.** *Opt. Express*, 26:A99. <https://doi.org/10.1364/OE.26.000A99>

Günter Kewes, Felix Binkowski, Sven Burger, Lin Zschiedrich, Oliver Benson (2018). **Heuristic Modeling of Strong Coupling in Plasmonic Resonators.** *ACS Photonics*, 5:4089. <https://doi.org/10.1021/acsp Photonics.8b00766>

Marco Klindt (2018). **Implications of Perceived Preservation Levels.** In *Proceedings of iPres - 15th International Conference on Digital Preservation/Boston, MA, USA/September 24-27, 2018*, Volume 15 of *iPres Proceedings*. <https://doi.org/10.17605/OSF.IO/M6QA5>

Torsten Klug (2018). **Freight Train Routing.** In *Handbook of Optimization in the Railway Industry*, Volume 268 of *International Series in Operations Research & Management Science*, Pages 73-92. Springer International Publishing. <https://doi.org/10.1007/978-3-319-72153-8>

Stefan Klus, Feliks Nüske, Peter Koltai, Hao Wu, Ioannis Kevrekidis, Christof Schütte, Frank Noé (2018). **Data-Driven Model Reduction and Transfer Operator Approximation.** *Journal of Nonlinear Science*, 28(3):985-1010. <https://doi.org/10.1007/s00332-017-9437-7>

Stefan Klus, Andreas Bittracher, Ingmar Schuster, Christof Schütte (2018). **A Kernel-Based Approach to Molecular Conformation Analysis.** *Journal of Chemical Physics*. (accepted for publication 2018-12-04)

S. Klus, P. Gelß, S. Peitz, Ch. Schütte (2018). **Tensor-Based Dynamic Mode Decomposition.** *Nonlinearity*, 31(7). <https://doi.org/10.1088/1361-6544/aabc8f> (accepted for publication 2018-04-09)

David Knötel, Carola Becker, Gerhard Scholtz, Daniel Baum (2018). **Global and Local Mesh Morphing for Complex Biological Objects from μ CT Data.** In *Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*, 179-183. <https://doi.org/10.2312/vcbm.20181243>

Peter Koltai, Hao Wu, Frank Noé, Christof Schütte (2018). **Optimal Data-Driven Estimation of Generalized Markov State Models for Non-equilibrium Dynamics.** *Computation*, 6(1). <https://doi.org/10.3390/computation6010022>

Péter Koltai, Christof Schütte (2018). **A Multiscale Perturbation Expansion Approach for Markov State Modeling of Nonstationary Molecular Dynamics.** *SIAM J. Multiscale Model. Simul.*, 16(4):1455–1485. <https://doi.org/10.1137/17M1146403> (accepted for publication 2018-08-15)

F. Kork, C. Spies, T. Conrad, B. Weiss, T. Roenneberg, K.-D. Wernecke, F. Balzer (2018). **Associations of Postoperative Mortality with the Time of Day, Week and Year.** *Anaesthesia*. (accepted for publication 2018-02-23)

*Tobias Kramer, Matthias Noack, Daniel Baum, Hans-Christian Hege, Eric J. Heller (2018). **Dust and Gas Emission from Cometary Nuclei: the Case of Comet 67P/Churyumov–Gerasimenko.** *Advances in Physics: X*, 3(1):1404436. <https://doi.org/10.1080/23746149.2017.1404436> (Joint publication: Visual Data Analysis, Distributed Algorithms and Supercomputing)

Tobias Kramer, Matthias Noack, Alexander Reinefeld, Mirta Rodríguez, Yaroslav Zelinsky (2018). **Efficient Calculation of Open Quantum System Dynamics and Time-Resolved Spectroscopy with Distributed Memory HEOM (DM-HEOM).** *Journal of Computational Chemistry*, 39(22):1779–1794. <https://doi.org/10.1002/jcc.25354>

Tobias Kramer, Matthias Noack, Jeffrey Reimers, Alexander Reinefeld, Mirta Rodríguez, Shiwei Yin (2018). **Energy Flow in the Photosystem I Supercomplex: Comparison of Approximative Theories with DM-HEOM.** *Chemical Physics*, 515:262–271. <https://doi.org/10.1016/j.chemphys.2018.05.028>

Martin Krämer, Marta B Maggioni, Christoph von Tycowicz, Nick Brisson, Stefan Zachow, Georg N Duda, Jürgen R Reichenbach (2018). **Ultra-short Echo-Time (UTE) Imaging of the Knee with Curved Surface Reconstruction-Based Extraction of the Patellar Tendon.** In *ISMRM (International Society for Magnetic Resonance in Medicine), 26th Annual Meeting 2018, Paris, France*. (accepted for publication 2018-02-03)

Alexander Lange, Robert Schwieger, Julia Plöntzke, Stefan Schäfer, Susanna Röblitz (2018). **Follicular Competition in Cows: the Selection of Dominant Follicles as a Synergistic Effect.** *Journal of Mathematical Biology*. <https://doi.org/https://doi.org/10.1007/s00285-018-1284-0>

Jakob Rosenkrantz de Lasso, Lars Hagedorn Frandsen, Philipp Gutsche, Sven Burger, Oleksiy S. Kim, Olav Breinbjerg, Aliaksandra Ivanskaya, Fengweng Wang, Ole Sigmund, Teppo Häyrynen, Andrei V. Lavrinenko, Jesper Mork, Niels Gregersen (2018). **Benchmarking Five Numerical Simulation Techniques for Computing Resonance Wavelengths and Quality Factors in Photonic Crystal Membrane Line Defect Cavities.** *Opt. Express*, 26:11366. <https://doi.org/10.1364/OE.26.011366>

Han Cheng Lie, T. J. Sullivan (2018). **Erratum: Equivalence of Weak and Strong Modes of Measures on Topological Vector Spaces (2018 Inverse Problems 34 115013).** *Inverse Problems*, 34(12):129601. <https://doi.org/10.1088/1361-6420/aae55b>

Han Cheng Lie, T. J. Sullivan (2018). **Quasi-invariance of Countable Products of Cauchy Measures Under Non-unitary Dilations.** *Electronic Communications in Probability*, 23(8):1–6. <https://doi.org/10.1214/18-ECP113>

Han Cheng Lie, T. J. Sullivan (2018). **Equivalence of Weak and Strong Modes of Measures on Topological Vector Spaces.** *Inverse Problems*, 34(11):115013. <https://doi.org/10.1088/1361-6420/aaef2>

Han Cheng Lie, T. J. Sullivan, Aretha Teckentrup (2018). **Random Forward Models and Log-Likelihoods in Bayesian Inverse Problems.** *SIAM/ASA Journal on Uncertainty Quantification*, Volume 6 of *ZIB-Report 18-03*, Pages 1600–1629. <https://doi.org/10.1137/18M1166523>

Norbert Lindow, Daniel Baum, Morgan Leborgne, Hans-Christian Hege (2019). **Interactive Visualization of RNA and DNA Structures.** *IEEE Transactions on Visualization and Computer Graphics*, 25(1):967–976. <https://doi.org/10.1109/TVCG.2018.2864507> (epub ahead of print 2018-10-17)

Norbert Lindow, Daniel Baum, Hans-Christian Hege (2018). **Atomic Accessibility Radii for Molecular Dynamics Analysis.** In *Workshop on Molecular Graphics and Visual Analysis of Molecular Data*. <https://doi.org/10.2312/molva.20181101>

Philipp Lorenz-Spreen, Frederik Wolf, Jonas Braun, Gourab Ghoshal, Natasa Djurdjevac Conrad (2018). **Tracking Online Topics over Time: Understanding Dynamic Hashtag Communities.** *Computational Social Networks*, 5(9). <https://doi.org/10.1186/s40649-018-0058-6> (accepted for publication 2018-09-28)

Matthias Läter, Tobias Kramer, Martin Rubin, Kathrin Altwegg (2019). **Surface Localization of Gas Sources on Comet 67P/Churyumov–Gerasimenko Based on DFMS/COPS Data.** *Monthly Notices of the Royal Astronomical Society*, 483:852–861. <https://doi.org/10.1093/mnras/sty3103> (epub ahead of print 2018-11-15)

Franz Löchner, Anna Fedotova, Sheng Liu, Gordon Keeler, Gregory Peake, Sina Saravi, Maxim Shcherbakov, Sven Burger, Andrey Fedyanin, Igal Brener, Thomas Pertsch, Frank Setzpfandt, Isabelle Staude (2018). **Polarization-Dependent Second Harmonic Diffraction from Resonant GaAs Metasurfaces.** *ACS Photonics*, 5:1786. <https://doi.org/10.1021/acsp Photonics.7b01533>

X. Ma, T. Conrad, M. Alchikh, J. Reiche, B. Schweiger, B. Rath (2018). Can We Distinguish Respiratory Viral Infections Based on Clinical Features? A Prospective Pediatric Cohort Compared to Systematic Literature Review. *Medical Virology*, 28(5):1-12. <https://doi.org/10.1002/rmv.1997> (accepted for publication 2018-06-14)

