

ANNUAL REPORT 2022



Created by an AI

CONTENTS

4	Executive Summary
8	Organization
10	ZIB Structure
12	Mathematics of Complex Social Systems: Past, Present, and Future
16	Economic Situation in 2022
22	Taking Math to the Heart
36	Numerics meets Quantum Technology
48	Hydrogen Blending
59	Computer-Assisted Proofs in Extremal Combinatorics
66	The Tropical Geometry of Periodic Timetabling
77	Programming on the Data Highway
86	Addressing Energy Challenges in HPC Data Centers
96	References
100	Publications
108	Imprint

EXECUTIVE SUMMARY

After the outbreak of the COVID-19 pandemic in 2020, the year 2022 saw a gradual return to a scientific culture of collaboration through face-to-face meetings and direct exchanges, as opposed to relying solely on flat screens. At ZIB, researchers were able to reunite in person once again, with group seminars, workshops, and conferences resuming as part of everyday life from mid-2022.

Despite the ongoing pandemic and the uncertainties brought about by the Russian war against Ukraine, we are proud to say that 2022 was a year filled with scientific highlights and positive developments at ZIB. These include exceptional scientific achievements, a variety of scientific awards, and the successful acquisition of large third-party grants for collaborative research centers. Additionally, we formed new strategic partnerships and hosted engaging conferences. Here are just a few examples from our long list of accomplishments and events in 2022:

- ZIB researchers have optimized nanostructures for multijunction solar cells based on which, in collaboration with our partner Helmholtz Zentrum

Berlin, monolithic perovskite/silicon tandem solar cells could be built with demonstrated efficiencies toward 29.8%, which is a world record in efficiency for these devices.

- In collaboration with the Research Campus MODAL and MATH+, researchers at ZIB have developed a number of new mathematical methods combining discrete and continuous optimization that allow fuel efficiency in flight planning to be improved. Together with other algorithms developed at ZIB, the new methods now form the algorithmic core of the VOLAR optimization engine of the commercial flight route planning suite of Lufthansa Systems, helping to significantly reduce greenhouse gas emissions in the real world.
- Science & Startups' Research to Market Challenge is a competition for research-based business and startup ideas. The prize in the category "Artificial Intelligence" was awarded to ZIB researchers Jan-Patrick Clarner and Christine Tawfik for their planning tool for multi-objective optimization of investments and production portfolios in the energy sector.



- The extension proposals for the DFG collaborative research centers CRC 1114 "Scaling cascades in complex systems", CRC 1444 "Directed cellular self-organization to advance bone regeneration", and TRR 154 "Mathematical modeling, simulation, and optimization using the example of gas networks" were successful, and will fund a significant number of excellent research projects at ZIB during the next four years. This was complemented by the acquisition of a variety of additional third-party funded single and collaborative projects.
- The German Modeling Network for Severe Infectious Diseases (MONID), funded by BMBF, was established, comprising ten research consortia, one of which is also located at ZIB. MONID serves as the primary point of contact for inquiries from the German government regarding the current and future pandemic situations.
- The Thematic Einstein Semester "The Mathematics of Complex Social Systems: Past, Present, and Future" was held at ZIB in 2022. The event consisted of an opening and a final conference, three workshops, a summer school, and an Einstein lecture series, supported by the MATH+ and the Einstein Foundation Berlin. It brought together several hundred researchers from diverse fields, including sociology, archaeology, high-performance computing, and mathematics with a focus on sustainability science.
- ZIB formed an Intel Center of Excellence focused on using oneAPI cross-architecture programming for energy-efficient HPC computing by delivering portable implementations on GPUs and FPGAs, and entered a strategic partnership with Google to further the common interests in high-performance cloud computing as well as large-scale compute applications for AI, Optimization, and Simulation.

We are proud to share that ZIB remains committed to our highest priority – research excellence. The numerous positive developments that took place over the past year have bolstered our confidence in this mission. We are thrilled to see that our core objectives remain relevant and continue to attract increased interest. Our goal is to advance Computational Thinking through a data-driven approach centered on Computational Science, application-oriented Mathematics, and High-Performance Computing methods. With partners from science, technology, and industry, we strive to solve societal problems using high-end computing and data management infrastructures. These efforts have contributed to ZIB's reputation as a place for excellent research and first-rate scientific services and infrastructure, and we remain steadfast in our pursuit of cutting-edge advancements.

2022 was a year full of opportunities but also challenges. One of the major challenges we faced was the energy crisis in Germany, which was compounded by the onset of the conflict between Russia and Ukraine. The resulting surge in electricity prices had a significant impact on our operations, especially on the energy consumption costs of ZIB's HPC infrastructure. The doubling of these costs from 2021 to the beginning of 2023 put a considerable strain on our budget, requiring our experts and leadership to invest hundreds of working hours in crisis prevention. In response, we intensified our development and research programs on green HPC, which were successfully supported by additional funding. Although we remain committed to finding long-term solutions to reduce our energy consumption costs, we recognize that this issue will remain a significant part of our HPC strategy for the foreseeable future.

The pandemic is gradually receding, but its effects persist, along with the ongoing conflict in Ukraine and the accelerating threat of global warming. These challenges will present entirely new types of problems for the sciences, including the need to quickly find sustainable solutions to mitigate the impact of climate change. We believe that computational science is particularly relevant in combating these challenges as it allows for realistic scenario studies that are impossible, too costly, or time-consuming to perform in the real world. This, in turn, promotes evidence-based decision-making to address these critical issues. As a result, ZIB has begun to shift its research strategy and project portfolio towards computational sustainability. This includes supporting the transition to sustainable mobility, energy, and public health, as well as approaches to better understand complex social and societal processes.


This annual report provides a comprehensive overview of ZIB's organizational structure, research statistics, and key factors driving our successful development. The report offers valuable insights into the diverse range of research activities at ZIB through seven "Feature Articles", highlighting our accomplishments and progress but also the breadth of our research agenda.

1. "Taking math to the heart" explains our recent results on how diagnosis and therapy of cardiovascular diseases benefit from numerical simulation of cardiac excitation and the data analysis of resulting electrocardiograms. Research at ZIB contributes efficient numerical algorithms for simulation and parameter identification as well as model-free methods for the analysis of recorded data.

2. In “Numerics meets quantum technology” insight is provided on how several research groups at ZIB, including the newly created group “Quantum Computation and Optimization”, work together to develop numerical methods for modeling, simulating, and optimizing quantum technology applications.
3. “Hydrogen blending” reports on how research at ZIB contributes to the decarbonization of the European energy infrastructure. With a growing hydrogen market, pipeline-based transportation using the existing natural gas infrastructure becomes economically feasible, and accelerates the transition process. At Research Campus MODAL, mathematical methods were developed to analyze the maximum feasible injection of hydrogen with respect to regulatory limits of gas quality.
4. “Computer-assisted proofs in extremal combinatorics” explains how we used techniques from optimization to obtain novel results relating to some long-standing open problems in extremal combinatorics where the limit of human intuition and ability to both execute and verify complex proofs has met its limits and the increasing formalization and automation of proof techniques is required.
5. “The tropical geometry of periodic timetabling” reports on recent progress in analyzing the inherent geometry of periodic timetables that allows for new cutting-edge optimization algorithms for creating timetables in public transport that maximize passenger satisfaction while guaranteeing operational practicality.
6. “Programming on the data highway” starts from the insight that data parallelism, i.e., performing the same operations on a bunch of data in parallel, is a crucial factor to gain best performance from modern microprocessors, and illustrates how, at ZIB, we study and strive for leverage modern microprocessors’ data parallel capabilities in a variety of applications, aiming to use existing hardware more effectively, ultimately saving energy and accelerating algorithms.
7. Last but not least, “Addressing energy challenges in HPC data centers” reports on research at ZIB aimed at identifying potential areas in the NHR@ZIB supercomputer system for energy-saving measures, in particular regarding the question of whether reduced energy consumption can be achieved by slowing down CPU speeds and which specific code characteristics may be instrumental in this regard.



Christof Schütte
President
schuette@zib.de



Sebastian Pokutta
Vice President
pokutta@zib.de

ORGANIZATION

The Statutes

The ZIB Statutes, adopted by the Board of Directors at its meeting on April 28, 2021, redefines ZIB's bodies and organizational structure based on the ZIB law. Furthermore, it determines ZIB's research and development mission and its service tasks, and frame upon the composition of the Scientific Advisory Board and its role.

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. Sebastian Pokutta

The Board of Directors was composed as follows:

Prof. Dr. Julia von Blumenthal

President, Humboldt-Universität zu Berlin
since October 1, 2022

Prof. Dr. Geraldine Rauch

President, Technische Universität Berlin (Chairwoman)
since April 1, 2022

Prof. Dr. Günter Ziegler

President, Freie Universität Berlin

Christian Hingst

Senatsverwaltung für Wissenschaft, Gesundheit, Pflege
und Gleichstellung

Thomas Frederking

Helmholtz-Zentrum Berlin für Materialien und Energie
(HZB)

The Board of Directors met on May 30, 2022, and December 1, 2022.

Scientific Advisory Board

The Scientific Advisory Board advises ZIB on scientific and technical issues, supports ZIB's work, and facilitates ZIB's collaboration 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 (Chairman)
Carleton University, Ottawa, Canada

Prof. Dr. Frauke Liers
Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany

Prof. Dr. Michael Dellnitz
Universität Paderborn, Germany

Prof. Dr. Rolf Krause

Università della Svizzera italiana, Lugano, Switzerland

Ludger D. Sax

Grid Optimization Europe GmbH, Essen, Germany

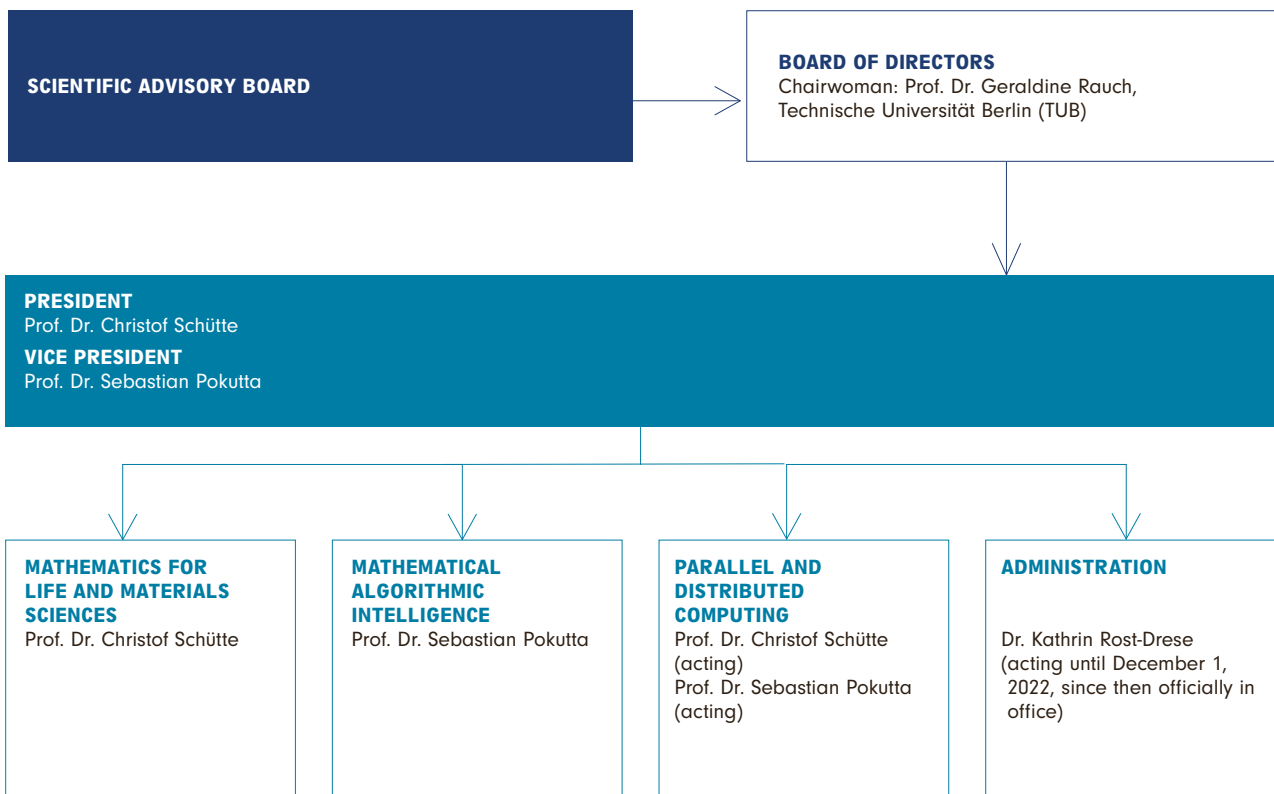
Prof. Dr. Reinhard Schneider

Université du Luxembourg, Luxembourg

Prof. Dr. Dorothea Wagner

Karlsruher Institut für Technologie (KIT), Karlsruhe, Germany
(Membership shall be held in abeyance during the continuing service as Chair of the German Science Council)

The Scientific Advisory Board met on July 4, and July 5, 2022.



ZIB STRUCTURE

Administrative Council

President (C. Schütte) | **Vice-President** (S. Pokutta) | **Managing Director** (K. Rost-Drese)

Research Departments

Mathematics of Complex Systems
(C. Schütte)

Mathematical Algorithmic
(S. Pokutta)

**Department
"Modelling and
Simulation of
Complex Processes"**
(M. Weiser)

**Department
"Visual and
Data-Centric
Computing"**
(T. Conrad)

**Department
"AI in Society,
Science, and
Technology"**
(S. Pokutta)

**Department
"Applied
Intelligence"**
(T. Koch)

Research Service Units

IT and Data Services
(C. Schäuble)

**Digital Data and
Information for
Society, Science,
and Culture**

Administration (K. Rost-Drese)

Ombudsperson/Good Scientific Practice
(J. Sack)

Staff Council Chairperson
(R. Skillen)

Scientific Advisory Board

Intelligence

Parallel and Distributed Computing
(C. Schütte, act. | S. Pokutta, act.)

**Algorithmic
Methods"**

**Department
"Network
Optimization"**
(R. Borndörfer)

**Department
"Distributed
Algorithms"**
(F. Schintke)

**Department
"Supercomputing"**
(T. Steinke)

(B. Rusch, T. Koch)

NHR Center
(T. Steinke, act.)

Women's Representative
(M. Kostre)

Data Protection Officer
(O. Paetsch)

Open Data Officer
(T. Hasler)

MATHEMATICS OF COMPLEX SOCIAL SYSTEMS: PAST, PRESENT, AND FUTURE

Thematic Einstein Semester


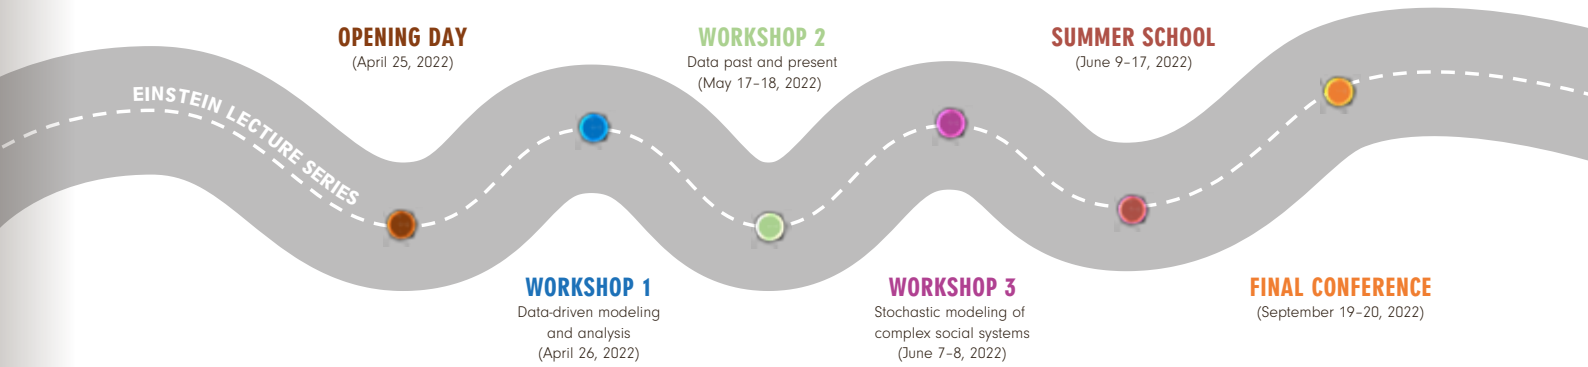
In the summer semester 2022, ZIB hosted the Thematic Einstein Semester (TES)
“The Mathematics of Complex Social Systems: Past, Present, and Future”.

This TES aimed at unlocking the potential for mathematical modeling and reasoning within the extremely large and diversified fields of study that constitute modern social sciences and the humanities. As such, the TES was of a strongly interdisciplinary and pioneering nature that brought together researchers from mathematics and other disciplines to develop synergies and to contribute to computational analysis of complex social systems.

The TES consisted of a series of events, starting on April 25 with the opening day, followed by three workshops, a summer school, and a final conference. In addition, there have been continuous activities over the semester: The Einstein lecture series with talks by experts from the field on hot topics and open challenges, a seminar for undergraduate students to enter the field, and the TES Data Challenge for practical work on real-world data. With a total of 62 invited speakers from all over the world and on average 50 participants in every event of this TES (most of them from the Berlin area), it was discussed how innovative ideas and concepts from mathematics can be developed and made accessible for the study of past and present social systems.