Stephen J. Maher, Guy Desaulniers, Francois Soumis (2018). The Daily Tail Assignment Problem under Operational Uncertainty Using Look-Ahead Maintenance Constraints. *European Journal of Operational Research*, 264(2):534-547. <https://doi.org/10.1016/j.ejor.2017.06.041>

Radu Malureanu, Jakob Rosenkrantz de Lasson, Lars Hagedorn Frandsen, Philipp Gutsche, Sven Burger, Oleksiy S. Kim, Olav Breinbjerg, Aliaksandra Ivinskaya, Fengwen Wang, Ole Sigmund, Teppo Häyrynen, Andrei Lavrinenko, Jesper Mørk, Niels Gregersen (2018). Which Computational Methods Are Good for Analyzing Large Photonic Crystal Membrane Cavities? In *20th International Conference on Transparent Optical Networks (ICTON)*. <https://doi.org/10.1109/ICTON.2018.8473951>

Toni Mancini, Federico Mari, Annalisa Massini, Igor Melatti, Ivano Salvo, Stefano Sinisi, Enrico Tronci, Rainald Ehrig, Susanna Röblitz, Brigitte Leeners (2018). Computing Personalised Treatments through In Silico Clinical Trials. A Case Study on Downregulation in Assisted Reproduction. In *Experimental Evaluation of Algorithms for Solving Problems with Combinatorial Explosion, Workshop of the Federated Logic Conference 2018 (FLoC 2018)*, Oxford, 1-16. <https://doi.org/10.29007/g864>

Phillip Manley, Klaus Jäger, Philipp Tockhorn, Sven Burger, Steve Albrecht, Christiane Becker (2018). Optimization of Nanostructured High Efficiency Perovskite/c-Si Tandem Solar Cells via Numerical Simulation. In *Proc. SPIE*, 10731:107310B. <https://doi.org/10.1117/12.2321240>

Phillip Manley, Fatwa F. Abdi, Sean P. Berglund, A. T. M. Nazmul Islam, Sven Burger, Roel van de Krol, Martina Schmid (2018). Absorption Enhancement for Ultrathin Solar Fuel Devices with Plasmonic Gratings. *ACS Appl. Energy Mater.*, 1:5810. <https://doi.org/10.1021/acsaem.8b01070>

Phillip J. Manley, Min Song, Sven Burger, Martina Schmid (2018). Efficient Determination of Bespoke Optically Active Nanoparticle Distributions. *J. Opt.*, 20:085003. <https://doi.org/10.1088/2040-8986/aad114>

Forough Marzban, Tim Conrad, Pouria Marban, Sahar Sodoudi (2018). Estimation of the Near-Surface Air Temperature during the Day and Nighttime from MODIS in Berlin, Germany. *International Journal of Advanced Remote Sensing and GIS*, 7(1):2478-2517. <https://doi.org/10.23953/cloud.ijarsg.337>

Matthias Miltenberger, Ted Ralphs, Daniel Steffy (2018). Exploring the Numerics of Branch-and-Cut for Mixed Integer Linear Optimization. In *Operations Research Proceedings 2017*, 151-157. <https://doi.org/10.1007/978-3-319-89920-6>

Victor Mireles, T. Conrad (2018). Reusable Building Blocks in Biological Systems. *Journal of the Royal Society Interface*, 15(149):1-10. <https://doi.org/10.1098/rsif.2018.0595> (accepted for publication 2018-11-23)

Matthias Noack, Alexander Reinefeld, Tobias Kramer, Thomas Steinke (2018). DM-HEOM: A Portable and Scalable Solver-Framework for the Hierarchical Equations of Motion. In *2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, 19th IEEE Int. Workshop on Parallel and Distributed Scientific and Engineering Computing (PDSEC 2018), 947-956. <https://doi.org/10.1109/IPDPSW.2018.00149>

Steffen Oeltze-Jaffra, Monique Meuschke, Matthias Neugebauer, Sylvia Saalfeld, Kai Lawonn, Gabor Janiga, Hans-Christian Hege, Stefan Zachow, Bernhard Preim (2018). Generation and Visual Exploration of Medical Flow Data: Survey, Research Trends, and Future Challenges. *Computer Graphics Forum*. <https://doi.org/10.1111/cgf.13394>

Wolfgang Peters-Kottig (2018). Anwendung von OPUS für Forschungsdaten - das Beispiel eines außeruniversitären Forschungsinstituts. *o-bib. Das offene Bibliotheksjournal*, 5(2):79-91. <https://doi.org/10.5282/o-bib/2018h2s79-91>

Julia Plöntzke, Mascha Berg, Susanna Röblitz (2018). A Mathematical Modelling Approach to Dynamic Networks: Potassium Homeostasis and Glucose-Insulin in Dairy Cows. *Journal of Dairy Science* 101 Suppl. 2, 226-227

Jannes Quer, Luca Donati, Bettina Keller, Marcus Weber (2018). An Automatic Adaptive Importance Sampling Algorithm for Molecular Dynamics in Reaction Coordinates. *SIAM Journal on Scientific Computing*, 40(2):A653-A670.

Ted Ralphs, Yuji Shinano, Timo Berthold, Thorsten Koch (2018). Parallel Solvers for Mixed Integer Linear Optimization. In Youssef Hamadi, editor, *Handbook of Parallel Constraint Reasoning*, 283-336. https://doi.org/10.1007/978-3-319-63516-3_8

Stefanie Redemann, Ina Lantzsch, Norbert Lindow, Steffen Prohaska, Martin Srayko, Thomas Müller-Reichert (2018). A Switch in Microtubule Orientation during *C. elegans* Meiosis. *Current Biology*. <https://doi.org/10.1016/j.cub.2018.07.012>

Daniel Rehfeldt, Thorsten Koch (2018). SCIP-Jack - a Solver for STP and Variants with Parallelization Extensions: An Update. In *Operations Research Proceedings 2017*, 191-196.

- Daniel Rehfeldt, Thorsten Koch (2018). **Transformations for the Prize-Collecting Steiner Tree Problem and the Maximum-Weight Connected Subgraph Problem to SAP.** *Journal of Computational Mathematics*, 36(3):459–468. <https://doi.org/10.1016/j.jco.2018.05.022>
- Marco Reidelbach, Marcus Weber, Petra Imhof (2018). **Prediction of Perturbed Proton Transfer Networks.** *PLoS ONE*, 13(12):e0207718–e0207718. <https://doi.org/10.1371/journal.pone.0207718>
- Malte Renken, Amin Ahmadi, Ralf Borndörfer, Guvenc Sahin, Thomas Schlechte (2018). **Demand-Driven Line Planning with Selfish Routing.** In *Operations Research Proceedings 2017*, 687–692. https://doi.org/10.1007/978-3-319-89920-6_91
- Bernhard Reuter, Marcus Weber, Konstantin Fackeldey, Susanna Röblitz, Martin E. Garcia (2018). **Generalized Markov State Modeling Method for Nonequilibrium Biomolecular Dynamics: Exemplified on Amyloid β Conformational Dynamics Driven by an Oscillating Electric Field.** *Journal of Chemical Theory and Computation*, 14(7):3579–3594. <https://doi.org/10.1021/acs.jctc.8b00079>
- Markus Reuther, Thomas Schlechte (2018). **Optimization of Rolling Stock Rotations.** In *Handbook of Optimization in the Railway Industry*, Volume 268 of *International Series in Operations Research & Management Science*, Pages 213–241, Springer International Publishing. <https://doi.org/10.1007/978-3-319-72153-8>
- Guillaume Sagnol, Daniel Schmidt genannt Waldschmidt, Alexander Tesch (2018). **The Price of Fixed Assignments in Stochastic Extensible Bin Packing.** In *WAOA 2018: Approximation and Online Algorithms*, Volume 11312 of *Lecture Notes in Computer Science*, Pages 327–347. https://doi.org/10.1007/978-3-030-04693-4_20 (accepted for publication 2018-07-23)
- Guillaume Sagnol, Christoph Barner, Ralf Borndörfer, Mickaël Grima, Mathees Seeling, Claudia Spies, Klaus Wernecke (2018). **Robust Allocation of Operating Rooms: a Cutting Plane Approach to Handle Lognormal Case Durations.** *European Journal of Operational Research*, 271(2):420–435. <https://doi.org/10.1016/j.ejor.2018.05.022>
- Guillaume Sagnol, Marco Blanco, Thibaut Sauvage (2018). **The Cone of Flow Matrices: Approximation Hierarchies and Applications.** *Networks*, 72(1):128–150. <https://doi.org/10.1002/net.21820>
- Guillaume Sagnol, Marco Blanco, Thibaut Sauvage (2018). **Approximation Hierarchies for the Cone of Flow Matrices.** In *INOC 2017 – 8th International Network Optimization Conference*, Volume 64 of *Electronic Notes in Discrete Mathematics*, Pages 275–284. <https://doi.org/10.1016/j.endm.2018.02.002>
- Xavier Garcia Santiago, Martin Hammerschmidt, Sven Burger, Carsten Rockstuhl, Ivan Fernandez-Corbato, Lin Zschiedrich (2019). **Decomposition of Scattered Electromagnetic Fields into Vector Spherical Wave Functions on Surfaces with General Shapes.** *Phys. Rev. B*, 99:045406. <https://doi.org/10.1103/PhysRevB.99.045406> (accepted for publication 2018-12-12)
- Xavier Garcia Santiago, Philipp-Immanuel Schneider, Carsten Rockstuhl, Sven Burger (2018). **Shape Design of a Reflecting Surface Using Bayesian Optimization.** *J. Phys.: Conf. Ser.*, 963:012003. <https://doi.org/10.1088/1742-6596/963/1/012003>
- Stanley Schade, Thomas Schlechte, Jakob Witzig (2018). **Structure-Based Decomposition for Pattern-Detection for Railway Timetables.** In *Operations Research Proceedings 2017*, 715–721. https://doi.org/10.1007/978-3-319-89920-6_95
- Robert Schmidtke, Kenny Erleben (2018). **Chunked Bounding Volume Hierarchies for Fast Digital Prototyping Using Volumetric Meshes.** *IEEE Transactions on Visualization and Computer Graphics*, 24(12):3044–3057. <https://doi.org/10.1109/TVCG.2017.2784441>
- Robert Schmidtke, Florian Schintke, Thorsten Schütt (2018). **From Application to Disk: Tracing I/O Through the Big Data Stack.** In *High Performance Computing ISC High Performance 2018 International Workshops, Frankfurt/Main, Germany, June 24–28, 2018, Revised Selected Papers, Workshop on Performance and Scalability of Storage Systems (WOPSSS)*, 89–102. https://doi.org/10.1007/978-3-030-02465-9_6
- Peter Schnauber, Johannes Schall, Samir Bounouar, Theresa Höhne, Suk-In Park, Geun-Hwan Ryu, Tobias Heindel, Sven Burger, Jin-Dong Song, Sven Rodt, Stephan Reitzenstein (2018). **Deterministic Integration of Quantum Dots into On-chip Multi-mode Interference Beamsplitters Using In Situ Electron Beam Lithography.** *Nano Lett.*, 18:2336. <https://doi.org/10.1021/acs.nanolett.7b05218>
- Philipp-Immanuel Schneider, Xavier Garcia Santiago, Felix Binkowski, Philipp Gutsche, Theresa Höhne, Martin Hammerschmidt, Lin Zschiedrich, Sven Burger (2018). **Light Management for Engineering Luminescence in Nanoscale Environments by Numerical Optimization.** In *The Electrochemical Society, Meeting Abstracts*, 16:1164.
- Philipp-Immanuel Schneider, Lin Zschiedrich, Xavier Garcia Santiago, Sven Burger (2018). **Optimization of Quantum Optical Devices with Machine Learning.** In *XXVI International Workshop on Optical Wave & Waveguide Theory and Numerical Modelling, Proceedings*, 6.

Philipp-Immanuel Schneider, Nicole Srocka, Sven Rodt, Lin Zschiedrich, Stephan Reitzenstein, Sven Burger (2018). Numerical Optimization of the Extraction Efficiency of a Quantum-Dot Based Single-Photon Emitter into a Single-Mode Fiber. *Opt. Express*, 26:8479. <https://doi.org/10.1364/OE.26.008479>

Katharina Schrade, Jessica Tröger, Adeep Eldashan, Kerstin Zühlke, Kamal R. Abdul Azees, Jonathan M. Elkins, Martin Neuenchwander, Andreas Oder, Mohamed Elkewedi, Sarah Jaksch, Karsten Andrae, Jinliang Li, Jaoa Fernandes, Paul Markus Müller, Stephan Grunwald, Stephen F. Marino, Tanja Vukicevic, Jenny Eichhorst, Burkhard Wiesner, Marcus Weber, Michael Kapiloff, Oliver Rocks, Oliver Daumke, Thomas Wieland, Stefan Knapp, Jens Peter von Kries, Enno Klussmann (2018). An AKAP-Lbc-RhoA Interaction Inhibitor Promotes the Translocation of Aquaporin-2 to the Plasma Membrane of Renal Collecting Duct Principal Cells. *PLoS ONE*, 13(1):e0191423–e0191423. <https://doi.org/10.1371/journal.pone.0191423>

Stephan Schwartz, Ralf Borndörfer, Leonardo Balestrieri (2018). On Finding Subpaths with High Demand. In *Operations Research Proceedings 2017*, 355–360.

Jonas Schweiger, Frauke Liers (2018). A Decomposition Approach for Optimal Gas Network Extension with a Finite Set of Demand Scenarios. *Optimization and Engineering*, 19(2):297–326.

Robert Schwieger, Heike Siebert, Susanna Röblitz (2018). Correspondence of Trap Spaces in Different Models of Bioregulatory Networks. *SIAM Journal on Applied Dynamical Systems*, 17(2):1742–1765. <https://doi.org/10.1137/17M1155302>

Feda Seblany, Ulrike Homberg, Eric Vincens, Paul Winkler, Karl Josef Witt (2018). Merging Criteria for Defining Pores and Constrictions in Numerical Packing of Spheres. *Granular Matter*, 20(37). <https://doi.org/10.1007/s10035-018-0808-z>

Felix Seibert, Mathias Peters, Florian Schintke (2018). Improving I/O Performance through Colocating Interrelated Input Data and Near-Optimal Load Balancing. *Proceedings of the IPDPSW; Fourth IEEE International Workshop on High-Performance Big Data, Deep Learning, and Cloud Computing (HPBDC)*, 2018:448–457. <https://doi.org/10.1109/IPDPSW.2018.00081>

Yuji Shinano, Timo Berthold, Stefan Heinz (2018). ParaXpress: An Experimental Extension of the FICO Xpress-Optimizer to Solve Hard MIPs on Supercomputers. *Optimization Methods & Software*. <https://doi.org/10.1080/10556788.2018.1428602> (accepted for publication 2018-01-13)

Yuji Shinano, Stefan Heinz, Stefan Vigerske, Michael Winkler (2018). FiberSCIP – a Shared Memory Parallelization of SCIP. *INFORMS Journal on Computing*, 30(1):11–30. <https://doi.org/10.1287/ijoc.2017.0762>

Maksim Shpak, Sven Burger, Ville Byman, Kimmo Saastamoinen, Mertsu Haapalainen, Antti Lassila (2018). Online Measurement of Optical Fibre Geometry During Manufacturing. *Proc. SPIE*, 10683:1068318. <https://doi.org/10.1117/12.2314762>

Viola Spahn, Giovanna Del Vecchio, Antonio Rodriguez-Gaztelumendi, Julia Temp, Dominika Labuz, Michael Klöner, Marco Reidelbach, Halina Machelska, Marcus Weber, Christoph Stein (2018). Opioid Receptor Signaling, Analgesic and Side Effects Induced by a Computationally Designed pH-Dependent Agonist. *Scientific Reports*, 8:8965.

Nicole Srocka, Anna Musiał, Philipp-Immanuel Schneider, Paweł Mrowiński, Paweł Holewa, Sven Burger, David Quandt, André Strittmatter, Sven Rodt, Stephan Reitzenstein, Grzegorz Sęk (2018). Enhanced Photon-Extraction Efficiency from InGaAs/GaAs Quantum Dots in Deterministic Photonic Structures at 1.3 μm Fabricated by In Situ Electron-Beam Lithography. *AIP Adv.*, 8:085205. <https://doi.org/10.1063/1.5038137>

Tom Streubel, Christian Strohm, Philipp Trunschke, Caren Tischendorf (2018). Generic Construction and Efficient Evaluation of Network DAEs and Their Derivatives in the Context of Gas Networks. *Operations Research Proceedings 2017*, 627–632. https://doi.org/10.1007/978-3-319-89920-6_83

Alexander Tack, Anirban Mukhopadhyay, Stefan Zachow (2018). Knee Menisci Segmentation Using Convolutional Neural Networks: Data from the Osteoarthritis Initiative. *Osteoarthritis and Cartilage*, 26(5):680–688. <https://doi.org/10.1016/j.joca.2018.02.907>

Alexander Tack, Stefan Zachow (2018). Accurate Automated Volumetry of Cartilage of the Knee Using Convolutional Neural Networks: Data from the Osteoarthritis Initiative. In *IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*. (accepted for publication 2018-12-18)

Alexander Tesch (2018). Improving Energetic Propagations for Cumulative Scheduling. In *Principles and Practice of Constraint Programming (CP 2018)*. (accepted for publication 2018-06-01)

Onur Teymur, Han Cheng Lie, T. J. Sullivan, Ben Calderhead (2018). Implicit Probabilistic Integrators for ODEs. In *Advances in Neural Information Processing Systems 31 (NIPS 2018)*.