The success of this semester lies in diverse presentations that combined traditional and modern approaches from different fields, for example, agent- and network-based modeling and simulation of social networks in the context of archaeology, epidemiology, or climate research. This resulted in fruitful discussions revealing potential future collaborations. Throughout the semester, the need for mathematical modeling of such systems was emphasized, because it can help in providing forecasts, support a better understanding of social processes, and facilitate structured reasoning about systems whose complexity defies intuition.

This TES has been supported by the Berlin Mathematics Research Center Math+ and by the Einstein Foundation Berlin. It has been organized by scientists from ZIB (Nataša Djurdjevac Conrad, Stefanie Winkelmann, Christof Schütte) together with cooperation partners from Freie Universität Berlin and Global Climate Forum (Sarah Wolf), University of Surrey (Stefan Klus), Deutsches Archäologisches Institut (Benjamin Ducke, Friederike Fless), and Potsdam Institute for Climate Impact Research (Jürgen Kurths).



Nataša Djurdjevac Conrad
Phone +49 30 841 85-439
natasa.conrad@zib.de



Stefanie Winkelmann
Phone +49 30 841 85-156
winkelmann@zib.de



Christof Schütte
Phone +49 30 841 85-101
schuette@zib.de

Making Berlin's scientific high-speed data network fit for the future

In the digital age, high-throughput data exchange and communication networks are of crucial importance. The science network BRAIN is operating Berlin's scientific high-speed data network. With its headquarters at ZIB, BRAIN is available to all scientific, cultural, and educational institutions based in Berlin. Since 2022, the core capacities have been crucially improved, including the fiber-optic transmission network which can now provide up to 400 Gbit/s.

With its more than 200 sites at more than 40 partner institutions, the Berlin Research Area Information Network (BRAIN) has a wide reach and a broad user base. BRAIN provides highly available core infrastructures, including, among other things, site networking (BRAIN-Verbund) and Internet supply (BRAIN-Versorger).

In 2022, major expansion measures could be completed. The BRAIN planning group implemented a new community Internet connection (BRAIN-Versorger) for the Berlin research landscape in collaboration with the German Research Network (DFN). For this purpose, new technology for the BRAIN core network and the BRAIN-Versorger (routers, fiber optic transmission technology) was procured and put into operation.

The BRAIN-Versorger service now ensures the Internet traffic of the core research institutions, including an increase from 30 Gbit/s to 100 Gbit/s. Simultaneously, performance optimization of the fiber optic lines without construction measures led to an increase from 100 to 400 Gbit/s for data exchange. The BRAIN-Versorger allows BRAIN participants to obtain Internet access directly from BRAIN instead of from the German Research Network (DFN). This allows the BRAIN network to

provide not only broadband communication but also secure exchange of data between Berlin's research institutions.

The Berlin Senate has decided to provide additional funds for BRAIN for the modernization of its core network with high-performance technology, in particular for the expansion and upgrading of the fiber optic network. This will allow for further expansion and the migration of further institutions in 2023. Thus, the BRAIN network becomes an active data interconnection network fit for integration of enhanced services in data-related research and teaching.

For 30 years, the Zuse Institute Berlin (ZIB) has been supporting universities in the field of supercomputing. This is now being extended by services for data science and teaching-supporting IT tools, such as Kubernetes Docker systems for Jupyter notebooks. There is also an increasingly high demand for the provision of high-performance AI systems for teaching and research, including corresponding data services. The expansion of the Berlin Research Area Information Network takes this dynamic development into account.



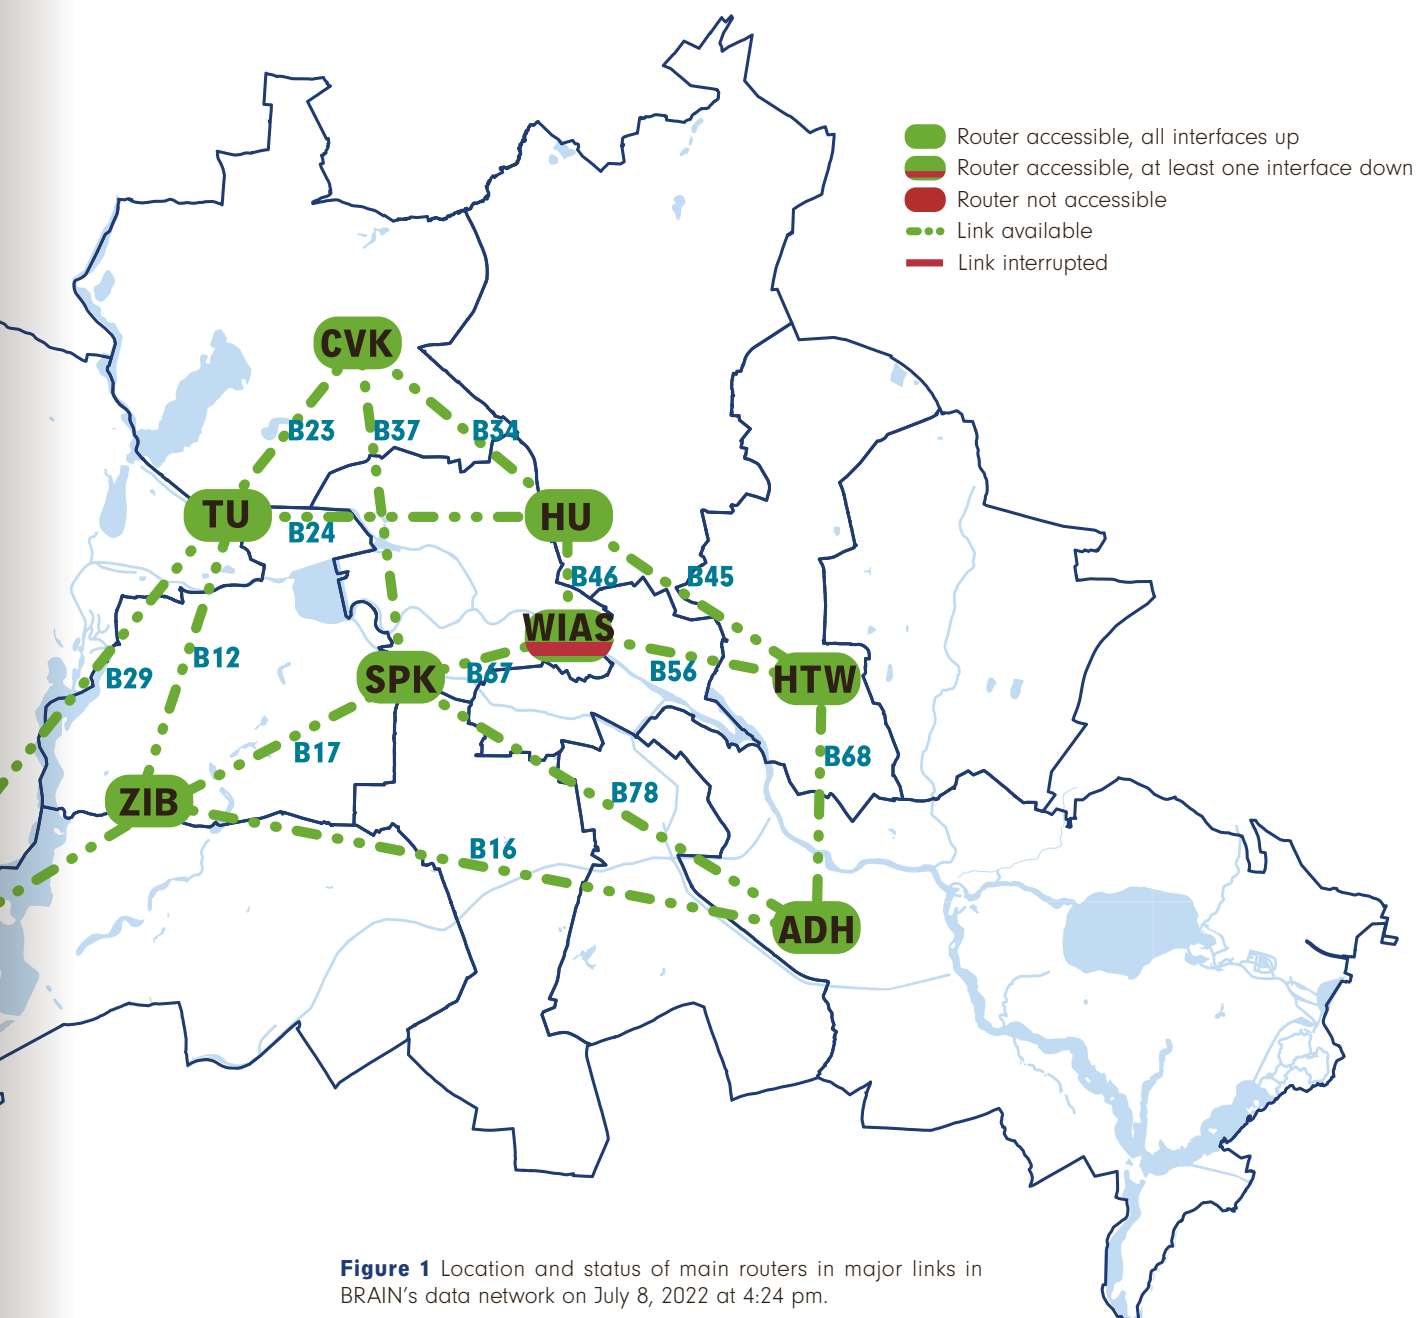



Figure 1 Location and status of main routers in major links in BRAIN's data network on July 8, 2022 at 4:24 pm.

ECONOMIC SITUATION IN 2022

In 2022, the total income of ZIB comprised 28.5 million euros. The main part of this was made available by the Federal State of Berlin as the basic financial stock of ZIB (9.2 million euros) including general investments. The basic financial stock was complemented by the budget of the NHR center at ZIB (7.8 million in 2022, made available by the Federal State of Berlin but co-funded by the federal government and the German Federal States participating in HLRN), and by third-party funds (9.7 million euros in total) 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 research service budget of KOBV, adding up to almost 1.7 million euros in total.


ZIB INCOME



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 2022, ZIB raised third-party funding through a large number of projects. Project-related public third-party funds increased from 6.8

million euros in 2021 to more than 7.7 million euros in 2022. Funding resulting from industrial third-party projects also increased from 1.3 million euros to 1.7 million euros. In total, more than 9.7 million euros in third-party funding marked an increase of almost 20% compared to 2021 – a new all-time record at ZIB.

ZIB THIRD-PARTY FUNDS BY SOURCE




31% BMBF
incl. FC Modal


18%
Industry

30% DFG
incl. Math+

21%
Other Public Funds

ZIB THIRD-PARTY FUNDS IN EUROS





■ Industry
■ Public Funds

Spin-Offs

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

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

Lenné 3D GmbH

2005 | www.lenne3d.com

3D 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

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

Exazyme

2021 | www.exazyme.com

Exazyme provides AI solutions for protein engineering

Number of Employees

In 2022, 242 people were employed at ZIB; of these, 158 positions were financed by third-party funds. The number of employees grew in comparison to 2021, mainly because of newly starting third-party funded projects. At the end of 2022, the number of members in permanent positions is almost 9% less than at the end of 2021 but this is a statistical artefact resulting from an unusually high number of retirements towards the end of the year.


1/1/2022			1/1/2023			
2	0	2*	3	0	3	MANAGEMENT
31	89	120	29	97	126	SCIENTISTS
43	3	46	40	3	43	SERVICE PERSONNEL
14	2	16	12	5	17	KOBV HEADQUARTERS
0	52	52	0	53	53	STUDENTS
90	146	236	84	158	242	Total
Permanent	Temporary	Total	Permanent	Temporary	Total	



TAKING MATH TO THE HEART

Numerical Simulation and Data
Analysis Foster Physiological
Understanding and Support
Diagnosis and Therapies

Diagnosis and therapy of cardiovascular diseases benefit from numerical simulation of cardiac excitation and the data analysis of resulting electrocardiograms. Research at ZIB contributes efficient numerical algorithms for simulation and parameter identification as well as model-free methods for the analysis of recorded data.



Math Under the Hood

Cardiac Electrophysiology

During the lifetime of a human, the heart beats about three billion times, reliably pumping blood through the body every day. If it does not, the personal consequences, as well as those for our societies' medical systems, are serious. Cardiovascular diseases are the most frequent reason for death in developed countries. The understanding of disease mechanisms, onset, and progression, as well as diagnosis and individual therapy planning, are of the utmost importance.

Malfunctions of the excitation conduction system are an important class of heart diseases, and can lead to arrhythmia, tachycardia, heart failure, and cardiac arrest. Diagnoses based on acquired data such as electrocardiograms and on corresponding parameter identification benefit from simulation and data analysis, e.g., for the detection of scars and infarction. Similarly, therapies like ablation, resynchronization, pacemakers, and defibrillation benefit from simulation of cardiac excitation for prediction and identification.



Tim Conrad
Phone +49 30 841 85-250
conrad@zib.de



Martin Weiser
Phone +49 30 841 85-170
weiser@zib.de

Simulation Model Hierarchy

Electrical excitation of the myocardium, ultimately leading to Ca^{2+} release and contraction, propagates due to a complex interplay of ion diffusion within cells and in the extracellular space on the one hand, and ion transport across the cell membranes on the other. The latter is described by ordinary differential equations (ODEs) for ion concentrations and the state of ion channels (Fig. 1). As the human heart comprises about a billion myocytes, modeling individual cells is extremely compute-intensive. Homogenized models averaging over several cells are therefore ubiquitous. The standard bidomain model results in a reaction-diffusion equation coupled to pointwise ODEs and an elliptic constraint, which is eliminated in the simpler monodomain model.

The solutions exhibit a quite narrow moving depolarization front of activation, requiring high spatial resolution and small time steps for a faithful numerical simulation, and leading to a high computational effort calling for efficient algorithms. Research at ZIB focuses on designing highly efficient numerical schemes for simulation and identification.

Several mechanisms and impact factors, however, in particular related to substrate heterogeneity, myocyte distribution, and their geometrical arrangement and connectivity as shown in Fig. 2, can only be studied when resolving the cells individually. Despite the tremendous computational effort induced, calling for high-performance computing in the exascale range, such models like the extracellular-membrane-intracellular (EMI) model gain attention. Again, ZIB research contributes efficient solution methods for these challenging problems.




Figure 1: Computed activation times color-coded on a ventricular geometry.




Figure 2: A 3D geometry of 48 myocytes and their interconnection.

Efficient Adaptivity

Locality of Features

Solutions of electrophysiology models exhibit strong local features, in particular the traveling activation fronts (Fig. 3). Away from these fronts, the dynamics are rather slow and smooth, and could be captured with much coarser spatial grids and longer time steps. Saving computational work by local refinement of finite element meshes and time discretization is complicated by the fact that the fronts are moving in space. Established mesh adaptivity turned out to be inefficient due to the incurred overhead of error estimation, mesh modification, and frequent reassembly of mass and stiffness matrices, such that most

simulations are performed on fixed uniform grids with constant time steps.

Within the Microcard project, co-funded by EuroHPC and BMBF, a novel route to space-time adaptivity has been developed at ZIB by moving adaptivity completely to the algebraic level, thus eliminating most of the overhead. The method combines higher order integration with spectral deferred correction methods and a cheap degree of freedom subset selection, and thus realizes a multirate-integration type approach [6,12].




Figure 3:
Excitation propagation in a cell-by-cell micromodel of the myocardium.

Spectral Deferred Corrections




Figure 4: Quantitative support of SDC corrections for an elliptic excitation front spreading out. Left: first iteration. Right: twelfth iteration. Not only are the corrections restricted to a neighborhood of the front, but their support is shrinking over the SDC iteration.

The approach is based on higher order integration with spectral deferred correction (SDC) methods, a stationary iteration for collocation time-stepping using low order integrators as preconditioners. Being iterative solvers, SDC methods interact well with all kinds of perturbations usually introduced by inexact and adaptive schemes, such as operator splitting, truncation of solvers for implicit schemes, and spatial adaptivity [13,14]. SDC methods converge with reasonable contraction factors $\rho \approx 0.2$, and gain one convergence order in time per iteration.

Algebraic Adaptivity

Applying SDC to electrophysiology models, one observes that far from the front, where the dynamics are slow, the corrections contributed by later SDC iterations essentially vanish, as shown in Fig. 4. The core idea is therefore to restrict the later iterations to a spatial region close to the front, and to perform this subdomain selection on the algebraic level based on local error estimates for the SDC convergence. Thus, the systems to be solved in later iterations are significantly smaller, reducing the computational effort (Fig. 5). Moreover, this adaptivity can be done with negligible overhead, and leads to moderate but significant speedups between 2 and 4. This has been demonstrated both on established monodomain models and on recent EMI models resolving individual myocytes.