Nicholas J. Tierney, Antonietta Mira, J. Jost Reinhold, Martin Weiser, Roman Burkart, Claudio Benvenuti, Angelo Auricchio (2018). Novel Relocation Methods for Automatic External Defibrillator Improve Out-of-Hospital Cardiac Arrest Coverage under Limited Resources. *Resuscitation*, 83–89. <https://doi.org/10.1016/j.resuscitation.2018.01.055>

Jürgen Titschack, Daniel Baum, Kei Matsuyama, Karin Boos, Claudia Färber, Wolf-Achim Kahl, Karsten Ehrig, Dietmar Meinel, Carmen Soriano, Stuart R. Stock (2018). Ambient Occlusion – a Powerful Algorithm to Segment Shell and Skeletal Intrapores in Computed Tomography Data. *Computers and Geosciences*, 115:75–87. <https://doi.org/10.1016/j.cageo.2018.03.007> (epub ahead of print 2018-03-10)

Christoph von Tycowicz, Felix Ambellan, Anirban Mukhopadhyay, Stefan Zachow (2018). An Efficient Riemannian Statistical Shape Model Using Differential Coordinates. *Medical Image Analysis*, 43(1):1–9. <https://doi.org/10.1016/j.media.2017.09.004>

Brygg Ullmer, Philip Paar, Liviu Coconu, Jing Liu, Miriam Konkel, Hans-Christian Hege (2018). An 1834 Mediterranean Garden in Berlin – Engaged from 2004, 2018, 2032, and 2202. In *PerDis '18 – 7th ACM International Symposium on Pervasive Displays, Munich, Germany, June 6–8, 2018*, No. 12. <https://doi.org/10.1145/3205873.3205894>

Stefan Vigerske, Ambros Gleixner (2018). SCIP: Global Optimization of Mixed-Integer Nonlinear Programs in a Branch-and-Cut Framework. *Optimization Methods & Software*, 33(3):563–593. <https://doi.org/10.1080/10556788.2017.1335312>

Sabine Wagner, Carlos Zapata, Wei Wan, Kornelia Gawlitza, Marcus Weber, Knut Rurack (2018). Role of Counterions in Molarly Imprinted Polymers for Anionic Species. *Langmuir*, 34(23):6963–6975. <https://doi.org/10.1021/acs.langmuir.8b00500>

Marcus Weber (2018). Implications of PCCA+ in Molecular Simulation. *Computation*, 6(1):20. <https://doi.org/10.3390/computation6010020>

Tobias Weber, Sebastian Sager, Ambros Gleixner (2018). Solving Quadratic Programs to High Precision Using Scaled Iterative Refinement. *Mathematical Programming Computation*. (accepted for publication 2018-12-22)

Marcus Weber (2018). Transformationsprodukte im Klärwerk: Mathematische Ansätze der Bewertung. *KA Korrespondenz Abwasser, Abfall*. (accepted for publication 2018-06-04)

*Martin Weiser, Bodo Erdmann, Sebastian Schenk, Holger Muggenthaler, Michael Hubig, Gita Mall, Stefan Zachow (2018). Uncertainty in Temperature-Based Determination of Time of Death. *Heat and Mass Transfer*, 54(9):2815–2826. <https://doi.org/10.1007/s00231-018-2324-4> (Joint publication: Numerical Mathematics, Visual Data Analysis)

Martin Weiser, Yvonne Freytag, Bodo Erdmann, Michael Hubig, Gita Mall (2018). Optimal Design of Experiments for Estimating the Time of Death in Forensic Medicine. *Inverse Problems*, 34(12):125005. <https://doi.org/10.1088/1361-6420/aae7a5>

Martin Weiser, Sunayana Ghosh (2018). Theoretically Optimal Inexact SDC Methods. *Communications in Applied Mathematics and Computational Science*, 53–86. <https://doi.org/10.2140/camcos.2018.13.53>

Daniel Werdehausen, Isabelle Staude, Sven Burger, Jörg Petschulat, Toralf Scharf, Thomas Pertsch, Manuel Decker (2018). Design Rules for Customizable Optical Materials Based on Nanocomposites. *Opt. Mater. Express*, 8:3456. <https://doi.org/10.1364/OME.8.003456>

Daniel Werdehausen, Isabelle Staude, Sven Burger, Jörg Petschulat, Toralf Scharf, Thomas Pertsch, Manuel Decker (2018). Using Effective Medium Theories to Design Tailored Nanocomposite Materials for Optical Systems. In *Proc. SPIE*, 10745:107450H. <https://doi.org/10.1117/12.2320525>

Jon Wilson, Philipp Gutsche, Sven Herrmann, Sven Burger, Kevin McPeak (2019). Correlation of Circular Differential Optical Absorption with Geometric Chirality in Plasmonic Meta-Atoms. *Opt. Express*, 27:5097. <https://doi.org/10.1364/OE.27.005097> (accepted for publication 2018-12-26)

*Jakob Witzig, Isabel Beckenbach, Leon Eifler, Konstantin Fackeldey, Ambros Gleixner, Andreas Grever, Marcus Weber (2018). Mixed-Integer Programming for Cycle Detection in Non-reversible Markov Processes. *Multiscale Modeling and Simulation*, 16(1):248–265. <https://doi.org/10.1137/16M1091162> (Joint publication: Mathematical Optimization, Numerical Mathematics)

Ryohei Yokoyama, Yuji Shinano, Yuji Wakayama, Tetsuya Wakui (2018). Optimal Design of a Gas Turbine Cogeneration Plant by a Hierarchical Optimization Method with Parallel Computing. In *ASME Turbo Expo 2018: Turbomachinery Technical Conference and Exposition*, 3.

NOT PEER-REVIEWED (7)

Ryohei Yokoyama, Yuji Shinano, Yuki Wakayama, Tetsuya Wakui (2018). **Model Reduction by Time Aggregation for Optimal Design of Energy Supply Systems by an MILP Hierarchical Branch and Bound Method.** In *Proceedings of the 31st International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy Systems (ECOS 2018)*.

Wei Zhang, Carsten Hartmann, Max von Kleist (2018). **Optimal Control of Markov Jump Processes: Asymptotic Analysis, Algorithms and Applications to the Modeling of Chemical Reaction Systems.** *Communications in Mathematical Sciences*, 293–331. <https://doi.org/10.4310/CMS.2018.v16.n2.a1>

Lin Zschiedrich, Felix Binkowski, Niko Nikolay, Oliver Benson, Günter Kewes, Sven Burger (2018). **Riesz-Projection-Based Theory of Light-Matter Interaction in Dispersive Nanoresonators.** *Phys. Rev. A*, 98:043806. <https://doi.org/10.1103/PhysRevA.98.043806>

Felix Binkowski, Lin Zschiedrich, Sven Burger (2018). **A Riesz-Projection-Based Method for Nonlinear Eigenvalue Problems.** *arXiv*. <https://arxiv.org/abs/1811.11624>

Anne Dennstedt, Fabrice Gaslain, Marion Bartsch, Vincent Guipont, Vincent Maurel (2018). **Crack Morphology in a Columnar Thermal Barrier Coating System.** In .

P. Lalanne, W. Yan, A. Gras, C. Sauvan, J.-P. Hugonin, M. Besbes, G. Demésy, M. D. Truong, B. Gralak, F. Zolla, A. Nicolet, F. Binkowski, L. Zschiedrich, Sven Burger, J. Zimmerling, R. Remis, P. Urbach, H. T. Liu, T. Weiss (2018). **Quasinormal Mode Solvers for Resonators with Dispersive Materials.** *arXiv*. <https://arxiv.org/abs/1811.11751>

Annemarie Lang, Andrea Volkamer, Laura Behm, Susanna Röblitz, Rainald Ehrig, Marlon Schneider, Lisbet Geris, Joerg Wichard, Frank Buttgerit (2018). **In Silico Methods – Computational Alternatives to Animal Testing.** *ALTEX*, 35(1):126–128. <https://doi.org/10.14573/altex.1712031>

Norbert Lindow, Stefanie Redemann, Florian Brünig, Gunar Fabig, Thomas Müller-Reichert, Steffen Prohaska (2018). **Quantification of Three-Dimensional Spindle Architecture.** In *Methods in Cell Biology Part B*, 145:45–64, Academic Press. <https://doi.org/10.1016/bs.mcb.2018.03.012>

*Mehmet Neset Ozel, Abhishek Kulkarni, Amr Hasan, Josephine Brummer, Marian Moldenhauer, Ilsa-Maria Daumann, Heike Wolfenberger, Vincent Dercksen, Ferdi Ridvan Kiral, Martin Weiser, Steffen Prohaska, Max von Kleist, Peter Robin Hiesinger (2018). **Serial Synapse Formation through Filopodial Competition for Synaptic Seeding Factors.** *bioRxiv preprint*. <https://doi.org/10.1101/506378> (Joint publication: Numerical Mathematics, Visual Data Analysis)

Farouk Salem, Florian Schintke, Thorsten Schütt, Alexander Reinefeld (2018). **Data-Flow Scheduling for a Scalable FLEsnet.** *CBM Progress Report 2017*, 130–131. <https://doi.org/10.15120/GSI-2018-00485>

BOOKS (1)

Ralf Borndörfer, Sabine Storandt, Editors (2018). *Proceedings of the 18th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems*. Dagstuhl Publishing, Saarbrücken/Wadern, Germany.