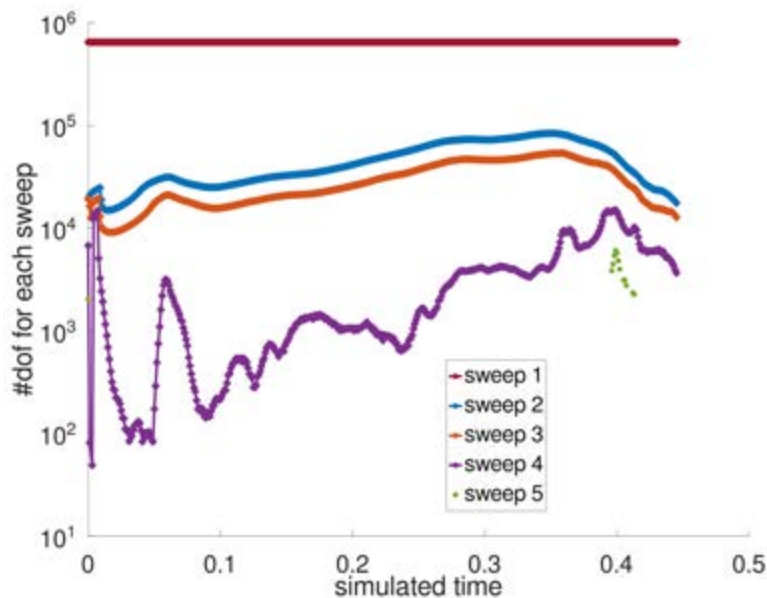


Figure 5: Number of active degrees of freedom used in the SDC iterations versus simulated time in a 3D monodomain problem.

Improving the Condition

Stiff Systems

Due to the diffusion of ions, electrophysiology models are stiff. In monodomain equations with the usual ratio of time-step and mesh width, the stiffness is only moderate due to fast reactions. In bidomain and cell-by-cell models, electric potentials are represented explicitly, which introduces algebraic conditions of elliptic type and renders the models

infinitely stiff. This makes implicit time-stepping schemes mandatory. Due to the problem size, direct elimination methods are a poor choice for solving the arising linear equation systems. Instead, iterative solvers such as preconditioned conjugate gradients with optimal, mesh-independent preconditioners are the methods of choice.

Substructuring Preconditioners

Non-overlapping domain decomposition preconditioners are known to be efficient and highly effective for elliptic equations and scale to thousands of nodes in high-performance computing systems. In collaboration with partners from Pavia and Milan, ZIB researchers have investigated balanced domain decomposition by constraints (BDDC) preconditioners tailored towards the special problem structure arising from EMI models [9]. In particular, the new preconditioner respects the geometrical partitioning of the

domain into natural subdomains given by the myocytes and the extracellular space, and splits the domain along the membranes. The degrees of freedom located on the membranes are duplicated and taken into account on both sides of the subdomain partitioning. With a standard averaging, we were able to show the usual condition bound independent of the number of subdomains also for this problem type.




Figure 6: BDDC preconditioners solve elliptic problems, here a Poisson equation, exactly on a broken space by enforcing continuity only at, e.g., the corners of subdomains. Projecting this back onto the continuous ansatz space in an energy-minimal way yields a very good approximate solution.

Mesh-Independent Convergence

On preliminary 2D computations, see Fig. 7, the new BDDC preconditioner proved to be optimal, with iteration count bounded independently of the number of myocytes and of the time-step size, and with a weak logarithmic dependence on the mesh width. The condition number is not only independent, but also absolutely rather small with values below 20; see Fig. 7. In comparison, simpler preconditioners such as using an unpreconditioned Schur complement system, or using no preconditioners at all, suffer from a clearly growing and much larger condition number in the order of 10^5 .




Figure 7: A snapshot of excitation propagation in a rectangular domain consisting of 14 by 14 myocytes surrounded by a fringe of extracellular space. The electrical potentials are color-coded. Excitation was started in the bottom left corner and spreads to the right. Clearly visible is the sign change of transmembrane voltage, i.e., the potential difference between intra- and extracellular domain, at the activation front in the middle of the domain.




Figure 8: Compared to simpler preconditioners, the new BDDC preconditioner for EMI models requires significantly fewer CG iterations.

Identifying Scars

Inverse Problem and Scar Modeling

Detection and quantification of myocardial scars are helpful for the diagnosis of heart diseases and for personalized simulation models. Scar tissue is generally characterized by different excitation propagation, and in particular by reduced conductivity. We aim at estimating distributed conductivity-related parameters from endocardial mapping data as shown in Fig. 9. Computing the maximum posterior point estimate in a Bayesian setting leads to the optimization problem of minimizing the misfit between simulated and measured electrical activity on the endocar-

dial surface, subject to the monodomain model. The prior knowledge of scar tissue electrophysiology enters as a regularization term. The parameter reconstruction depends crucially on the amount of acquired data (Fig. 10).

Solving this optimization problem requires many computationally expensive monodomain simulations on fine discretizations. ZIB researchers designed novel heterogeneous multilevel methods for accelerating the estimation process.




Figure 9a: Identification of scar tissue (shown in blue) within a ventricular geometry. This is the “ground truth” used for simulating measured endocardial data.

Figure 9b: Simulated activation times at a discrete set of endocardial measurement locations. The color shows the local activation time (in ms). From this data, the scar region is to be reconstructed.




Figure 10a: Cutout of the ventricular geometry with data positions and scar region marked.

10b: Distributed conductivity identified from the provided activation time data as a marker for scar tissue. The identifiability of scars depend crucially on the number and position of the data points.

Multilevel Identification of Scars

Eikonal models of electrophysiology are much faster to solve, but scar reconstructions based on eikonal and monodomain models, respectively, differ significantly in several situations, suggesting the use of the more accurate monodomain model for this purpose.

We considered grid hierarchies and monodomain-eikonal model hierarchies in a recursive multilevel trust-region method [7,8]. For computing a maximum posterior estimate, we investigated different optimization approaches based on adjoint gradient computa-

tion: steepest descent, limited memory BFGS [2], and recursive multilevel trust region methods. The multilevel approach accelerates identification considerably, enabling the use of complex electrophysiology models in estimating myocardial scars if sufficient data is acquired [3,4]. While grid hierarchies turned out to be more effective in reducing the iteration count, model hierarchies show a promising reduction of computing time due to fast eikonal models [5,3], as shown in Fig. 11 and 12.

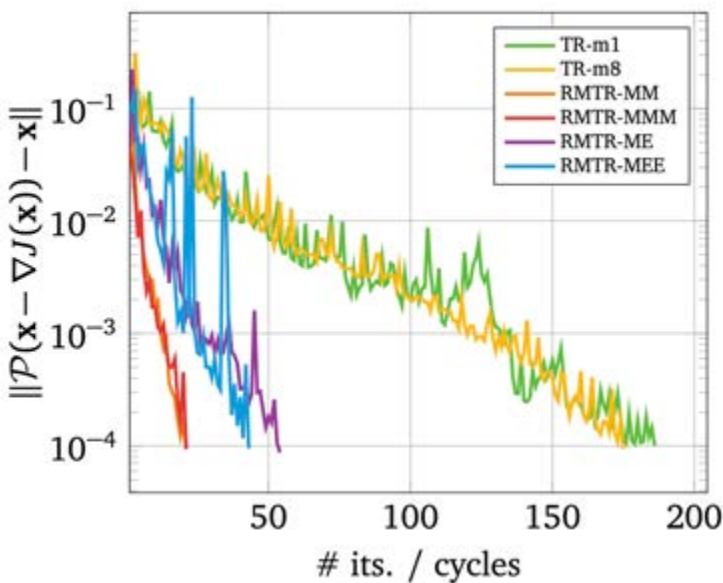


Figure 11: Convergence of different identification algorithms in terms of the projected gradient. Single-level methods (green/yellow) require considerably more iterations than multilevel ones.

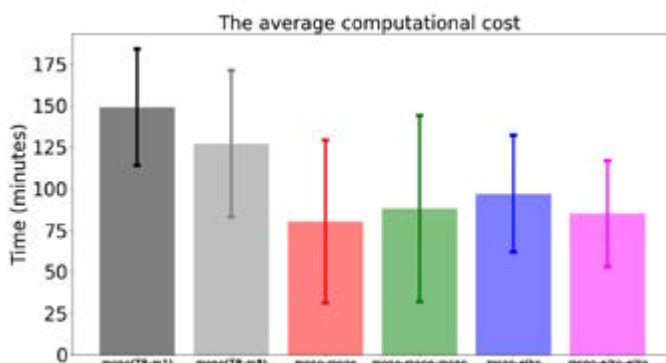


Figure 12: Wall-clock time for identification of distributed conductivity. Multilevel optimization methods show a clear benefit over single-level optimization methods (gray). Heterogeneous model hierarchies show a performance comparable to mesh hierarchies.




Figure 13: An electrocardiogram (ECG) interpreted by a medical doctor to diagnose heart-related diseases.

Reading the Heart – How Regular Are You?

A healthy heart beats about 60 times per minute, in a quite regular fashion, from when we are born. Abnormalities in this cardiac rhythm, such as extra beats or a too slow or too fast rhythm can be caused not only by a heart disease but also by stressful lifestyles or long-term mental stress. One way of observing the way a heart beats in detail over time is to collect data about the electrical activity

associated with the cardiac cycle. This can be done using a high-resolution electrocardiogram (ECG) recorder or a wearable device, which can be built into a modern smartwatch or implanted next to the heart for permanent monitoring of high-risk patients. The resulting ECG can be interpreted by trained medical staff to diagnose heart-related diseases.

Wearable Devices: Omnipresent but Hard to Analyze

Compared to the high-resolution medical devices, wearable ECG devices are of lower resolution, contain more noise, and overall less information, resulting in more difficult diagnostics. Due to the wide distribution of these devices on the mass market, there is an ever-increasing number of ECG episodes that need to be analyzed. Therefore, there is a growing need to assist the physicians with the interpretation of ECG recordings through automated algorithms.

Current state-of-the-art Deep Learning-based methods have attained cardiologist-level classification performance. However, these types of algorithms require very large human-annotated training sets to learn the intended task. Most available ECG data sets from wearables are often much smaller than needed and additionally of rather low quality, which makes it difficult to achieve a desirable performance level and calls for new approaches.

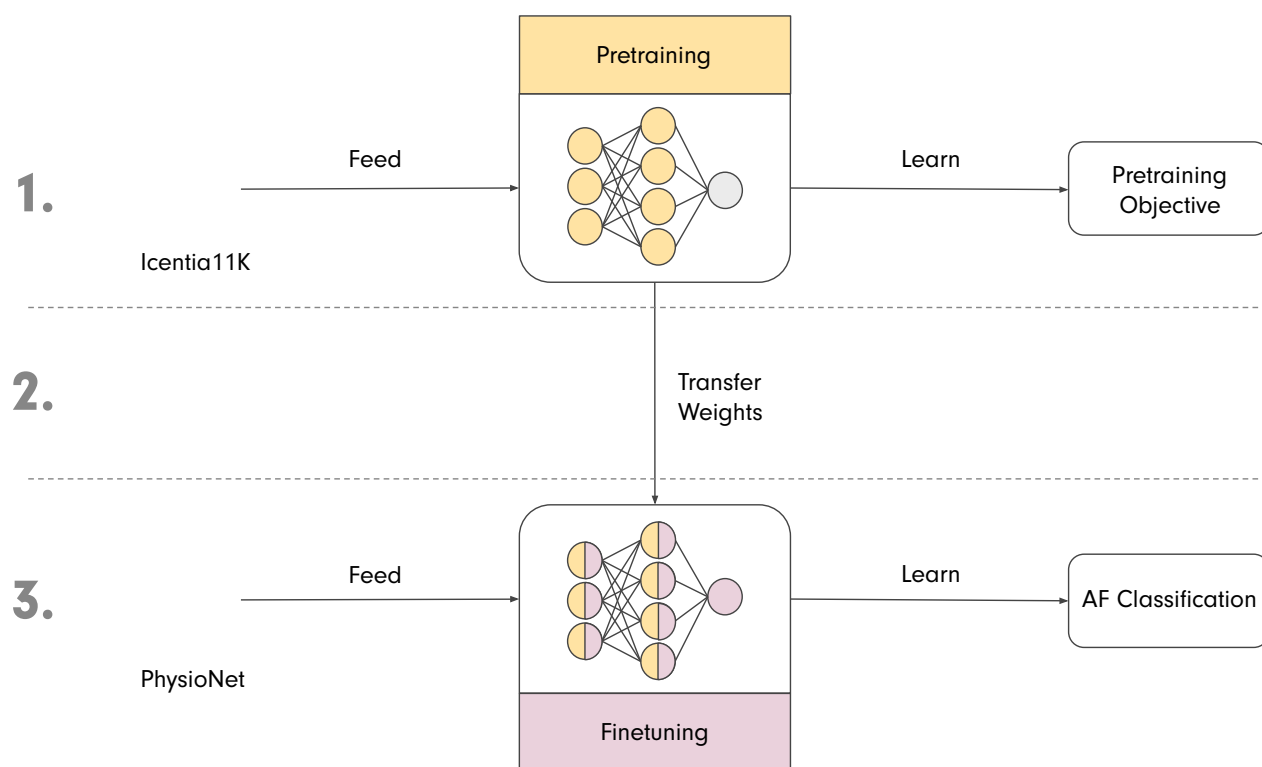


Figure 14: Visualization of the three main steps of the new method: (1) deep convolutional neural network (CNN) is pre-trained on the Icentia11K data set for a selected pretraining objective, e.g., classification of heart rate; (2) the pre-trained weights are used as initial weights of a new CNN; (3) this CNN is fine-tuned on the PhysioNet/CinC Challenge 2017 data set for a particular disease class (here: Atrial Fibrillation (AF)).

What if the Data Is Not Good Enough?

Researchers at ZIB have developed a new method [11] for the classification of ECG data within the BMBF-funded project Forschungscampus MODAL. This new method can deal with this situation where the data quality is low and no large annotated data sets are available.

The new method is based on three main steps which are shown in Fig. 14. First, we pretrain deep convolutional neural networks (CNN) on the Icentia11K data set, the largest data set available to date with 630,000 h of ECG signals with over 2,700,000,000 labeled beats. Next, we fine-tune the pretrained CNNs on the PhysioNet/CinC 2017 data set using available annotated ECG

episodes. Our main contribution is a successful large-scale pre-training of CNNs on the largest public ECG data set to date. By pretraining CNNs, we improve their performance on the target task and effectively reduce the number of expensive annotations required to achieve the same performance level as CNNs that are not pretrained (see Fig. 15). Our pretraining methods are robust to changes in the properties of ECG signals and can be applied to different models and data sets. Notably, we show how a CNN, which was pretrained on single lead ECG data typical for a wearable device, can be fine-tuned on 12 lead ECG data from a typical clinical device.

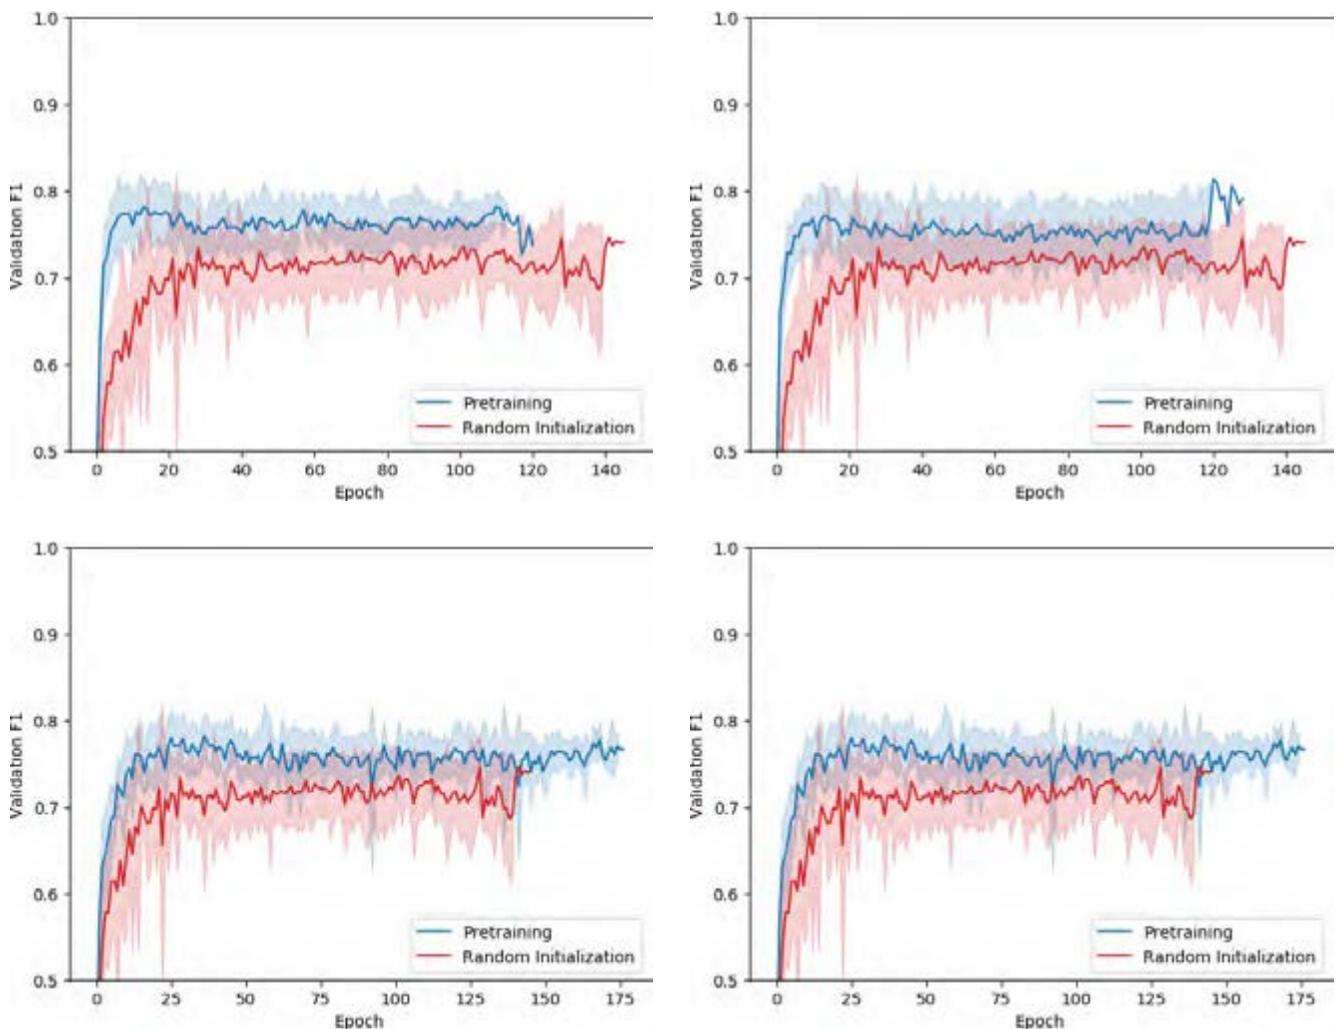


Figure 15: Average macro F1 on the validation set during fine-tuning. At each epoch, the average validation score of a selected pretraining method (blue) and random weight initialization (red) is reported. Pretraining improves validation performance on the downstream task and accelerates the training.

Outlook: More Data for a More Detailed Understanding

In the future, we might see even more diagnostic methods to describe the status of the heart that are based on so-called omics data sources, such as single-cell transcriptomics data. Researchers at ZIB have developed a new method [10] that can classify the evolution of single cells connected to heart status that can be harvested from the peripheral blood stream (e.g., immune cells, see [1]). Based on this data, dynamical changes in the molecular (RNA) constitution of these cells can be dynamically monitored and might help to identify upcoming cardiovascular diseases in very early stages.

[illegible]

QUANTUM TECHNOLOGY



Numerical Methods and Models for Quantum Computing and Quantum Communications

Quantum technology is an emerging field of physics which is expected to yield disruptive and revolutionary new engineering applications. This is possible due to the exploitation of quantum mechanical properties of matter which so far could not be used technologically. Examples of envisioned applications are in quantum computation, quantum metrology, and quantum communications. While quantum computers are expected to drastically outperform current methods for certain computation tasks, like factorizing large numbers, quantum sensors in principle allow for the highest sensitivities, and quantum communication systems allow to transmit information without any chance for unwanted eavesdropping.

Several groups at ZIB develop numerical methods for modeling, simulating, and optimizing quantum technology applications. This includes a research group on Quantum Computation and Optimization (QOPT) which was newly

established in 2022 and which focuses on quantum algorithms and quantum-inspired methods for combinatorial optimization and supervised learning, as well as the Computational Nano Optics (CNO) group.




Patrick Gelß
Phone +49 30 841 85-285
gelss@zib.de



Sven Burger
Phone +49 30 841 85-302
burger@zib.de


Figure 1: Artistic view of an experimentally realized plug & play telecom wavelength single photon source for quantum key distribution applications [1]. The nanophotonic setup of the single photon source has been numerically optimized by using methods developed at ZIB. Single photon sources are essential components of quantum communication systems, quantum sensing systems, and of quantum computing networks.



Tensor-Based Simulation of Quantum Circuits

Despite recent advances indicating that various optimization problems can be solved efficiently on quantum computers, developing new quantum algorithms that go beyond the capabilities of classical computers remains a complex challenge. The fundamental question is which classical mathematical problems can be translated into a quantum circuit framework.

In collaboration between Heriot-Watt University, TU Berlin, and ZIB, researchers from the QOPT group recently described how to use so-called matrix product states and operators – also known under the name tensor-train (TT) format – to represent not only quantum states but also quantum gates and entire quantum circuits [2]. It was shown that tensor decompositions can be used to represent complex



quantum circuits acting on exponentially large tensor product spaces and simulate them more efficiently than conventional techniques under certain conditions which is critical for understanding the computational advantages of quantum computing. Fig. 2 shows a form of tensor network typically used for sampling quantum measurements from a wave function given in TT format.

As the number of potential application areas as well as the complexity of quantum algorithms are rapidly increasing, there is a tremendous need for improved methods for simulating quantum systems. The overall goal is to design, extend, and implement tensor-based methods for quantum simulation in order to facilitate and accelerate the development of new methods.




Figure 2: Generative sampling in TT format. In graphical notation, the components of a tensor train are depicted by circles with different arms indicating the set of dimensions. Here, conditional probability distributions are constructed from segments of the tensor network. By successively drawing the bit values from the conditional probabilities, we generate samples according to the probability distribution given by the wave function.

Transfer of Dynamical Systems to Quantum Systems

One of the most important tasks of applied mathematics is the numerical simulation of dynamical systems. This problem class includes the prediction of the spread of infections or the modeling of planet trajectories. In fact, a wide variety of mathematical algorithms can also be interpreted as a discretization of an ordinary differential equation. One approach to this problem class is by Koopman and Perron-Frobenius operators. Here, the trajectories are identified with time-dependent probability measures over the domain of all possible configurations. This establishes a shift from ordinary to partial differential equations.

The Koopman-von Neumann approach takes this view as a starting point and derives a Schrödinger equation involving a Hamiltonian that is tailored to the target dynamical system. The probability density associated with the generated wave function is then the solution of the Perron-Frobenius

equation. For most dynamical systems one is interested in the development of a particular starting value for every time, but with regard to optimization one is only interested in accumulation points. Hence, the use of an a priori distribution over a wide range of starting values is of interest. In combination with the transition to partial differential equations, this procedure is more expensive in classical computing. As quantum computers can exploit superpositions, this approach might be advantageous.

The QOPT group set one of its research focuses on the analysis of the Koopman von Neumann framework, aiming at both a firm theoretical understanding and the realization on a quantum computer. In the course of our research, various techniques from the numerical discretization of PDEs with particular emphasis on Trotter-type product formulas will be of high relevance.




Figure 3:

Approximation of the evolved wave function associated with Koopman-von Neumann mechanics for the Himmelblau potential computed by FEM with a resolution of 2^7 .


Bell Nonlocality

In a famous article published in 1935, Einstein, Podolsky, and Rosen brought out the seemingly paradoxical feature that quantum theory allows correlations to happen at an arbitrary distance. Thirty years later, in 1964, Bell phrased this contradiction in a more quantitative way, paving the way to experimental demonstrations of this counterintuitive phenomenon. The Nobel Prize in Physics 2022 was then awarded to researchers having conducted these implementations.

In a Bell scenario, Alice and Bob are asked simultaneous questions and try to correlate their answers, for instance by giving the same reply if they were asked the same question. If they can have access to entangled quantum states, they can outperform any classical strategy, which is referred to as quantum nonlocality. In a sense, the precise relationship between entanglement and nonlocality is, however, not completely solved: In the simplest case of two-qubit states and for a simple class of symmetric states, the entanglement threshold is simple to compute, whereas the nonlocality one is still unknown after several decades of research. We take advantage of the recent developments of optimization algorithms (Frank-Wolfe and Quadratic Unconstrained Binary Optimization) to improve on the best bounds known on this problem.

The method is fully general and can be applied to any quantum state as well as more general scenarios, e.g., involving more than two parties.

Figure 4: Representation of a set of 2,161 quantum measurements in the Bloch sphere which approximate the full set of all projective measurements with a 99.94% precision. We use such approximations to construct our local model and hence our bound.



Numerical Tools for Optimizing Modal Fields in Nanophotonic Resonators

By miniaturizing optical setups, experiments approach so-called nanophotonic regimes. In such setups, typical dimensions are smaller or of the order of the wavelength of light. This results, e.g., in a clear separation of the discrete energy eigenstates of resonant structures. These eigenstates relate to mode fields of electromagnetic energy with specific eigenfrequencies which are complex due to the coupling to an unbounded exterior.

Theoretical and numerical models of light-matter interaction are needed to understand and optimize the performance of such nanophotonic devices. We have recently developed methods for solving nonlinear eigenvalue problems, which also include the computation of sensitivities by including direct differentiation methods in the contour integration-based algorithm [3], and we have developed methods for exact resonance state expansions, so-called Riesz-projection expansions, which we have also made accessible through an open source software package [4]. Further recent development of numerical methods includes advancements of Bayesian optimization algorithms [5,6].

In collaboration with Yuri Kivshar of the Australian National University we have used these methods for investigating and optimizing setups where a point source is coupled to specific resonance states, known as quasi-bound states in the continuum [7]. Fig. 5 shows a visualization of the energy distribution in a corresponding setup. We could show that the emission enhancement (known as the Purcell effect) in such devices is strong, when compared to other state-of-the-art nanophotonic resonators. This will allow for improved future generation quantum devices.




Figure 5: Visualization of a cylindrical nanophotonic resonator with a localized point-like light source which excites a bound state in the continuum resonance supported by the resonator [7].

Efficient Sources for Flying Q-Bits

Quantum networks for communication offer a possibility to transmit information without any chance of unwanted eavesdropping [1]. This offers a great advantage for the information society, and is in contrast to classical communication channels. In quantum communications, single photons at a well-defined quantum state, flying Q-bits, are used to transmit information. Nanophotonic resonances are used to boost the brightness of sources of such quantum states.

We have transferred our numerical methods to research groups at the Technical University of Denmark and at the Technical University Berlin, and collaborated with these in order to realize optimized numerical designs for miniaturized on-chip single-photon sources [8,9]. Within the Berlin Joint Laboratory for Optical Simulations in Energy Research (BerOSE) we have applied the developed mode expansion techniques for analyzing experimental results on the far field radiation pattern of quantum emitters in diamond nanoresonators [10]. Fig. 6 shows a microscopic image of such a device, where nitrogen vacancy centers are included in a thin layer in the upper part of the needle-like diamond crystal.




Figure 6: SEM micrograph of a diamond nanoresonator (top diameter 200 nm) realized experimentally at Helmholtz Zentrum Berlin [10]. Photon extraction and polarization state properties of the emitted photons are modified by the resonator structure.

Hot Carriers for Quantum Chemistry


Hot electron generation is a rising topic in solid state physics and quantum chemistry. A well-controllable excitation of single-charge carriers on surfaces and in bulk materials allows in principle the driving of chemical reactions and physical processes. Efficiency of such processes can be greatly enhanced by performing hot carrier generation on nanostructured surfaces. These modify and enhance both the local electromagnetic fields which excite hot electrons, and the rate at which these excitations can be extracted and supplied to the chemical reactants from the solid state. For modeling, simulation, and optimization of such systems we combine quantum formalisms and classical wave equations.

Plasmons are collective oscillations of electrons and they are solutions to classical electromagnetic wave equations. When plasmons are absorbed at metal surfaces, then individual electrons can escape from the surfaces. This is a quantum effect and the excited charge carriers are called hot electrons. Fig. 7 shows a schematic of hot electron generation on a structured metal surface. By generating hot electrons, the energy of visible light can be harvested and used in, e.g., photocatalytic processes.

We have studied theoretically the generation of hot electrons in dependence on the excitation of a plasmonic resonance [11]. A combination of a quantum model and the solutions of the classical Maxwell's equations was used to predict hot electron generation. We found that the generation rate was overestimated for very strong excitations

of the resonance. To address this, we developed and implemented a numerical method to correct the underlying material model. This resulted in a self-consistent approach for modeling hot electron generation in nanoresonators on metal surfaces.

In collaborations with researchers at University of Turin, Ohio University, Ludwig-Maximilians-Universität München, University of Bordeaux, and others, we have studied further theoretical aspects of hot electron generation [12,13], and we have applied our numerical approaches in order to analyze experimental results on plasmonics and hot carrier induced photochemistry [14,15,16,17,18]. Experimental findings in this context range from DNA-assembled nanoresonators to chiral light-matter interactions and hot-carrier induced photocatalysis.



Quantum Sensing

Quantum technologies are one of the most relevant contributors to innovation and advanced technologies. However, there is an important range of quantum-enhanced measurements that are not yet exploited by national metrology institutes because single-photon and entangled-photon sources with the required performance parameters are not readily available. Therefore, the collaborative project “Single- and entangled-photon sources for quantum metrology” (SEQUME) develops high-brightness, high-efficiency entangled-photon sources, and it will exploit these to demonstrate a quantum advantage for specific measurements. The European partners within this European Metrology Program for Innovation and Research (EMPIR) project are listed in Table 1.

ZIB is a partner institution in the SEQUME project, with specific tasks in modeling, simulation, and optimization of quantum light sources for metrology applications.

EUROPEAN SEQUME PROJECT
Physikalisch-Technische Bundesanstalt, Germany
Aalto-korkeakoulusäätiö sr, Finland
Cesky Metrologický Institut, Czech Republic
Dansk Fundamental Metrologi A/S, Denmark
Istituto Nazionale di Ricerca Metrologica, Italy
Justervesenet, Norway
AS Metrosert, Estonia
Türkiye Bilimsel ve Teknolojik Arastırma Kurumu, Turkey
Consiglio Nazionale delle Ricerche, Italy
Agencia Estatal Consejo Superior de Investigaciones Científicas, Spain
Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany
Istituto Nazionale di Fisica Nucleare, Italy
Keemilise ja Bioloogilise Füüsika Instituut, Estonia
Technische Universität Berlin, Germany
Universität des Saarlandes, Germany
Università degli Studi di Torino, Italy
Universität Stuttgart, Germany
Zuse Institute Berlin (ZIB), Germany
Table 1: Partner institutions of the European SEQUME project, which develops and investigates quantum technology infrastructure for quantum-sensing applications.

Figure 7: Schematic representation of hot-electron generation in a plasmonic nanoresonator [11]. A visualization of the energy distribution of a localized plasmonic resonance driving the carrier generation is overlaid.

HYDROGEN BLENDING



What is the capacity offered by the natural gas infrastructure?

The European energy infrastructure is being disruptively transformed to decarbonize. With a growing hydrogen market, pipeline-based transportation using the existing natural gas infrastructure becomes economically feasible, and accelerates the transition process. At Research Campus MODAL, we analyzed the maximum feasible injection of hydrogen with respect to regulatory limits of gas quality.



Ramping Up the Hydrogen Market

Hydrogen – a Prospective Renewable Energy Carrier

In the effort of decarbonizing the European energy system, it is necessary to mitigate the volatility of renewable energy sources, such as wind and solar power, and to prevent greenhouse gas emissions which are hard to reduce, e.g., from the industrial or heating sectors. To tackle these challenges, green hydrogen as a flexible, regenerative, and storable energy carrier is going to play a key role in the future energy system.

The German government has established the first official framework for the development of a broad hydrogen infrastructure with the National Hydrogen Strategy (NWS) 2020. The NWS forecasts a hydrogen demand of 90 to 110 TWh in 2030. However, an installed electrolysis capacity of 5 GW in 2030 and an assumed 4,000 full-load hours at an efficiency of 70% would lead to a production of only 14 TWh. Thus, Germany needs hydrogen imports to meet these demands [1]. Now, considering in particular the impact of past dependence on Russian natural gas, time is of the essence to stimulate hydrogen production, distribution, and offtake.



Janina Zittel
Phone +49 30 841 85-5146
zittel@zib.de



Jaap Pedersen
Phone +49 30 841 85-5437
pedersen@zib.de

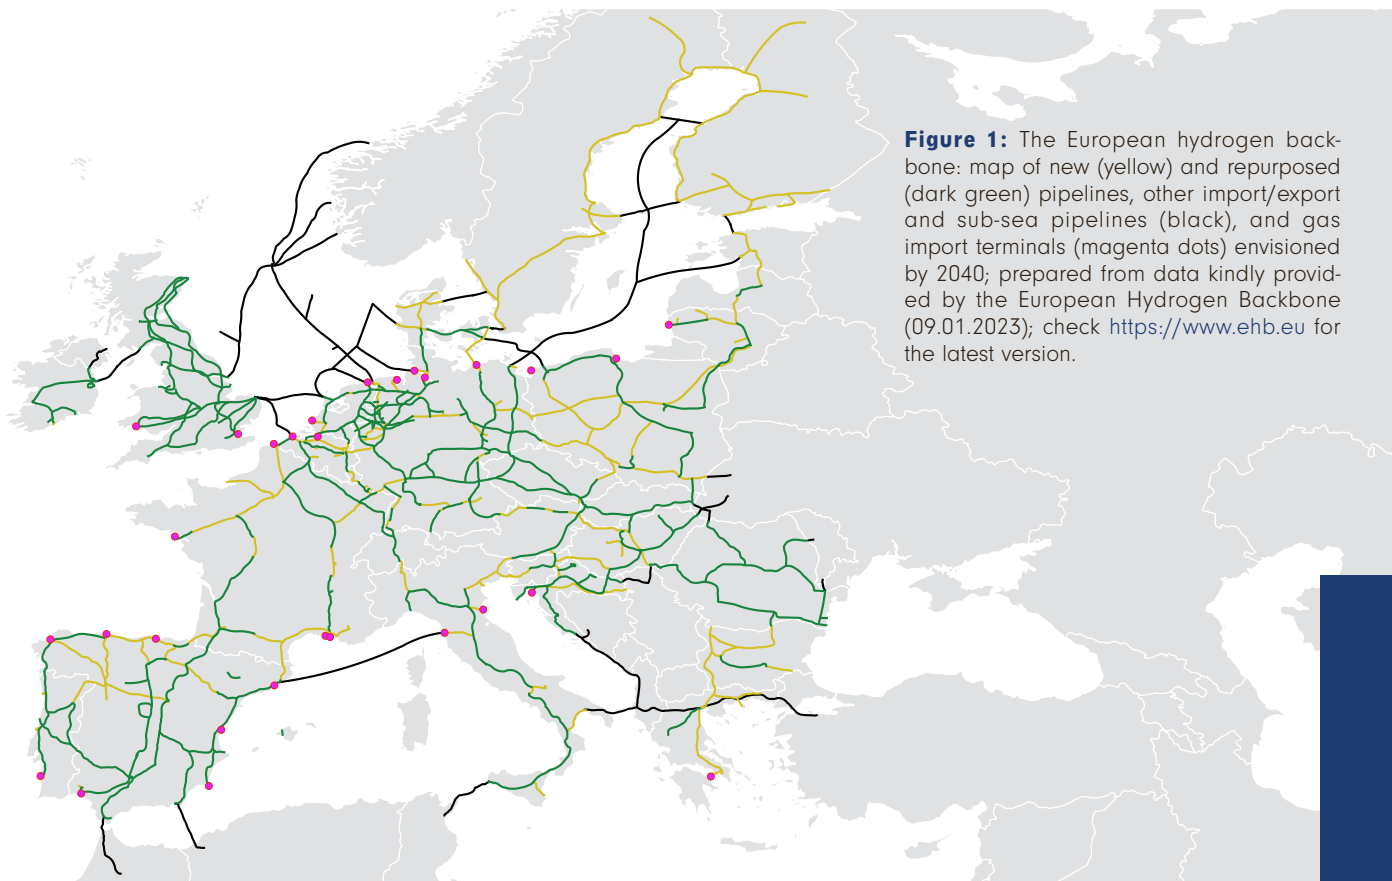


Figure 1: The European hydrogen backbone: map of new (yellow) and repurposed (dark green) pipelines, other import/export and sub-sea pipelines (black), and gas import terminals (magenta dots) envisioned by 2040; prepared from data kindly provided by the European Hydrogen Backbone (09.01.2023); check <https://www.ehb.eu> for the latest version.