DISSERTATIONS (8)

Carlo Barth (2018). *Analysis of Photonic Crystals for Interaction with Near-Surface Emitters*. Technische Universität Berlin.

Christoph Gorgulla (2018). *Free Energy Methods Involving Quantum Physics, Path Integrals, and Virtual Screenings*. Freie Universität Berlin.

Martina Klimm (2018). *New Strategies in Conformation Dynamics*. Freie Universität Berlin.

Stefan Heinz (2018). *Presolving Techniques and Linear Relaxations for Cumulative Scheduling*, Technische Universität Berlin.

Jannes Quer (2018). *Importance Sampling for Metastable Dynamical Systems in Molecular Dynamics*. Freie Universität Berlin.

Bernhard Reuter (2018). *Generalisierte Markov-Modellierung von Nichtgleichgewichtssystemen – Simulation und Modellierung der Amyloid-beta(1-40)-Konformationsdynamik unter Mikrowelleneinfluss*. Universität Kassel.

Borong Shao (2018). *Phenotype Relevant Network-based biomarker discovery Integrating multiple Omics data – EMT network-based lung cancer prognosis prediction*. Freie Universität Berlin

José Villatoro (2018). *A Combined Approach for the Analysis of Biomolecules Using IR-MALDI Ion Mobility Spectrometry and Molecular Dynamics Simulations of Peptide Ions in the Gas Phase*. Universität Potsdam.

ZIB REPORTS (55)

Lovis Anderson, Benjamin Hiller (2018). *A Sweep-Plane Algorithm for the Computation of the Volume of a Union of Polytopes*. ZIB-Report 18-37.

Isabel Beckenbach, Meike Hatzel, Sebastian Wiederrecht (2018). *The Tight Cut Decomposition of Matching Covered Uniformable Hypergraphs*. ZIB-Report 18-61.

Kai Helge Becker, Benjamin Hiller (2018). *ASTS Orientations on Undirected Graphs: Structural Analysis and Enumeration*. ZIB-Report 18-31.

Timo Berthold, Peter J. Stuckey, Jakob Witzig (2018). *Local Rapid Learning for Integer Programs*. ZIB-Report 18-56.

Ralf Borndörfer, Heide Hoppmann, Marika Karbstein, Niels Lindner (2018). *Separation of Cycle Inequalities in Periodic Timetabling*. ZIB-Report 18-16.

Ralf Borndörfer, Marika Karbstein, Christian Liebchen, Niels Lindner (2018). *A Simple Way to Compute the Number of Vehicles that Are Required to Operate a Periodic Timetable*. ZIB-Report 18-38.

Ralf Borndörfer, Ricardo Euler, Marika Karbstein, Fabian Mett (2018). *Ein mathematisches Modell zur Beschreibung von Preissystemen im öV*. ZIB-Report 18-47.

Ralf Borndörfer, Ascan Egerer, Marika Karbstein, Ralf Messerschmidt, Marc Perez, Steven Pfisterer, Petra Strauß (2018). *Kombilösung: Optimierung des Liniennetzes in Karlsruhe*. ZIB-Report 18-45.

Leon Eifler, Ambros Gleixner, Jonad Pulaj (2018). *Chvátal's Conjecture Holds for Ground Sets of Seven Elements*. ZIB-Report 18-49.

Konstantin Fackeldey, Alexander Sikorski, Marcus Weber (2018). *Spectral Clustering for Non-reversible Markov Chains*. ZIB-Report 18-48.

Alexander Georges, Ambros Gleixner, Gorana Gojic, Robert Lion Gottwald, David Haley, Gregor Hendel, Bartłomiej Matejczyk (2018). **Feature-Based Algorithm Selection for Mixed Integer Programming**. ZIB-Report 18-17.

Ambros Gleixner, Michael Bastubbe, Leon Eiffler, Tristan Gally, Gerald Gamrath, Robert Lion Gottwald, Gregor Hendel, Christopher Hojny, Thorsten Koch, Marco E. Lübbecke, Stephen J. Maher, Matthias Miltenberger, Benjamin Müller, Marc E. Pfetsch, Christian Puchert, Daniel Rehfeldt, Franziska Schlösser, Christoph Schubert, Felipe Serrano, Yuji Shinano, Jan Merlin Viernickel, Matthias Walter, Fabian Wegscheider, Jonas T. Witt, Jakob Witzig (2018). **The SCIP Optimization Suite 6.0**. ZIB-Report 18-26.

Carl Martin Grewe, Gabriel Le Roux, Pilz Sven-Kristofer, Stefan Zachow (2018). **Spotting the Details: the Various Facets of Facial Expressions**. ZIB-Report 18-06.

Richard Hasenfelder, Lutz Lehmann, Manuel Radons, Tom Streubel, Christian Strohm, Andreas Griewank (2018). **Computational Aspects of the Generalized Trapezoidal Rule**. ZIB-Report 18-23.

Gregor Hendel, Matthias Miltenberger, Jakob Witzig (2018). **Adaptive Algorithmic Behavior for Solving Mixed Integer Programs Using Bandit Algorithms**. ZIB-Report 18-36.

Gregor Hendel (2018). **Adaptive Large Neighborhood Search for Mixed Integer Programming**. ZIB-Report 18-60.

Benjamin Hiller, Kai Helge Becker (2018). **Improving Relaxations for Potential-Driven Network Flow Problems via Acyclic Flow Orientations**. ZIB-Report 18-30.

Benjamin Hiller, Tom Walther (2018). **Improving Branching for Disjunctive Polyhedral Models Using Approximate Convex Decompositions**. ZIB-Report 17-68.

Benjamin Hiller, René Saitenmacher, Tom Walther (2018). **Analysis of Operating Modes of Complex Compressor Stations**. ZIB-Report 17-65.

Benjamin Hiller, Thorsten Koch, Lars Schewe, Robert Schwarz, Jonas Schweiger (2018). **A System to Evaluate Gas Network Capacities: Concepts and Implementation**. ZIB-Report 18-11.

Benjamin Hiller, Tom Walther (2018). **Modeling Compressor Stations in Gas Networks**. ZIB-Report 17-67.

Satoshi Ito, Yuji Shinano (2018). **Calculation of Clinch and Elimination Numbers for Sports Leagues with Multiple Tiebreaking Criteria**. ZIB-Report 18-51.

Leonardo Agudo Jácome, Hans-Christian Hege, Olaf Paetsch, Kai Pöthkow (2018). **Three-Dimensional Reconstruction and Quantification of Dislocation Substructures from Transmission Electron Microscopy Stereo Pairs**. ZIB-Report 18-50.

David Knötel, Carola Becker, Gerhard Scholtz, Daniel Baum (2018). **Global and Local Mesh Morphing for Complex Biological Objects from microCT Data**. ZIB-Report 18-53.

Thorsten Koch, Klaus Ceynowa, Konstanze Söllner, Jürgen Christof, Roland Bertelmann (2018). **DeepGreen - Open Access Transformation Etablierung und Weiterentwicklung rechtssicherer Workflows zur effizienten Umsetzung von Open-Access-Komponenten in Lizenzvereinbarungen für wissenschaftliche Publikationen**. ZIB-Report 18-39.

Ralf Lenz, Kai Helge Becker (2018). **Optimization of Capacity Expansion in Potential-Driven Networks Including Multiple Looping - a Comparison of Modelling Approaches**. ZIB-Report 18-44.

Norbert Lindow, Stefanie Redemann, Gunar Fabig, Thomas Müller-Reichert, Steffen Prohaska (2018). **Quantification of Three-Dimensional Spindle Architecture**. ZIB-Report 18-07.

Norbert Lindow, Daniel Baum, Hans-Christian Hege (2018). **Atomic Accessibility Radii for Molecular Dynamics Analysis**. ZIB-Report 18-18.

Norbert Lindow, Daniel Baum, Morgan Leborgne, Hans-Christian Hege (2018). **Interactive Visualization of RNA and DNA Structures**. ZIB-Report 18-33.