Blending Hydrogen into Natural Gas Grids

Germany has a well-developed natural gas network, which will be gradually converted to transport hydrogen in the medium and longer term. In the European Hydrogen Backbone initiative, more than 30 energy system operators are advancing a European hydrogen network [2], and issue road maps about the construction of new pipelines and repurposing of existing gas infrastructure (Fig. 1). In the transition phase to a pure hydrogen infrastructure, blending hydrogen into the natural gas grid offers a guaranteed purchase for hydrogen, and can thus additionally motivate hydrogen production. At the same time, hydrogen blending entails technical and regulatory challenges. In particular, issues such as the embrittlement of pipelines, the more flammable nature of the gas mixture, and gas quality are primary concerns.

The MODAL EnergyLab Capacity Study

How Much Hydrogen Can the Natural Gas Grid Cope With?

Given both the large potential for hydrogen injection and its technical and regulatory challenges, it is critical to understand the capacities the gas infrastructure offers. The network's capacities are not only determined by its technical properties but also depend on the respective flow situation and chemical gas properties. This is because blending has technological and regulatory limits in terms of the volumetric ratio and the Wobbe Index.

In a case study on one of the largest European transmission grids, the EnergyLab at Research Campus MODAL investigated blending capacities based on the observed flow situation in 2022 [3]. The study addressed the following questions:

- How much hydrogen can be injected into the grid?
- How high is the storage capacity of the grid with regard to hydrogen?
- How is the hydrogen distributed in the grid?

Critical Limits




Figure 2: Wobbe Index of the hydrogen gas mixture – calculated daily at 2 p.m. based on the hydrogen injection in the scenario with the highest blending potential.

In the context of the MODAL EnergyLab hydrogen blending study, we define two categories of critical network elements, at which restrictions regarding hydrogen have to be observed: The first category consists of exits, the second of compressor stations:

1. Exits reflect the various customers and network connections for which different blending limits apply.
2. Compressor stations play a particularly important and central role in controlling the network and cannot easily be operated at higher H₂ levels.

Under current regulatory and technical framework conditions, a proportion of “biogases” – which also include hydrogen – of a maximum of 10 vol.-% is allowed, at natural gas filling stations even only up to 2 vol.-%. Another important criterion for gas quality is the Wobbe Index. For H-gas, this must be between 13.6 kWh/m³ and 15.4 kWh/m³.

Case Study for a Large German Network

Network and Data

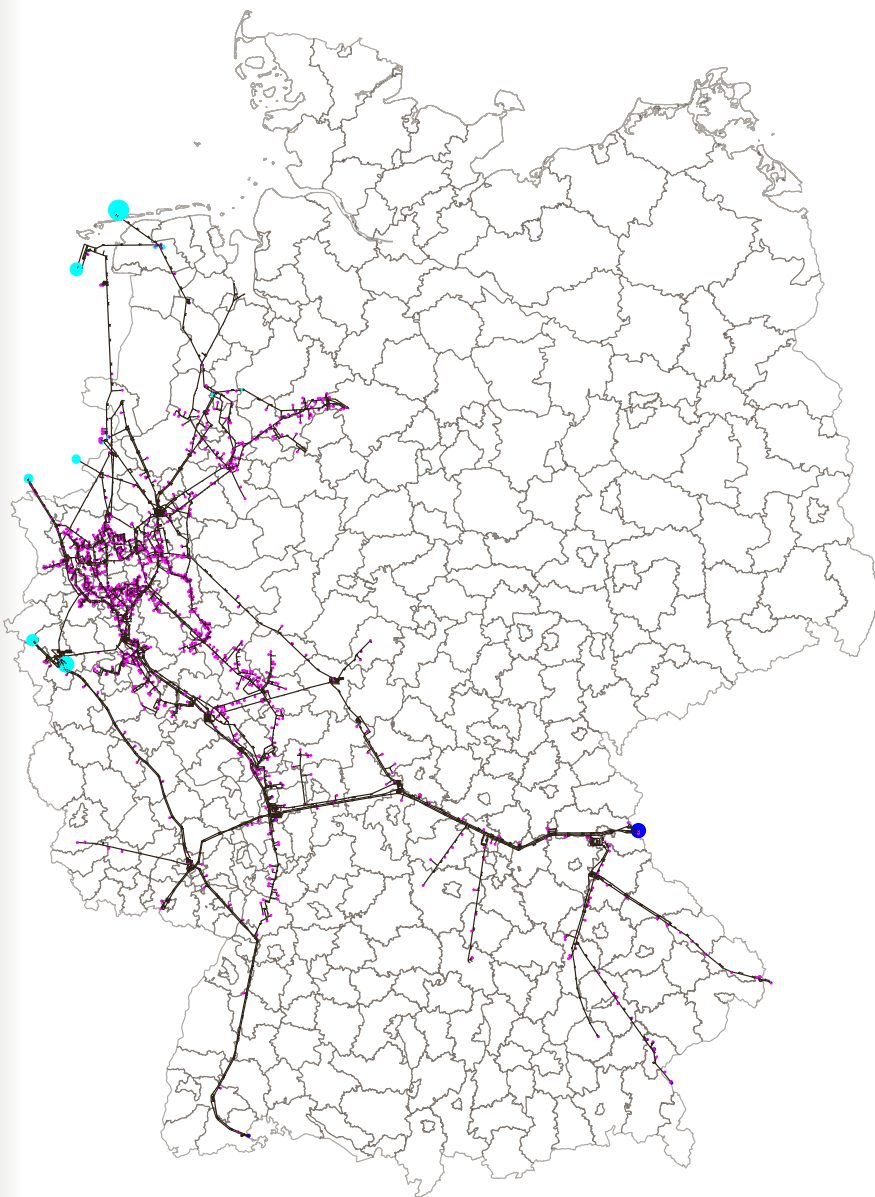


Figure 3: The Open Grid Europe network, study setup: hydrogen admixture at feed-in points in light blue, feed-in points without hydrogen injection in dark blue, exit points in magenta; size of entries relative to inflow during studied period.

Based on hourly measurement data of the transport network of our project partner Open Grid Europe GmbH (OGE), the potential for hydrogen admixture is investigated.

With a length of around 12,000 km and more than 50 entries, and over 1,100 exits, the OGE network is one of the largest in Europe (Fig. 3).

At 20 selected entries in the OGE network, it is assumed that hydrogen can be added to the natural gas stream. The selection was made in coordination with OGE and includes the northwest German entries of the OGE network. This fits with the fact that, predominantly in this region, green gas projects are located [4], which act as an upstream source for hydrogen.

The data basis for this study is hourly recorded data for gas flow, pressure, temperature, standard density, and calorific value for all points, pipes, and elements such as compressor stations, valves, and regulators throughout the OGE network. The data cover the period Jan. 23, 2022 to Aug. 13, 2022.

The EnergyLab Hydrogen Propagation Model

Based on the historical measured pressures and temperatures, the amount of gas in each pipe is calculated using the thermal equation of state for ideal gases. To represent the gas mixture, the model solves a pooling problem. Mathematically, this is an optimization problem with a linear objective, maximized subject to linear restrictions [5].

In the model, the gas flows through the network from the entries to the exits. In addition, modeling includes the change

in linepack over time. Further constraints beyond the usual constraints of the pooling problem ensure that scenario-dependent limits at critical network elements are met.

Finally, the theoretical hydrogen injection is maximized. As a result, the variable assignment of a solution of the model reflects the maximum possible hydrogen feed-in as well as the hydrogen content in the gas mixture at each node and arc of the graph modeling the gas grid.

The Gas Grid as a Guaranteed Customer for Green Hydrogen from Germany

To account for the ongoing discussions on regulatory

Moreover, even with rigorous limits, the gas grid offers

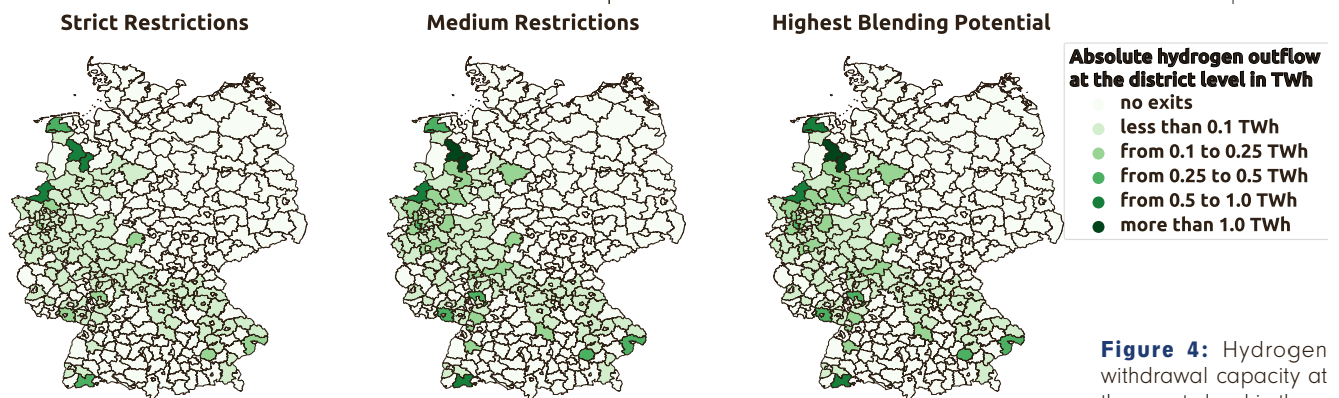


Figure 4: Hydrogen withdrawal capacity at the county level in three scenarios with different regulatory limits.

restrictions, the study looked at three scenarios, which assume a different level of restrictions on the volumetric ratio of hydrogen depending on the types of customers or network connections at the exit points. The results show that strict restrictions would still allow for 7.15 TWh of hydrogen injection from the end of January to mid-August 2022. This value increases up to 16.0 TWh for the most relaxed restrictions and 14.5 TWh for the medium restriction cases. In comparison, the NWS assumes a green hydrogen production capacity of 14 TWh by 2030. Thus, the calculated possible quantity for the latter two scenarios would be significantly higher than the quantity projected in the NWS.

Hydrogen withdrawal is largest at connection points, for example towards Switzerland or other German networks, due to large flow values and comparatively low restrictions. By contrast, the scenario-dependent limits of industrial customers, for example in the Ruhr area, explain the differences between the scenarios' total capacity (Fig. 4).

enough capacity to serve as a guaranteed customer for much of the green hydrogen production planned by 2030. Today's discourse tends to focus on using hydrogen in pure form. Large industrial sectors, such as the steel or chemical industries, will depend not only on the energetic but also on the material use of hydrogen, and will therefore act as large-scale consumers of hydrogen. Thus, as only a small part of the generation will be fed into the gas grid at all, no bottlenecks are to be expected here. The grid's capacity will increase significantly, assuming higher tolerances up to 10 vol.-% or 18 vol.-% hydrogen, as already tested at the distribution grid level.

Note that, on the other hand, the blending capacity decreases when less natural gas stays in the system as, over time, more and more customers replace gas as their energy source. Thus, grid transition planning needs to thoroughly track capacities for natural gas and pure hydrogen as well as gas mixtures over time to avoid bottlenecks.

Impact of the Unprecedented Gas Flow Situation in 2022 – Comparison to 2020




Figure 5: General gas inflow pattern change from 2020 (left) to 2022 (right).

A Fundamentally New Gas Import Structure in 2022

Due to Russia's attack on Ukraine and the consequent changes in Europe's gas import situation, we find a fundamentally changed flow pattern in our studied network. The flow via the Ukraine pipeline has dramatically decreased,

while North Sea gas takes a much larger share than before (Fig. 5). Most prominently, the import volume via Belgium is more than four times as high.

Impact on Hydrogen Injection Capacities

These changes in gas flow patterns also had a massive influence on hydrogen blending capacities. In fact, they substantially increased when compared to a study based on observed flows in 2020 [5].


The reasons are three-fold:

1. Since the volume fraction of hydrogen in the gas mixture is the determining factor for the capacity, the natural gas flow in the area of the feed-in points is relevant. With a preferred location in northern Germany, the injection potential increases with a stronger transport of North Sea gas.
2. North Sea gas also has a higher caloric value than Russian gas, which in turn is relevant for compliance with Wobbe Index limits for H-gas.
3. Finally, in the analyzed month in 2022, the total entry flow was also about 10% higher compared to the 2020 study.

In summary, extrapolations for a 365-day injection capacity reflect an increase of 43–52% from the first study [5] based on April 2020–March 2021 to the current from January 2022–August 2022 (Table 1).

Scenario	Extrapolated annual injection capacity (TWh)	
	2020	2022
Strict restrictions	9.0	12.9
Medium restrictions	17.4	26.2
Highest blending potential	18.9	28.9

Table 1: Comparison of extrapolated annual hydrogen injection capacity in TWh in three scenarios for 2020 and 2022.



The Long Shot: Towards a Pure Hydrogen Infrastructure

New Challenges

Finally, hydrogen blending is a bridging technology on the way to a pure hydrogen infrastructure. Such a pure hydrogen infrastructure brings new challenges for sustainable planning and safe and efficient operation. For example, due to the lower energy density of hydrogen compared to methane, the main component of natural gas, the volumetric flow needs to increase in order to transport the same amount of energy. Thus, the amount of energy need-

ed for compression significantly increases. In a feasibility study on pure hydrogen transport, we virtually repurposed a transmission grid in western Germany. The increase in compression energy was found to be close to 440% (Fig. 6) on average for scenarios based on historic flow situations [6]. At the same time, the control of the grid becomes more dynamic, and more network switches are needed.




Figure 6: Comparison of compression energy used for transport of natural gas (NG-TC) and hydrogen with the energy equivalent of NG-TC (H2-HC-EQ); based on data from [6].

Algorithmic Intelligence for Decision Support on Future Energy Systems

Based on the previous work on natural gas networks, the MODAL EnergyLab develops solutions that will assist decision makers in the future by means of algorithmic intelligence. Whether for network planning and capacity evaluations [7] or for network operations [8,9], mathematical optimization and machine learning [10] have proven to provide valuable decision support. The integration of new fuels, as prototypically shown here for hydrogen blending and pure hydrogen transport, demonstrates the usefulness of these methods for future networks as well. These decision support systems bridge the gap between gut feeling of system operators together with small-scale pilot engineering studies and large-scale energy system analysis, which entail an extremely simplified view of gas physics. Only the interplay of all these elements, together with landmark policy decisions, will enable the decarbonization of the energy system.



COMPUTER- ASSISTED PROOFS IN EXTREMAL COMBINATORICS

How to Challenge Erdős
Without Being Erdős

The limit of human intuition and ability to both execute and verify complex proofs has motivated the increasing formalization and automation of proof techniques. Following this direction, we used techniques from optimization to obtain novel results relating to some long-standing open problems in extremal combinatorics.

Introduction

Computers and artificial intelligence have always been an important tool for mathematicians to conjecture, prove, and verify mathematical statements. They were crucial in the formulation of the Birch and Swinnerton-Dyer conjecture in the 1960s, established the four-color theorem through proof-by-exhaustion in the 1970s, provided rigorous numerical estimates to prove a conjecture of Feigenbaum in the 1980s, contributed solutions to around 100,000 linear programming problems to a proof of the Kepler conjecture in the 1990s, and can even execute fully automated reasoning, for example to prove that Robbins algebras are boolean. Lately, even machine learning-based approaches have been proposed, giving new insights by formulating novel conjectures and claiming to be capable of answering university-level problems. Computer-assisted proofs can often involve significant computational challenges: Determining the Schur number $S(5)$, for example, required a massive two petabytes of storage.

The use of computer assistance has been particularly fruitful in extremal combinatorics, an area full of hard results and ingenious proofs that nevertheless seems to draw upon only a handful of essentially distinct techniques, at least when obtaining asymptotic results. According to Alexander Razborov [1]: “the difficulty (as well as beauty) lies in finding the ‘right’ relations using instruments from this relatively small toolbox.” Funded through Math+ projects EF1-12 and EF1-21, researchers at the Zuse Institute, partially in cooperation with the Graph Theory group at the Freie Universität Berlin, explored how ideas and techniques from the world of optimization can contribute to this toolbox. In the process, we used conic optimization and SOS hierarchies as well as tools from discrete optimization and operations research, while also exploring new approaches based on machine learning. As a result, we were able to obtain significant novel results relating to some long-standing open problems in extremal combinatorics.



Christoph Spiegel
Phone +49 30 841 85-209
spiegel@zib.de



Sebastian Pokutta
Phone +49 30 841 85-436
pokutta@zib.de

The Goal of Extremal Combinatorics

The field of extremal combinatorics is concerned with how large or small specific parameters of combinatorial objects, such as graphs, families of sets, and additive structures, can be while satisfying certain restrictions. One pillar of this is Ramsey theory, based on the British mathematician Frank P. Ramsey's well-known result from 1930 that any large-enough graph must contain either a large clique or a large independent set. A slightly more general form of this asserts

that any coloring of the edges of a large-enough complete graph with a finite number of colors must contain large monochromatic cliques. Determining precisely how large is large enough has occupied mathematicians for close to a century, with no clear answer in sight and results addressing both the case of specific sizes of the cliques as well as asymptotic statements.




Figure 1: Two colorings avoiding monochromatic cliques: The coloring on the left avoids monochromatic triangles in both colors while the coloring on the right avoids triangles in one and cliques of size four in the other color. Both examples are extremal in the sense that the edges of the next larger complete graphs cannot be colored while maintaining these properties.

A Problem of Erdős

We looked at the closely related Ramsey multiplicity problem, where given the fact that monochromatic cliques are unavoidable, we now ask precisely how many we need to have and what extremal graphs attaining the lowest possible value look like. Formally, for c colors and sizes of the cliques t_1, \dots, t_c , we are interested in determining

$$m_{t_1, \dots, t_c} = \lim_{n \rightarrow \infty} \max_{G \in \mathcal{G}_n} k_{t_1}(G_1) + \dots + k_{t_c}(G_c),$$

where \mathcal{G}_n denotes the set of all c -colorings of the edges of the complete graph on n vertices and $k_{t_i}(G_i)$ the fraction of t_i -cliques monochromatically colored with color i . For the case of triangles and two colors, that is $m_{3,3}$, this was resolved by Goodman in 1959 [3], establishing that asymptotically at least $1/4$ of them need to be monochromatic.

One way of seeing that extremal constructions attaining the bound of Goodman exist is by simply considering the

average behavior of colorings: If we color the edges independently and uniformly at random, each triangle has a likelihood of $1/4$ being monochromatic. This led to a famous conjecture of the Hungarian mathematician Paul Erdős in 1962 [4], stating that this should hold more generally whenever considering two colors and cliques of equal size, that is that $m_{t,t} = 2^{1-t(t-1)/2}$.

Surprisingly, this conjecture was disproven by Thomason [5] for any $t \geq 4$ and even the value of $m_{4,4}$ remains unresolved. Up until now, the only other Ramsey multiplicity value besides $m_{3,3}$ that had been resolved was obtained in 2016 by Cummings et al. [6], who showed that $m_{3,3,3} = 1/25$ while also giving a characterization of the extremal sequences of colorings. We give the best current upper and lower bounds for $m_{4,4}$ and $m_{5,5}$ while also settling the precise values of $m_{3,4}$, $m_{3,5}$ and $m_{3,3,3,3}$.

“Suppose aliens invade the earth and threaten to obliterate it in a year’s time unless human beings can find the Ramsey number for red five and blue five. We could marshal the world’s best minds and fastest computers, and within a year we could probably calculate the value. If the aliens demanded the Ramsey number for red six and blue six, however, we would have no choice but to launch a preemptive attack.”
– Paul Erdős as quoted in [2].




Figure 2: Many constructions in extremal combinatorics rely on so-called blow-up constructions that allow one to derive infinite sequences from finite objects with desirable properties. Here the 2-fold blowup of a triangle is illustrated.

Constructions beyond Human Intuition

Random constructions are out given that Erdős' conjecture was refuted, but explicit so-called blow-up constructions, replacing vertices with a certain number of copies while inheriting colors from the original coloring, have proven a powerful source of upper bounds, often resulting in tight results. Extending previous notation, we let \mathcal{G}_n^o denote the set of all c -colorings of the complete looped graph on n vertices and $k_{t_i}^o(G_i)$ the fraction of not necessarily in-

jective maps from the complete graph on vertices to the (possibly looped) graph given by color i in $G \in \mathcal{G}_n^o$ that are strong graph homomorphisms. The central statement that we will use to computationally derive upper bounds is the simple fact that, through the sequence of blow-ups, any coloring $G \in \mathcal{G}_n^o$ defines an upper bound through $m_{t_1, \dots, t_c} \leq k_{t_1}^o(G_1) + \dots + k_{t_c}^o(G_c)$.

Heuristics and Symmetries

The right-hand-side of the previous equation is easily computable even for large G and therefore allows one to formulate the search for constructive upper bounds as a simple combinatorial optimization problem for any order of the underlying graph G . This problem is amenable to a variety of approaches and in the past both by-hand as well as exhaustive computer searches over specific (small) constructions, such graphs defined through powerset relations or XOR graph products, have found use.

One underlying property of these optimization problems is that the quality of the obtained results takes precedence over any guarantee of optimality; since we a priori do not know that an optimal solution to the problem gives the best possible upper bound, such guarantees are of little value in this application. Instead, it is not uncommon that we already have derived a lower bound through some other means that we are trying to match and that already can provide a strong guarantee of optimality. We therefore found that using simple and well-established metaheuristics such as simulated annealing [7] and tabu search [8] to directly search for constructions to be a very successful approach.

Applying this for $n = 27$, we [9] were able to derive novel bounds that had so far eluded mathematicians, determining the first two-color Ramsey multiplicities since Goodman's result from 1959 by showing that $m_{3,4} = 689/3^8$ and $m_{3,5} = 24011/3^{12}$. In the former case we were also able to show the stability of extremal sequences with respect to the sequence of symmetric blow-ups based on the coloring defined by the Schläfli graph. Noting that the previous best constructions for $m_{t,t}$ all exhibit very specific symmetries, namely that the edge relations are dictated by the binary operation of specific groups, we also biased our search space to exclusively search for Cayley graphs. As the number of variables is decreased from quadratic to linear in the number of vertices of the underlying graph, this allowed us to explore significantly larger constructions and obtain the bounds $m_{4,4} \leq 0.03012$ and $m_{5,5} \leq 0.001707$ from the groups $C_3 \times C_2^{*8}$ and $C_3 \times C_2^{*6}$. Both values pose the first significant improvements for these well-studied problems in almost 30 years.

A Classic Tool Done Right

When studying existing by-hand approaches to problems like the Ramsey multiplicity, one inevitably encounters a relatively small set of existing techniques, in particular double-counting arguments in combination with an application of the Cauchy-Schwarz inequality. Formalizing these techniques and putting them to a common denominator had been a common goal for decades. Alexander Razborov developed the most consequential of these approaches by introducing flag algebras, an algebraic formulation of the limit behavior of combinatorial objects, and viewing existing proof techniques as linear operators and algebra homomorphisms acting between these objects.

For our purposes, a σ -flag will be a partially labeled coloring of the complete graph where the labeled part induces a particular type σ . We will use \emptyset to denote the empty type, i.e., the coloring of the empty graph. We

define the flag algebras $\mathcal{A}^\sigma = \mathbb{R}\mathcal{F}^\sigma / \mathcal{K}^\sigma$ by taking linear combinations of σ -flags, factoring out the relations of a chain rule through \mathcal{K}^σ and defining an appropriate product. The fundamental realization by Razborov [1], as well as previously Lovász and Szegedy [10], was that the limit densities in convergent sequences of coloring are in a one-to-one correspondence with algebra homomorphisms in $\text{Hom}^+(\mathcal{A}^\sigma, \mathbb{R}) = \{\phi \in \text{Hom}(\mathcal{A}^\sigma, \mathbb{R}) : \phi|_{\mathcal{K}^\sigma} \equiv 0\}$. This means that determining the maximum of the limit behavior of any linear combination of flags $f^* \in \mathcal{A}^\sigma$, e.g., the sum of all monochromatically colored triangles, can be equivalently stated as the optimization problem $\max\{\lambda \in \mathbb{R} : f^* - \lambda \emptyset \in \mathcal{S}\}$ over the semantic cone $\mathcal{S} = \{f \in \mathcal{A}^\sigma : \phi(f) \geq 0 \ \forall \ \phi \in \text{Hom}^+(\mathcal{A}^\sigma, \mathbb{R})\}$. Information from other algebras \mathcal{A}^σ where $\sigma \neq \emptyset$ gets mapped into \mathcal{S} through a downward operator.

$$\begin{aligned}
 \triangle_{\text{blue}} + \triangle_{\text{orange}} &= \frac{3}{2} \left(\left(\frac{1}{3} \triangle_{\text{blue}} + \triangle_{\text{blue}} \right) + \left(\frac{1}{3} \triangle_{\text{orange}} + \triangle_{\text{orange}} \right) - \frac{1}{3} \right) \\
 &= \frac{3}{2} \left(\left(\triangle_{\text{blue}} + \triangle_{\text{blue}} \right) + \left(\triangle_{\text{orange}} + \triangle_{\text{orange}} \right) - \frac{1}{3} \right) \rightarrow \frac{3}{2} \left(\bullet^2 + \bullet^2 - \frac{1}{3} \right) \\
 &\geq \frac{3}{2} \left(\bullet^2 + \left(1 - \bullet^2 \right)^2 - \frac{1}{3} \right) = 3 \left(\bullet - \frac{1}{2} \right)^2 + \frac{1}{4} \geq \frac{1}{4}.
 \end{aligned}$$

Figure 3: A symbolic representation of the proof of Goodman using flag algebras. It states that asymptotically at least a quarter of all triangles in any coloring of the complete graph must be monochromatic.

“If you are double-counting: do it right!”

– Alexander Razborov when introducing flag algebras

How to Leverage Symmetries

Optimizing over \mathcal{S} is hard, but one can approximate it through sum-of-squares (SOS) hierarchies. For each integer N , this results in the concrete semidefinite programming (SDP) formulations

$$\max_{Q \succeq 0} \min_{H \in \mathcal{G}_N} f_H^* - \langle Q, D_H \rangle$$

where f_H^* is the parameter f^* evaluated at H and the matrix D_H captures specific density computations of the flags in $H \in \mathcal{G}_N$. Increasing N generally improves the resulting bound while making the SDP harder to solve.

To push the envelope of results obtainable through this framework, we [11] described specific and easily computable symmetries in the SDP formulations that stem from both the invariances of the element f^* under permutations of the colors as well from the (generalized) automorphism groups of the types σ . Using these symmetries, we formulated two complementary approaches to significantly reduce

the difficulty of the SDP formulations while also improving their conditioning.

Method 1. Using the symmetries of the element f^* , we reduce the number of constraints and blocks in the SDP formulation by combining constraints stemming from colorings of the complete graph of order N that are isomorphic under color permutations.

Method 2. Using the symmetries of the types σ , we reduce the number of variables through block diagonalization [12] of the block in our SDP formulation associated with σ , i.e., we apply Schur’s lemma.

The combination of these two methods allowed us to achieve novel results that were previously unattainable using existing computational resources. In particular, we established that $m_{4,4} \geq 0.02961$ and $m_{5,5} \geq 0.001557$, complementing the upper bounds established in the previous section, and were able to show that $m_{3,3,3} = 1/256$.





Figure 4: Block diagonalization uses symmetries stemming from the underlying combinatorial problem to reduce the problem formulation by transforming large semidefinite matrices into a smaller block-diagonal form.

THE TROPICAL GEOMETRY OF PERIODIC TIMETABLING

Improving Public Transport using Geometric Insights

What rhythm is to music, the timetable is for a public transportation system. Creating timetables that maximize passenger satisfaction while guaranteeing operational practicality is key for the success of public transport. Analyzing the inherent geometry of periodic timetables allows for new cutting-edge timetable optimization algorithms.





Periodic Timetabling and Combinatorial Optimization

A high-quality timetable is an indispensable asset of an attractive public transport system and thus an important tool to increase the share of sustainable transport modes. Since timetables are often operated periodically, it is therefore desirable to determine periodic timetables with mathematical optimization methods.



Ralf Borndörfer
Phone +49 30 841 85-243
borndorfer@zib.de



Enrico Bortoletto
Phone +49 30 841 85-312
bortoletto@zib.de



Niels Lindner
Phone +49 30 841 85-374
lindner@zib.de



Berenike Masing
Phone +49 30 841 85-223
masing@zib.de

Periodic Event Scheduling Problem

The features of the periodic timetabling problem in public transport are typically expressed in the language of the Periodic Event Scheduling Problem (PESP) [15]. PESP is a combinatorial optimization problem based on a digraph, the event-activity network. In this network, vertices correspond to departure and arrival events of, say, trains at stations, while arcs model activities, e.g., driving between two adjacent stations or passenger transfers [8]. A periodic timetable is simply an assignment of timings to all events. For an activity, the difference of event timings at its endpoints fixes its duration modulo the period time. The aim is to find a periodic timetable that minimizes the (weighted) activity durations, e.g., the total travel time of all passengers.




Figure 1: Line network ("metro map") of three lines intersecting at a common station.


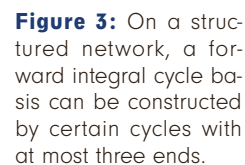


Figure 2: An event-activity network corresponding to the line network including driving, dwell, turn, and transfer activities.



Heuristics

On the primal side, several heuristics have been developed, such as the modulo network simplex algorithm [14], a generalization of the dual network simplex algorithm for minimum cost flows that runs in polynomial time, but gets stuck in local optima quite early. Although non-polynomial time, a variety of combinatorial optimization techniques have been successfully applied to PESP, e.g., maximum cut [2], graph partitioning [9], merging [11], or SAT approaches [5].



The Geometry of Periodic Timetables

Tensions

One standard approach to understand a mixed-integer linear program is to investigate the structure of the polyhedron generated by its feasible solutions. In the context of periodic timetabling, this boils down to an analysis of the periodic tension polytope, the convex hull of all feasible

activity durations. The modulo network simplex algorithm can be interpreted as a partial simplex algorithm on this polytope. Moreover, several families of cutting planes have been singled out; however their separation is typically NP-hard [1, 10].

Timetables

While most approaches focus on the tension polytope, the space of feasible periodic timetables has been neglected in previous research. It turns out that this space carries an interesting geometry: When event timings are allowed to be arbitrary real numbers, the space is tiled by a periodic pattern of polytopes. Moreover, these polytopes are not only convex in the ordinary sense, they are also the tropical convex hull of finitely many points in the world of tropical geometry. These objects are called polytropes [6], and they correspond to weighted digraph polyhedra, which in turn arise in the study of feasible potentials on a weighted digraph [7].

Taking into account the symmetry introduced by periodicity, one may as well embed the space of feasible periodic timetables into a torus, which is then still tiled by a hyper-plane arrangement whose cells are polytropes. As for the periodic tension polytope, the vertices of those cells correspond to spanning tree structures. Moreover, each polytrope cell can be embedded into the tension polytope, and the tension polytope is the convex hull of the embeddings of all polytropes [4].




Figure 5: The tension polytope of a PESP instance. The blue regions are the feasible activity durations, the facets stem from bounds on the activity durations, or from so-called flip inequalities.




Figure 6: Periodic tiling of the timetable space in the Euclidean plane by polytropes.

Cycle Offsets

From the viewpoint of the MIP formulation of PESP, the third player that can be analyzed, are the technical integer cycle offset variables. Projecting the fractional periodic tension polytope to the space of the cycle offset variables results in a zonotope, whose integer points actually correspond bijectively to the polytropes in the tiling of the periodic timetabling torus. Moreover, the tropical vertices of the polytropes are connected to zonotopal tilings – a remarkable duality [4].




Figure 7: Periodic timetabling torus and its tiling by polytropes.




Figure 8: Cycle offset zonotope with marked integer points.

Tropical Neighborhood Search

Idea

As the space of feasible periodic timetables decomposes into a disjoint union of polytropes, solving PESP restricted to each polytrope will solve the full instance. In fact, such a restricted problem is equivalent to a minimum cost uncapacitated network flow problem, and hence efficiently solvable. However, the number of polytropes is typically large: A sharp upper bound is the number of spanning trees of the event-activity network [4].