James Luedtke, Claudia D'Ambrosio, Jeff Linderoth, Jonas Schweiger (2018). **Strong Convex Nonlinear Relaxations of the Pooling Problem**. ZIB-Report 18-12.

James Luedtke, Claudia D'Ambrosio, Jeff Linderoth, Jonas Schweiger (2018). **Strong Convex Nonlinear Relaxations of the Pooling Problem: Extreme Points**. ZIB-Report 18-13.

Esfandiar Nava-Yazdani, Hans-Christian Hege, Christoph von Tycowicz, T. J. Sullivan (2018). **A Shape Trajectories Approach to Longitudinal Statistical Analysis**. ZIB-Report 18-42.

Arne Nägel, Peter Deuffhard, Gabriel Wittum (2018). **Efficient Stiff Integration of Density Driven Flow Problems**. ZIB-Report 18-54.

Manuel Radons, Lutz Lehmann, Tom Streubel, Andreas Griewank (2018). **An Open Newton Method for Piecewise Smooth Systems**. ZIB-Report 18-43.

Stefanie Redemann, Ina Lantzsch, Norbert Lindow, Steffen Prohaska, Martin Srayko, Thomas Müller-Reichert (2018). **A Switch in Microtubule Orientation during C. elegans Meiosis.** ZIB-Report 18-34.

Daniel Rehfeldt, Thorsten Koch (2018). **SCIP-Jack – a Solver for STP and Variants with Parallelization Extensions: an Update.** ZIB-Report 18-05.

Daniel Rehfeldt, Thorsten Koch (2018). **Reduction-Based Exact Solution of Prize-Collecting Steiner Tree Problems.** ZIB-Report 18-55.

Guillaume Sagnol, Marco Blanco, Thibaut Sauvage (2018). **Approximation Hierarchies for the Cone of Flow Matrices.** ZIB-Report 18-20.

Guillaume Sagnol, Daniel Schmidt genannt Waldschmidt, Alexander Tesch (2018). **The Price of Fixed Assignments in Stochastic Extensible Bin Packing.** ZIB-Report 18-19.

Stephan Schwartz, Leonardo Balestrieri, Ralf Borndörfer (2018). **On Finding Subpaths with High Demand.** ZIB-Report 18-27.

Jonas Schweiger (2018). **Exploiting Structure in Non-convex Quadratic Optimization.** ZIB-Report 18-35.

Robert Schwieger, Heike Siebert, Susanna Röblitz (2018). **Correspondence of Trap Spaces in Different Models of Bioregulatory Networks.** ZIB-Report 18-10.

Feda Seblany, Ulrike Homberg, Eric Vincens, Paul Winkler, Karl Josef Witt (2018). **Merging Criteria for the Definition of a Local Pore and the CSD Computation of Granular Materials.** ZIB-Report 18-09.

Feda Seblany, Ulrike Homberg, Eric Vincens, Paul Winkler, Karl Josef Witt (2018). **Merging Criteria for Defining Pores and Constrictions in Numerical Packing of Spheres.** ZIB-Report 18-25.

Felipe Serrano (2018). **Intersection Cuts for Factorable MINLP.** ZIB-Report 18-59.

Yuji Shinano, Daniel Rehfeldt, Thorsten Koch (2018). **Building Optimal Steiner Trees on Supercomputers by Using up to 43,000 Cores.** ZIB-Report 18-58.

Tom Streubel, Caren Tischendorf, Andreas Griewank (2018). **Piecewise Polynomial Taylor Expansions – the Generalization of Faà di Bruno’s Formula.** ZIB-Report 18-24.

Vikram Sunkara (2018). **Algebraic Expressions of Conditional Expectations in Gene Regulatory Networks.** ZIB-Report 18-52.

Alexander Tack, Anirban Mukhopadhyay, Stefan Zachow (2018). **Knee Menisci Segmentation Using Convolutional Neural Networks: Data from the Osteoarthritis Initiative.** ZIB-Report 18-15.

Alexander Tesch (2018). **A Polyhedral Study of Event-Based Models for the Resource-Constrained Project Scheduling Problem.** ZIB-Report 17-79.

Alexander Tesch (2018). **Improving Energetic Propagations for Cumulative Scheduling.** ZIB-Report 18-29.

Jürgen Titschack, Daniel Baum, Kei Matsuyama, Karin Boos, Claudia Färber, Wolf-Achim Kahl, Karsten Ehrig, Dietmar Meinel, Carmen Soriano, Stuart R. Stock (2018). **Ambient Occlusion – a Powerful Algorithm to Segment Shell and Skeletal Intrapores in Computed Tomography Data.** ZIB-Report 18-14.

Tobias Weber, Sebastian Sager, Ambros Gleixner (2018). **Solving Quadratic Programs to High Precision Using Scaled Iterative Refinement.** ZIB-Report 18-04.

Martin Weiser, Yvonne Freytag, Bodo Erdmann, Michael Hubig, Gita Mall (2018). **Optimal Design of Experiments for Estimating the Time of Death in Forensic Medicine.** ZIB-Report 18-08.

Jakob Witzig, Timo Berthold, Stefan Heinz (2018). **A Status Report on Conflict Analysis in Mixed Integer Nonlinear Programming.** ZIB-Report 18-57.

REFERENCES

PHOTONIC RESONANCES BOOST QUANTUM TECHNOLOGY

- [1] N. Srocka et al. Enhanced Photon-Extraction Efficiency from InGaAs/GaAs Quantum Dots in Deterministic Photonic Structures at 1.3 μm Fabricated by In Situ Electron-Beam Lithography. *AIP Adv.* 8, 085205, 2018.
- [2] P. Schnauber et al. Deterministic Integration of Quantum Dots into On-Chip Multimode Interference Beamsplitters Using In Situ Electron-Beam Lithography. *Nano Lett.* 18, 2336, 2018.
- [3] L. Zschiedrich et al. Riesz-Projection-Based Theory of Light-Matter Interaction in Dispersive Nanoresonators. *Phys. Rev. A*, 98, 043806, 2018.
- [4] C. Becker et al. Nanophotonic Enhanced Two-Photon Excited Photoluminescence of Perovskite Quantum Dots. *ACS Photonics* 5, 4668, 2018.
- [5] G. Kewes et al. Heuristic Modeling of Strong Coupling in Plasmonic Resonators. *ACS Photonics* 5, 4089, 2018.
- [6] P. Lalanne et al. Quasinormal Mode Solvers for Resonators with Dispersive Materials. *J. Opt. Soc. Am. A*, 2018

HOW MATHEMATICAL INTELLIGENCE CAN HELP WITH YOUR KNEE PAIN

- B. Shao, M. Bjaanæs, Å. Helland, Ch. Schütte, T. O. F. Conrad. EMT Network-Based Feature Selection Improves Prognosis Prediction in Lung Adenocarcinoma. Accepted to *PLoS ONE*, 2018.
- F. Jayrannejad, T. O. F. Conrad. Better Interpretable Models for Proteomics Data Analysis Using Rule-Based Mining. *Springer Lecture Notes in Artificial Intelligence*, 2018.
- B. Shao, C. Cannistraci, T. O. F. Conrad. Epithelial Mesenchymal Transition Network-Based Feature Engineering in Lung Adenocarcinoma Prognosis Prediction Using Multiple Omic Data. *Genomics and Computational Biology*, 3 (3). ISSN 2365-7154, 2018.
- T. O. F. Conrad, M. Genzel, N. Cvetkovic, N. Wulkow, J. Vybiral, G. Kutyniok, Ch. Schütte. Sparse Proteomics Analysis – a Compressed Sensing-Based Approach for Feature Selection and Classification of High-Dimensional Proteomics Mass Spectrometry Data. *BMC Bioinformatics*, 18 (160). ISSN 1471-2105, 2018.
- IT'S ALL IN THE IMAGES FOR DIAGNOSIS & THERAPY PLANNING

- [1] A. Tack, S. Zachow: Accurate Automated Volumetry of Cartilage of the Knee using Convolutional Neural Networks: Data from the Osteoarthritis Initiative. To appear in *Proc. of 16th IEEE International Symposium on Biomedical Imaging (ISBI)*, 2019.
- [2] H.H. Gidey: Automated Hip-Knee-Ankle Angle Determination using Convolutional Neural Networks. Master thesis (ZIB/Zachow & Otto-von-Guericke Universität Magdeburg/Preim), 2019.
- [3] F. Ambellan, A. Tack, M. Ehlke, S. Zachow: Automated Segmentation of Knee Bone and Cartilage combining Statistical Shape Knowledge and Convolutional Neural Networks: Data from the Osteoarthritis Initiative. *Medical Image Analysis*, 52(2):109-118, 2019.

- [4] Tack A, Mukhopadhyay A, Zachow S: Knee Menisci Segmentation Using Convolutional Neural Networks: Data from the Osteoarthritis Initiative. *Osteoarthritis and Cartilage*, 26(5):680-688, 2018.

- [5] Zachow: An Efficient Riemannian Statistical Shape Model using Differential Coordinates. *Medical Image Analysis*, 43(1):1-9, 2018.

- [6] R. Joachimsky, F. Ambellan, S. Zachow: Computerassistierte Auswahl und Platzierung von interpositionalen Spacern zur Behandlung früher Gonarthrose. *Proc. of Jahrestagung der Deutschen Gesellschaft für Computer- und Roboterassistierte Chirurgie (CURAC)*, 16:106-111, 2018.

- [7] D. Wilson, C. Anglin, F. Ambellan, C. M. Grewe, A. Tack, H. Lamecker, M. Dunbar, S. Zachow: Validation of Three-Dimensional Models of the Distal Femur Created from Surgical Navigation Point Cloud Data for Intraoperative and Postoperative Analysis of Total Knee Arthroplasty. *International Journal of Computer Assisted Radiology and Surgery*, 12(12):2097-2105, 2018.

- [8] M. Akbari-Shandiz, P. Boulos, S. Sævarsson, H. Ramm, C. Fu, S. Miller, S. Zachow, C. Anglin: Changes in Knee Shape and Geometry Resulting from Total Knee Arthroplasty. *Journal of Engineering in Medicine*, 232(1):6-79, 2018.

OPIOIDS: MANAGING THE PAIN

- [1] Drugcom.de, news "Fentanyl-Pflaster vermutlich für Todesfälle verantwortlich," 2018-06-01.
- [2] V. Spahn, G. Del Vecchio, D. Labuz, A. Rodriguez-Gaztelumendi, N. Massaly, J. Temp, V. Durmaz, P. Sabri, M. Reidelbach, H. Machel-ska, M. Weber, C. Stein: A Nontoxic Pain Killer Designed by Modeling of Pathological Receptor Conformations. *Science*, 355(6328):966-969, March 2018.

DIGITAL CULTURAL ASSETS – DIGIS OR THE BERLIN MARVEL

[3] V. Spahn, G. Del Vecchio, A. Rodriguez-Gaztelumendi, J. Temp, D. Labuz, M. Kloner, M. Reidelbach, H. Machelska, M. Weber, C. Stein: Opioid Receptor Signaling, Analgesic and Side Effects Induced by a Computationally Designed pH-Dependent Agonist. *Scientific Reports*, 8, Art. no. 8965, 2018.

[4] M. Klimm. Dissertation thesis: “New Strategies in Conformation Dynamics,” FU Berlin/SALSA, Mathematics, 2018.

[5] M. Reidelbach, M. Weber, P. Imhof: Prediction of Perturbed Proton Transfer Networks. *PLoS ONE*, 13(12): e0207718, 2018.

[6] K. Schrade, J. Tröger, A. Eldahshan, K. Zühlke, K. R. Abdul Azeez, J.M. Elkins, M. Neuenschwander, A. Oder, M. Elkewedi, S. Jaksch, K. Andrae, J. Li, J. Fernandes, P.M. Müller, S. Grunwald, S.F. Marino, T. Vukicevic, J. Eichhorst, B. Wiesner, M. Weber, M. Kapiloff, O. Rocks, O. Daumke, T. Wieland, S. Knapp, J.P. von Kries, E. Klussmann: An AKAP-Lbc-RhoA Interaction Inhibitor Promotes the Translocation of Aquaporin-2 to the Plasma Membrane of renal Collecting Duct Principal Cells. *PLoS ONE*, 13(1):e0191423, 2018.

[7] L. Donati, M. Heida, B. G. Keller, M. Weber: Estimation of the Infinitesimal Generator by Square-Root Approximation. *J. Phys.: Condens. Matter*, 30(42):425201, 2018.

STEP BY STEP

[1] A. Burchardt, C. Abicht, O. Sander: Simulating Wear on Total Knee Replacements, arXiv: 1704.08307, 2018.

[2] M. Weiser, S. Ghosh: Theoretically Optimal Inexact SDC Methods. *Comm. Appl. Math. Comp. Sci.* 13(1): 53–86, 2018.

[1] Refer to <http://www.friedenspreis-des-deutschen-buchhandels.de/1244997/> (last accessed 2019-01-18).

[2] Refer to “Gemeinsam gute Daten schaffen” – the 2018 digiS film, <https://www.digis-berlin.de/wissenswertes/digis-film-2018/> (last accessed 2019-01-18).

[3] Refer to Christina Riesenweber, 2018-12-06, https://twitter.com/c_riesen/status/1070632935903162368 (last accessed 2019-01-21).

[4] Refer to the digiS website, <https://www.digis-berlin.de/> (last accessed 2019-01-18).

[5] Refer to <https://www.digis-berlin.de/digis-veranstaltungen/workshops/> (last accessed 2019-01-18). Participation in the workshops is open to all interested parties and is free of charge. digiS conducts eight to ten workshops per year, which are attended by an average of 20 participants. In 2018, 217 interested people took part in the workshops.

[6] Materials from our workshop offers are available for download here: <https://www.digis-berlin.de/wissenswertes/> (last accessed 2019-01-18).

[7] Refer to DDB Pro <https://pro.deutsche-digitale-bibliothek.de/> (last accessed 2019-01-18).

[8] Refer to topic “Langzeitverfügbarkeit” <https://www.digis-berlin.de/wissenswertes/langzeitverfuegbarkeit/> (last accessed 2019-01-18) as well as KOBV pages <https://www.kobv.de/services/archivierung/lza/> (last accessed 2019-01-18).

[9] Convention for the Protection of Cultural Property in the Event of Armed Conflict with Regulations for the Execution of the Convention 1954 http://portal.unesco.org/en/ev.php-URL_ID=13637&URL_DO=DO_TOPIC&URL_SECTION=201.html (last accessed 2019-01-18).

[10] <https://www.unesco.de/kultur-und-natur/kulturgutschutz> (last accessed 2019-01-18).

[11] Refer to Georg Schelbert: Bildgeschichte digital greifbar. Die Glasdiasammlung des Instituts für Kunst- und Bildgeschichte der Humboldt-Universität zu Berlin. Report about a work in progress. DOI <https://doi.org/10.18452/19463>, 2018-10-09.

[12] Refer to Dr. Peter Jammerthal’s final report https://www.digis-berlin.de/wp-content/uploads/2018/04/Projektabschlussbericht_FU-Berlin_web.pdf and Julian Nordhues’ project documentation https://www.digis-berlin.de/wp-content/uploads/2018/09/Projektdokumentation_Regiebuch_Max-Reinhardt_2018-04-06.pdf (last accessed 2019-01-18).

[13] Refer to project partners from 2012 to 2019 <https://www.digis-berlin.de/projektpartner/> (last accessed 2019-01-18).

[14] Refer to <https://codingdavinci.de/> (last accessed 2019-01-19); Coding da Vinci was founded in 2014 as a cooperation of the German Digital Library, digiS, Open Knowledge Foundation Deutschland e. V., and Wikimedia Deutschland e. V. In 2019, Coding da Vinci received a four-year grant from the German Federal Cultural Foundation. Refer to https://www.kulturstiftung-des-bundes.de/de/projekte/film_und_neue_medien/detail/kultur_digital.html and <https://www.digis-berlin.de/kulturstiftung-des-bundes-foerdert-kultur-hackathon-coding-da-vinci/> (last accessed 2019-01-20).

[15] The data from the digiS project of the Stiftung Berliner Mauer were used in the MauAR app by Peter Kolski, who promptly attracted the interest of Apple CEO Tim Cook. <https://www.berliner-zeitung.de/digital/-mauar---eine-app-mit-der-ein-historisches-hauptstadtgefuehl-ensteht-31544208> (last accessed 2019-01-19).

[16] Digital cultural data have also found their way into the State of Berlin's 2015 Open Access strategy. In order to achieve the broadest possible accessibility, the State of Berlin advocates the widest possible access – within the framework of what is legally possible – to digital cultural data. Refer to <https://www.parlament-berlin.de/ados/17/IIIPlen/vorgang/d17-2512.pdf> (last accessed 2019-01-20).

[17] Refer to <http://www.friedenspreis-des-deutschen-buchhandels.de/1244997/> (last accessed 2019-01-19).