Nevertheless, this insight motivates a local search heuristic, called tropical neighborhood search [3]: The tiling of the periodic timetabling torus by polytropes induces a polyhedral complex. On this complex, two polytropes are neighbors if both are full-dimensional and share a common facet. It turns out that the number of neighbors of a given polytrope is linear, namely at most twice the number of activities. Moreover, the integer points of the cycle offset zonotope labeling two neighboring polytropes are related by a simple formula: They differ by a column of the cycle offset zonotope, i.e., the matrix that defines the cycle offset zonotope.




Figure 9: Polyhedral subdivision inducing the neighborhood relation on polytropes.

Algorithm

The tropical neighborhood search heuristic is an improving heuristic for the Periodic Event Scheduling Problem. Starting with an initial periodic timetable, it identifies the polytrope to which the timetable belongs, and then scans all neighbors for a timetable with a better objective value. If such a timetable has been found, then the procedure iterates using that timetable as an initial solution.

Advantages

Tropical neighborhood search runs in polynomial time and benefits from warm starts, as the minimum cost network flow problems associated to two neighbors differ only in their cost functions. Moreover, the search is based on polytropes, not only on vertices, and has therefore a wider scope than the modulo network simplex method. In general, the search space of these two heuristics is different. In practice, tropical neighborhood search is often able to escape local optima for the modulo network simplex algorithm.




Figure 10: Comparison of the search spaces of tropical neighborhood search (left) and modulo network simplex (right). Squares are non-global local optima, red means better, and purple means optimal.

ConcurrentPESP

A State-of-the-Art Solver

Since 2018, the MobilityLab group at Zuse Institute Berlin has been developing ConcurrentPESP, a software for optimizing instances of the Periodic Event Scheduling Problem [2]. This solver has served as a base for various research on periodic timetabling algorithms [2, 3, 9, 10, 11, 12], and computed the current incumbent primal and dual bounds for all 22 instances of the benchmarking library PESPlib¹. This library is composed of real-world railway and bus timetabling instances of medium to large size.

The setup of the solver is concurrent, so that several periodic timetabling algorithms can run in parallel and communicate solutions with each other. To this end, the ConcurrentPESP solver combines – among other methods – branch and cut via interfaces to MIP solvers, and implementations of the modulo network simplex heuristic and tropical neighborhood search.




Figure 11: Relative improvement of an initial solution by the heuristics implemented in ConcurrentPESP.

¹ <http://num.math.uni-goettingen.de/~m.goerigk/pesplib/>

Performance of Tropical Neighborhood Search

Tropical neighborhood search is available with several configuration options within ConcurrentPESP. For example, one parameter determines whether all potential neighbors should be explored, or only those that are neighboring the current optimal polytrope vertex. The first option is slower, but more careful. Another parameter enables adjusting the sorting of the neighbors. An extensive computational study revealed that tropical neighborhood search exploits so far unexplored optimization potential, in particular when the other local heuristics implemented in ConcurrentPESP are stuck in local optima. For five out of eight considered PESLib instances, tropical neighborhood search could improve the previously known primal bounds. This practical application of tropical methods to periodic timetabling has been awarded the ATMOS 2022 Best Paper Award [3].




Figure 12: Objective value evolution of modulo network simplex (black) vs. several configurations of tropical neighborhood search; lower is better.





PROGRAMMING ON THE DATA HIGHWAY

How SIMD Techniques Accelerate Non-Numerical Applications

Data parallelism, i.e., performing the same operations on a bunch of data in parallel, is a crucial factor to gain best performance from modern microprocessors. This parallelism can typically easily be found in numerical algorithms and applications, but is much harder to reveal in combinatorial codes that include many conditions and decisions along their execution path. Often, we need to adjust the basic algorithms carefully to reveal a good level of data parallelism.

At ZIB, we study and strive for leveraging modern microprocessors' data parallel capabilities in a variety of applications, aiming to use existing hardware more effectively, ultimately saving energy and accelerating algorithms.

SIMD and Implied Challenges for Combinatorial Algorithms

Modern high-performance processors draw most of their computational power from the presence of SIMD (single instruction multiple data) instruction set extensions. In contrast to ordinary (scalar) instructions, operating on individual numbers, SIMD instructions operate on **vectors** holding

4–16 numbers depending on data type and SIMD extension. Thus, each SIMD instruction performs up to 16 times as much work as a similar conventional (scalar) instruction, significantly accelerating the speed of data processing (see Fig. 1 for an illustration).




Figure 1: An array $A[]$ is added to an array $B[]$. Scalar execution (top) versus SIMD execution (bottom). SIMD permits us to turn four individual additions into one SIMD addition, saving time and energy.

Although required to access a large part of a CPU's computational power, use of SIMD instructions is largely restricted to specialized applications such as numerical computations, cryptography, and multimedia codecs. This is largely a consequence of three factors:

- a) Use of SIMD instructions requires the presence of multiple data items that can be processed independently at the same time. Many applications do not have a structure where multiple data items are processed at the same time, making it hard to use SIMD approaches without significant redesign of the algorithms involved.
- b) Most SIMD instructions are **vertical**, processing each vector element independently and in the same way. Operations that process the elements of a vector with one another or that change their order are tricky to use and often inefficient. On a high level, this also means that, ideally, all data items are processed the same way with as few special cases as possible. This makes it hard to use SIMD instructions for problems where data items are processed in dependence to one another or where individual items need special processing or cause a diversion in the course of the algorithm.

- c) Support for the use of SIMD instructions by high-level language is poor to nonexistent. And although a growing field of research (auto vectorization), compilers are largely incapable of leveraging SIMD instructions for conventionally written code automatically. As per the current state of the art, SIMD-accelerated algorithms must either be written in assembly or with vendor-supplied language extensions to C, C++, and Fortran. In addition to being clunky, this also means that SIMD code is largely unportable and must be written anew for each architecture to be supported. This situation increases the entry barrier to SIMD to the point where it is simply not on the radar for most programmers.

Despite these difficulties, we believe in the tremendous potential of SIMD techniques for general-purpose computing, especially for applications that are combinatorial (i.e., based on decision-making)

instead of numerical in nature. Where possible, use of SIMD techniques often leads to a tremendous increase in performance, often by a factor of two or more.

In our research, we hope to develop and contribute a catalogue of techniques for the application of SIMD instructions to combinatorial algorithms, defining a new field of "Combinatorial SIMD Programming." We hope that through a systematic approach to SIMD techniques, these will grow more accessible to programmers and that through widespread adoption of SIMD techniques in general applications, existing hardware can be used more effectively, ultimately saving energy and accelerating algorithms.

As part of this research, we have taken a number of sample problems from various areas of computer science and tackled them with SIMD techniques. In the following sections, we present some of these problems and how SIMD techniques effected better solutions.



Florian Schintke
Phone +49 30 841 85-306
schintke@zib.de



Robert Clausecker
Phone +49 30 841 85-263
clausecker@zib.de

Text Processing:

Fast Unicode Transcoding

Unicode is the international standard for representing text on a computer. It is widely used across diverse domains such as scientific data sets, databases, and Web applications. Depending on the requirements of the application, Unicode text is represented in a variety of **Unicode Transformation Formats** (UTF), the most popular being UTF-8 and UTF-16.

In cooperation with the workgroup of Prof. Daniel Lemire of TÉLUQ University of Québec, Canada, we have investigated the application of combinatorial SIMD techniques to the translation between text encodings, specifically UTF-8 and UTF-16.

While this task appears mundane on first sight, it touches on many challenging aspects of combinatorial SIMD programming and has broad application in industry, such as in data transfer from/to database systems and in Javascript interpreters, both important cornerstones to the World Wide Web.

case	UTF-16	UTF-8
ASCII	0000 0000 0 <u>G</u> FE DCBA	----- 0 <u>G</u> FE DCBA
2-byte	0000 0 <u>L</u> KJ HGFE DCBA	----- <u>110</u> L KJHG <u>10</u> FE DCBA
3-byte	<u>R</u> QPN <u>M</u> LKJ HGFE DCBA	----- <u>1110</u> <u>R</u> QPN <u>10</u> M <u>L</u> KJHG <u>10</u> FE DCBA
4-byte	<u>1101</u> <u>10</u> <u>v</u> <u>u</u> <u>t</u> sRQ PNML <u>1101</u> <u>11</u> KJ HGFE DCBA	<u>1111</u> <u>0</u> <u>W</u> <u>V</u> <u>U</u> <u>10</u> <u>T</u> S <u>R</u> QPN <u>10</u> M <u>L</u> KJHG <u>10</u> FE DCBA

(a) Bit-by-bit correspondence between UTF-16 and UTF-8 encodings in the four possible cases. The bits are named A to W starting at the least significant bits with 0vuts = WVUTS – 1.

codepoint	UTF-16	UTF-8
U+0	0000 0000 0 <u>0</u> <u>0</u> <u>0</u> 0000	----- 0 <u>0</u> <u>0</u> <u>0</u> 0000
U+7F	0000 0000 0 <u>1</u> <u>1</u> <u>1</u> 1111	----- 0 <u>1</u> <u>1</u> <u>1</u> 1111
U+80	0000 0 <u>0</u> <u>0</u> <u>0</u> 1 <u>0</u> <u>0</u> <u>0</u> 0000	----- <u>110</u> 0 <u>00</u> <u>10</u> <u>10</u> <u>00</u> 0000
U+7FF	0000 0 <u>1</u> <u>1</u> <u>1</u> 1 <u>1</u> <u>1</u> <u>1</u> 1111	----- <u>110</u> <u>1</u> <u>1111</u> <u>10</u> <u>11</u> 1111
U+800	<u>0</u> <u>0</u> <u>0</u> <u>0</u> 1 <u>0</u> <u>0</u> <u>0</u> 0 <u>0</u> <u>0</u> <u>0</u> 0000	----- <u>1110</u> <u>0000</u> <u>10</u> <u>10</u> <u>0000</u> <u>10</u> <u>00</u> 0000
U+FFFF	<u>1</u> <u>1</u> <u>1</u> <u>1</u> <u>1</u> <u>1</u> <u>1</u> <u>1</u> 1 <u>1</u> <u>1</u> <u>1</u> 1111	----- <u>1110</u> <u>1111</u> <u>10</u> <u>11</u> <u>1111</u> <u>10</u> <u>11</u> 1111
U+10000	<u>1101</u> <u>10</u> <u>0</u> <u>0</u> <u>0</u> 0 <u>0</u> <u>0</u> <u>0</u> 0000 <u>1101</u> <u>11</u> <u>0</u> <u>0</u> 0 <u>0</u> <u>0</u> <u>0</u> 0000	<u>1111</u> <u>0</u> <u>0</u> <u>0</u> <u>0</u> <u>10</u> <u>0</u> <u>1</u> 0000 <u>10</u> <u>00</u> 0000 <u>10</u> <u>00</u> 0000
U+10FFFF	<u>1101</u> <u>10</u> <u>1</u> <u>1</u> 1 <u>1</u> <u>1</u> <u>1</u> 1 <u>1</u> <u>1</u> <u>1</u> <u>1101</u> <u>11</u> <u>1</u> <u>1</u> 1 <u>1</u> <u>1</u> <u>1</u> 1111	<u>1111</u> <u>0</u> <u>1</u> <u>00</u> <u>10</u> <u>00</u> <u>1111</u> <u>10</u> <u>11</u> <u>1111</u> <u>10</u> <u>11</u> 1111

(b) Examples of matched code-point values in UTF-32, UTF-16LE, and UTF-8. For U+10000 and U+10FFFF, UTF-16 requires a surrogate pair.

Figure 2: Correspondence between UTF-16 and UTF-8. Format-specific prescribed bits (tag bits) are underlined.

Untangling Characters with New Tricks

The main difficulty comes from both encodings being **variable length encodings**. A universal character is represented as a sequence of 1–4 UTF-8 bytes or 2–4 UTF-16 bytes (see Fig. 2). This makes it hard to process multiple characters in parallel: It is not a priori known at which byte boundaries UTF-8/UTF-16 sequences begin or end, and each universal character may result in a variable number of output bytes. Additionally, the transcoding steps differ depending on how long the encodings involved are.

It is thus not obvious how a string of universal characters can be transcoded using SIMD techniques, which demand that the characters be processed independently and in the same way. Previous approaches such as [5] and [7] simplified the problem by not considering characters that encode into four UTF-16 bytes (these are exactly the 4-byte UTF-8 sequences), thus reducing UTF-16 to a much simpler encoding where each character is encoded in 2 bytes. Nevertheless, getting from 1, 2, or 3 UTF-8 bytes to 2 UTF-16 bytes is still challenging.

Our approach improves on these results by finding an elegant way to treat 4-byte UTF-8 sequences as an overlapping pair comprising a 2-byte sequence and a 3-byte sequence, thus reducing the cumbersome 4-byte UTF-16 sequences back to the 2-byte base case. We have also made use of many new instructions of the AVX-512 instruction set extensions, eliminating some major roadblocks in unraveling UTF-8 sequences into individual universal characters. Lastly, through a novel validation algorithm, the correctness of the input can now be checked at little to no performance cost throughout the procedure.

Thanks to these and other ideas and techniques, we were able to improve transcoding speeds by a factor of up to ten compared to a conventional implementation, and by up to three compared to the previous best SIMD-accelerated implementation [7], transcoding, e.g., Latin text from UTF-8 to UTF-16 at 23 GB/s, Arabic text at 7.8 GB/s, and Chinese text at 5.8 GB/s (see also Fig. 3).

A paper on this research project is currently under review at Wiley's "Software: Practice and Experience" and is expected to be published in Summer 2023.




Figure 3: Transcoding speeds in gigabytes of input data per second for various test files. We compare against the industry standard ICU library and the previous AVX2 kernels from [7].

Bioinformatics: Positional Population Counts for Genome Analysis

The problem of **taxonomic profiling** asks to find from a set of RNA or DNA samples the set of possible species these samples may have come from. This is commonly implemented using techniques from string searching like k-mer matching.

The idea of k-mer matching is to compute all k-mers (length k substrings) occurring in the sample. This set of k-mers is often much smaller than the sample string itself, and identifies it like a fingerprint. Matching these k-mer fingerprints against a database of known genomes, we can quickly find gene sequences with high overlaps in their k-mers, identifying potential matches even in the presence of slight transcribing errors.




Figure 4: The DNA sequence ACTAGACTA and its 3-mer set {ACT, AGA, CTA, GAC, TAG}.

In Shen's taxonomic profiler KMCP [8], this algorithm idea is implemented by computing for each k-mer a bit mask, indicating at bit position i if this k-mer occurs in sample string i . To check which samples might correspond to a given DNA/RNA sequence, we gather the bit masks for all k-mers that occur in that sequence and count for each i how often the bit at position i is set. If this count is close to the number of incidence words in the set, we have a possible match and can investigate further.

Positional Population Counts: Histograms on Bit Sets

This step is effectively computed using a **positional population count** operation: of an array of bit masks, a histogram is created, holding for each bit position how often this bit is set. With each bit mask representing the set of samples that contain a given k-mer, we obtain for an array of such bit masks formed from a candidate DNA/RNA sequence how well each of the samples match it.

Unfortunately, there is no direct or obvious way to compute this operation. Although modern processors generally have instructions to compute ordinary population counts (i.e., summing up the bits of a bit mask), this operation is not really useful for this purpose. In response to a public request by the work group authoring the KMCP taxonomic profiling program [8] to optimize the positional population count step of this search algorithm, we have investigated ways to compute this operation quickly.

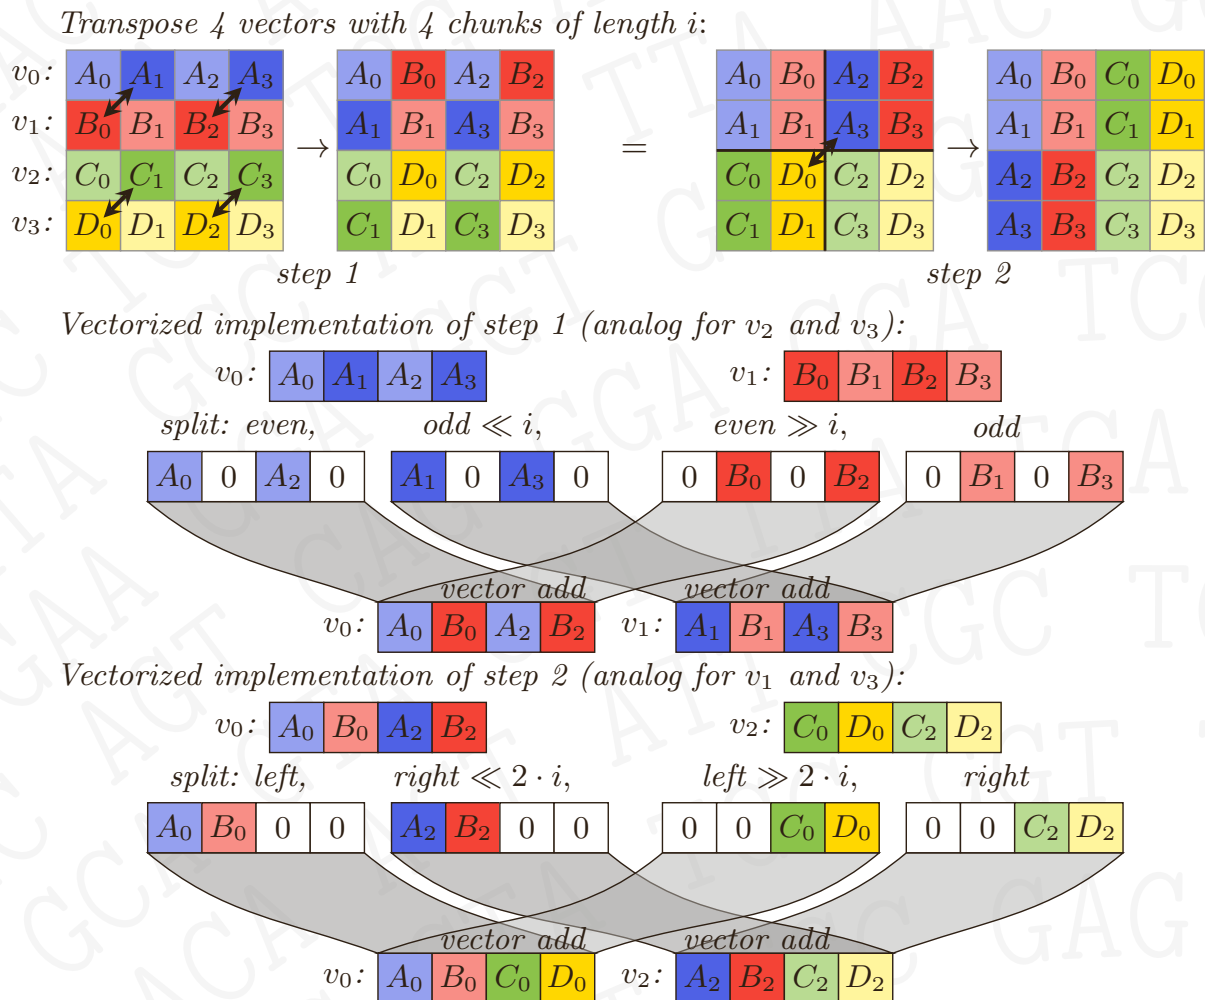


Figure 5: Vectorized matrix transposition.