IMAGE LICENSES

[1] CC0 1.0. <https://creativecommons.org/publicdomain/zero/1.0/>

[2] PDM 1.0. <https://creativecommons.org/publicdomain/mark/1.0/>

[3] CC BY-SA 4.0. <https://creativecommons.org/licenses/by-sa/4.0/>

[4] PDM 1.0. <https://creativecommons.org/publicdomain/mark/1.0/>

[5] PDM 1.0. <https://creativecommons.org/publicdomain/mark/1.0/>

[6] PDM 1.0. <https://creativecommons.org/publicdomain/mark/1.0/>

[7] CC0 1.0. <https://creativecommons.org/publicdomain/zero/1.0/>

[8] CC BY 3.0 DE. <https://creativecommons.org/licenses/by/3.0/de/>

[9] CC BY-SA 3.0 DE. <https://creativecommons.org/licenses/by-sa/3.0/de/>

OPTIMIZATION MEETS SUPERCOMPUTING

[1] Gerald Gamrath, Thorsten Koch, Stephen Maher, Daniel Rehfeldt, Yuji Shinano: SCIP-Jack – a Solver for STP and Variants with Parallelization Extensions. *Mathematical Programming Computation*, 9(2):231–296. <https://doi.org/10.1007/s12532-016-0114-x>, 2018.

[2] Ivana Ljubić, René Weiskircher, Ulrich Pfersch, Gunnar W. Klau, Petra Mutzel, Matteo Fischetti: An Algorithmic Framework for the Exact Solution of the Prize-Collecting Steiner Tree Problem. *Mathematical Programming*, 105(2–3):427–449. <https://doi.org/10.1007/s10107-005-0660-x>, 2006.

[3] <https://pacechallenge.org/>

[4] Yuji Shinano, Tobias Achterberg, Timo Berthold, Stefan Heinz, Thorsten Koch, Michael Winkler: Solving Open MIP Instances with ParaSCIP on Supercomputers using up to 80,000 Cores. 2016 IEEE International Parallel and Distributed Processing Symposium, <https://doi.org/10.1109/IPDPS.2016.56>, 2016.

[5] <http://dimacs11.zib.de/>

FAULT TOLERANCE IN DISTRIBUTED SYSTEMS

[1] Leslie Lamport. Paxos Made Simple. *ACM Sigact News* 32.4, 2001.

[2] Seth Gilbert, Nancy Lynch. Brewer's Conjecture and the Feasibility of Consistent, Available, Partition-Tolerant Web Services. *ACM Sigact News* 33.2, 2002.

[3] Michael J. Fischer, Nancy A. Lynch, Michael S. Paterson. Impossibility of Distributed Consensus with One Faulty Process. No. MIT/LCS/TR-282. Massachusetts Institute of Tech Cambridge Lab for Computer Science, 1982.

[4] Ashley Feinberg. An 83,000-Processor Supercomputer Can Only Match 1% of Your Brain, August 2013.

[5] Jack J. Dongarra, et al. The International Exascale Software Project Roadmap. *IJHPCA*, 25(1):3–60, 2011.

[6] John T. Daly. A Higher Order Estimate of the Optimum Checkpoint Interval for Restart Dumps. *Future Generation Computer Systems*, 22(3):303–312, 2006.

[7] John W. Young. A First Order Approximation to the Optimal Checkpoint Interval. *Commun. ACM*, 17(9):530–531, 1974.

[8] Sheng Di, Mohamed-Slim Bouguerra, Leonardo Arturo Bautista-Gomez, Franck Cappello. Optimization of Multi-level Checkpoint Model for Large Scale HPC Applications. *IPDPS 2014*, 1181–1190. IEEE Computer Society, 2014.

[9] Adam Moody, Greg Bronevetsky, Kathryn Mohror, Bronis R. de Supinski. Design, Modeling, and Evaluation of a Scalable Multi-level Checkpointing System. *International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2010*, 1–11. IEEE, 2010.

THE DEEP PARALLELISM CHALLENGE

[10] Leonardo Arturo Bautista-Gomez, Seiji Tsuboi, Dimitri Komatitsch, Franck Cappello, Naoya Maruyama, Satoshi Matsuoka. FTI: High Performance Fault Tolerance Interface for Hybrid Systems. International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2011, 32:1–32:32 IEEE, 2011.

[11] Masoud Gholami, Florian Schintke, Thorsten Schütt. Checkpoint Scheduling for Shared Usage of Burst-Buffers in Supercomputers. In ICPP Workshops, SRMPDS 2018, 44:1–44:10, ICPP, 2018.

[1] F. Wende, M. Marsman, J. Kim, F. Vasilev, Z. Zhao, T. Steinke. OpenMP in VASP: Threading and SIMD. *Int J Quantum Chem.*, e25851, <https://doi.org/10.1002/qua.25851>, 2018.

[2] G. Kresse, J. Furthmüller. Efficient Iterative Schemes for Ab Initio Total-Energy Calculations Using a Plane-Wave Basis Set, *Phys. Rev. B* 54, 11169, <http://link.aps.org/doi/10.1103/PhysRevB.54.11169>, 1996.

[3] G. Kresse, J. Furthmüller. Efficiency of Ab-Initio Total Energy Calculations for Metals and Semiconductors Using a Plane-Wave Basis Set, *Comput. Mater. Sci.* 6, 15–50, [https://doi.org/10.1016/0927-0256\(96\)00008-0](https://doi.org/10.1016/0927-0256(96)00008-0), 1996.

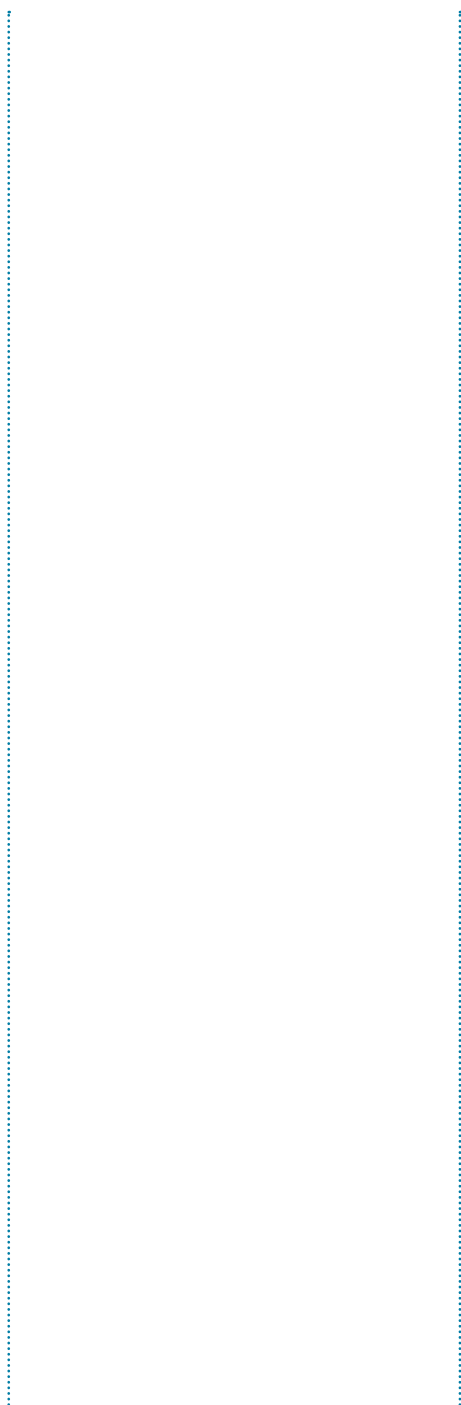
[4] M. Marsman, J. Paier, A. Stroppa, G. Kresse, Hybrid Functionals Applied to Extended Systems, *J. Phys. Condens. Matter* 20, 064201, <https://doi.org/10.1088/0953-8984/20/6/064201>, 2008.

[5] M. Kaltak, J. Klimes, G. Kresse, Cubic Scaling Algorithm for the Random Phase Approximation: Self-Interstitials and Vacancies in Si, *Phys. Rev. B: Condens. Matter Mater. Phys.* 90, 054115, <https://doi.org/10.1103/PhysRevB.90.054115>, 2014.

[6] P. Liu, M. Kaltak, J. Klimes, G. Kresse, Cubic Scaling GW: Towards Fast Quasiparticle Calculations, *Phys. Rev. B Condens. Matter* 94, 165109, <https://doi.org/10.1103/PhysRevB.94.165109>, 2016.

[7] P. E. Blöchl, Projector Augmented-Wave Method, *Phys. Rev. B: Condens. Matter Mater. Phys.* 50, 17953, <https://doi.org/10.1103/PhysRevB.50.17953>, 1994.

[8] G. Kresse, D. Joubert. From Ultrasoft Pseudopotentials to the Projector Augmented-Wave Method, *Phys. Rev. B* 59, 1758, <https://doi.org/10.1103/PhysRevB.59.1758>, 1999.



NOTES

IMPRINT

Published by

Konrad-Zuse-Zentrum für Informationstechnik Berlin
Takustrasse 7 | 14195 Berlin | Germany

Phone +49-30-841-850
Fax +49-30-8418-5125
E-mail info@zib.de
www.zib.de

Layout

SCHIEBE PREIL BAYER
werbung und kommunikation – Berlin

Printing

Lochmann grafische Produktionen – Berlin

ISSN

0934-5892

©2019 Konrad-Zuse-Zentrum für Informationstechnik Berlin



ZIB

WWW.ZIB.DE