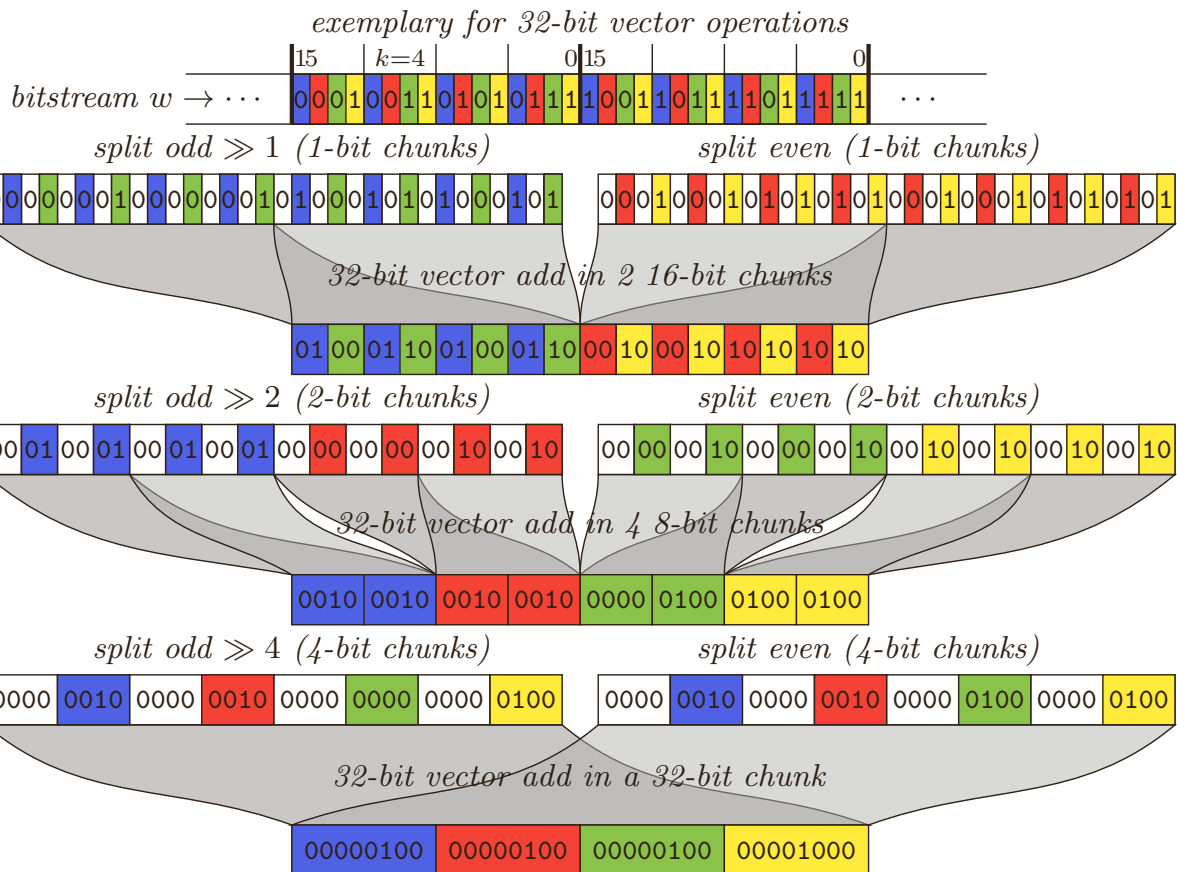


Figure 6: Vectorized positional population count on a bitstream for four counters.

In this research, we have discovered that the positional population count could effectively be thought of as being a transposition operation (see Fig. 5), followed by a series of ordinary population counts. Taking inspiration from the common high-speed recursive transposition algorithm, we have integrated into this transposition the classic bit-parallel population count algorithm [9], directly yielding a vector of population counts (see Fig. 6).

This novel transposition idea was then combined with the established carry-save adder algorithm skeleton from Klarqvist et al.'s approach [6], reducing the number of times even this fast kernel has to be executed to a minimum.

Special care was taken to optimize the algorithm to have low start-up and finalization times, leading to processing speeds in the order of multiple GB/s even for arrays as small as 1 kB. For larger arrays, speeds of up to 97 GB/s are reached, about twice as fast as the previous best algorithm.

Heuristic Search: Bit-Parallel Indexing for Faster PDB Lookups

As part of our award-winning research into the quality of heuristics for the IDA* algorithm [3], we implemented our own solver for the 24 puzzle, a standard test bed problem in this field. This solver makes use of powerful zero-aware pattern database (ZPDB) heuristics represented as tables of up to 20 GB in size each.

The heart of the solver is the **index function**, which decomposes a puzzle instance into a set of patterns (see Fig. 7), each of which is looked up in the corresponding pattern database. About half of the solver's runtime is spent here.

It is this index function that we optimized using combinatorial SIMD techniques, leading to a 30% performance improvement over the previous scalar code.

The main use of SIMD techniques lies in computing for each pattern where the tiles corresponding to the pattern are. Reducing the puzzle to just these tiles, we can quickly compute where in the ZPDB the pattern is found.

Lessons Learned and Future Directions

SIMD techniques are a powerful tool to make the most of the abilities of modern high-performance processors. Throughout our research we discovered and applied many novel tricks to make use of SIMD instruction set extensions even for problems that don't seem to lend themselves to these methods. This allowed us to achieve previously unimaginable performance gains for practically relevant algorithms.

Key tools in this process are accurate data about the performance of SIMD instruction set extensions, such as those published by Fog [4] and Abel [1], and accurate microarchitectural simulations built from these, such as uiCA [2]. Only with such accurate information is it possible to understand performance bottlenecks in SIMD code and to fix such shortcomings.

In future work, we plan to attack exciting and important algorithms such as SAT solvers and suffix array construction algorithms. We believe that even though these algorithms are highly sequential with little to no obvious potential for SIMD techniques, they may nevertheless be applied with smart ideas and clever programming.




Figure 7: A 24 puzzle instance and its partitioning into the {1, 2, 5, 6, 7, 10, 11}, {3, 4, 8, 9, 13, 14, 19, 24}, and {15, 16, 17, 18, 20, 21, 22, 23} patterns.

ADDRESSING ENERGY CHALLENGES IN HPC DATA CENTERS





Benchmarking energy savings on the NHR@ZIB supercomputer system "Lise"

The rising costs of energy consumption have become a severe issue in HPC. To identify potential areas in the NHR@ZIB supercomputer system for energy-saving measures, a representative mix of benchmark applications has been executed at reduced CPU speeds. We observe that individual code characteristics determine whether these measures reduce energy consumption.

Power Requirements in HPC

During the past decades, high-performance computing (HPC) research has mainly focused on computational performance. Since the mid-1990s, HPC has accumulated many new features, including the introduction of grid computing, multicore processors, and the development of the GPUs, which leads us to the gradual ending of Moore's law as high-performance computers become more powerful. However, with their increased performance, energy consumption has also increased; for scale, based on current trends a 1 exaflop machine's energy consumption can easily reach 20 MW within the next few years. Clearly, such developments are unsustainable in both an environmental and financial point of view. While the issue of energy consumption has been largely neglected in the past, as the HPC community reaches from petascale to exascale, power efficiency has already become the main concern for the future. Recent global developments and governmental pressures have especially highlighted the need for the design of energy efficient HPC.

The energy consumed by a supercomputer can be measured at vastly differing accuracies and granularities. In many cases, to determine a machine's efficiency it might be sufficient to measure the overall energy used by the whole system, e.g., at the main power distribution of the compute center; however, depending on the purpose, a more detailed measurement might be required. In increasing granularity, depending on the equipment used, measurements may be taken on the level of full compute racks, power distribution units within the racks, node chassis, or single nodes and even up to component level such as processors, memory modules, and accelerators. While measurements on the component level might require access to programs on the node or a vendor-specific system, most of the other devices can be read over the network via specific interfaces. In the present study we use power consumption values available for individual nodes of the NHR@ZIB supercomputer system "Lise" (shown in Fig. 1).



Team from left to right: Thomas Steinke, Anita Ragyanszki, Tobias Watermann, Ronaldo Rodrigues Pelá, Lewin Stein, Matthias Läuter, other team members not pictured Steffen Christgau and Christian Tuma.



Figure 1: NHR@ZIB supercomputer system “Lise” [27].
(Author: ITMZ, Universität Rostock)

Over the course of its lifetime, the energy costs for operating an HPC system reach or even exceed the initial costs of its hardware investment. Also, non-computational parts, such as cooling components, still consume a significant percentage of the overall energy usage. Besides its financial burden to institutions worldwide, HPC’s high energy consumption is becoming a major contributor to their carbon footprint. Environmental sustainability plays an important role in society, businesses, and scientific research; therefore, it is not surprising that the operational cost of cluster computing is one of the major concerns of data centers around the world. Institutions are actively searching for ways to reduce the operational cost and provide Green HPC resources while still expanding computational services. In 2008, the Green500 list [1] was introduced as a way to track the current most energy-efficient supercomputers (see Fig. 2). This initiative was successful in raising awareness of HPC’s high power consumption, and in promoting the development of performance metrics aimed at lowering the total cost of HPC.

The computationally intensive nature of HPC gives rise to exceptional challenges in curbing energy consumption. Fortunately, newly developing processor designs offer a promising start to improving energy efficiency, while global initiatives are making considerable progress toward a greener field in high-performance computing.

Trends in Energy-Efficient Computing

Significant costs for investment, energy, and maintenance lead to a tradeoff between time-to-solution, energy-to-solution, and total cost of ownership for HPC systems. Unified by the goal of advancing performance improvements, funding agencies have launched research initiatives to join low energy consumption and high-performance computing [2]. Since June 2022, the best computer on the TOP500 list of HPC has been an exascale supercomputer. Meanwhile, the same computer holds only the 6th place on the Green500 list. To obtain these improvements, energy-efficient computing is a combination of efforts with respect to the architecture on the one hand and to the software on the other.

Today, HPC systems consist of compute nodes with homogeneous or heterogeneous architectures, connected by a fast communication network, and attached to a fast parallel file system. Homogeneous architectures rely on one or more central processing units (CPU) in a single node while heterogeneous architectures enhance the node configuration with additional accelerators. On each of these microchips, power consumption is controlled by dynamic voltage and frequency scaling (DVFS) down to the limit of near-threshold voltage [3]. Dynamic power management acts as an interface to control energy usage and offers suggestions to the users to switch between power save




Figure 2: Historic trend of the energy efficiency for rank 1 of the Green500 list [1].

mode and performance mode to maintain energy efficiency. The power limit for a complete chip (thermal design power) allows for the trade between energy savings in one section for performance increases to another. Energy efficient solutions are based on many-core accelerators, like graphics processing units (GPU) [4], or specially designed CPUs, like the ARM A64FX of the Fugaku supercomputer [5]. Green500 energy efficiency started in the order of single-digit GFlops/W in 2013 and reached 21 GFlops/W in 2020. Today's highest performing computers per energy is supported by GPUs and reserve the top positions in the Green500 list (65 GFlops/W in the first position) for systems based on accelerators like NVIDIA H100 and AMD MI250X. Energy saving for Fugaku has placed the supercomputer in the 43rd place on the Green500 but 2nd on the TOP500 list. However, energy efficiency comes at a price. The top position in the Green500 is 405th on the TOP500 list, as not all application codes fully utilize the performance of accelerators, and the effort for code optimization is increased. Apart from that, most of the systems heading the TOP500 list are based on accelerators which make heterogeneous architectures the most promising design to establish high-performance computing systems in the near future [6].

Heterogeneous architectures can be accessed by functions in vendor-specific libraries [7]. Extended flexibility is available by unified programming models which consist of extensions for the predominating HPC languages, C/C++ and Fortran, such as SYCL and OpenMP. Monitoring and profiling do not focus on runtime and floating-point statistics only but also on energy consumption. For that one can rely on the Linux interface to the hardware's power controller, like Intel's Running Average Power Limit (RAPL), the Performance API (PAPI) library, and further software projects going back to these interfaces. Network access to these interfaces is available, for example by SNMP and IPMI interfaces. Dynamic tuning tools reflect runtime analysis of the code and improve parameter setups [8]. Thus the user is able to analyze the code with respect to execution time and energy consumption. Because accelerators and some applications, such as in the field of machine learning, are efficient especially for lower precision, the operation count is considered for double, single, and half precision floating point types. User codes can be rated with respect to these different metrics which gives indications for the properties of future system procurements.

Energy-Efficient Operation of ZIB's Supercomputer "Lise" in Practice

The previous section introduced the principal techniques to save energy during computation. Mostly, these techniques require a time-consuming (re)implementation by the software developers. In contrast, this section proposes hardware settings an HPC system operator can implement to run energy efficiently – without the need to reimplement individual software.

Energy efficiency depends both on the hardware (and its settings) and on the software. Hence, any increase in energy efficiency due to new hardware settings, as presented below, must be interpreted carefully. Results are valid for a given benchmark only, and a transfer of expectations to other applications requires at least similar code characteristics. To capture the impact of new hardware settings more generally we selected a number of benchmarks representing the application mix of the user community at NHR@ZIB. These are GROMACS [9] and Siesta [10] (chemistry), PALM (geoscience), OpenFOAM (engineering), and the representative basic linear algebra benchmark

DGEMM. For further details about the specific program setups please refer to Table 1 and to [11].

Powering off parts of the system is not discussed here as this will only postpone computations and the associated energy consumption. ZIB's supercomputing group evaluated two different power-saving approaches on "Lise" – introduction of a reduced limit for a compute node's power consumption (power capping), and enforcement of reduced CPU clock frequencies.

Power capping has the advantage that the compute node's management system can automatically and dynamically decide where to save energy. For example, core and memory frequencies can be controlled individually [12] according to the system's needs and configured limits. However, this setting caused most of our benchmarks to run much slower than expected, sometimes even consuming more energy than before. The reason was a desynchronization of the core frequencies and, hence, an unequal load distribution. Here we observed unpredictable and unreproducible

Benchmark	Version	Characteristic	Nodes x tasks per node	Benchmark description	Benchmark reference
DGEMM	N/A	compute bound	1 x 96 (serial)	matrix-matrix multiplication with 10,000 double-precision floating point numbers per row and column	N/A
GROMACS	2021.2	compute bound (memory bound)	1 x 96	benchRIB (ribosome in water)	[26]
OpenFOAM	2012	communication bound	8 x 96	SimpleFoam / HPC motorbike	[11]
PALM	21.10	communication bound	8 x 96	Domain size: 4,096 x 2,048 x 1,024, FFTW-based Poisson solver	N/A
Siesta	4.1.5	memory bound	1 x 96	512 molecules of H ₂ O in a box	[10]

Table 1: Details of the selected benchmark set.

differences between the clock frequencies assigned to the CPU cores of the first CPU socket and those assigned to the CPU cores of the second CPU socket of a given compute node, and also between compute nodes.

This experience led us to the second approach, the explicit specification of a reduced clock frequency valid for all CPU cores. Since the power consumption of a processor scales significantly with its frequency [13], clocking down is a promising way to save energy. We indeed observed energy savings, but to non-uniform extents depending on the characteristics of a given benchmark. Our results are shown in Table 2 and Figs. 3a,b.

CPU clock frequency [GHz]	Energy-saving (run-time penalty)				
	DGEMM	GROMACS	OpenFOAM	PALM	Siesta
1.5					4.6% (1.16x)
1.6	0.7% (1.36x)	11.0% (1.46x)	25.1% (1.11x)	22% (1.11x)	4.2% (1.13x)
1.7	0.6% (1.29x)	7.3% (1.28x)			3.7% (1.10x)
1.8	1.8% (1.22x)	5.5% (1.23x)	22.5% (1.08x)	15% (1.07x)	
1.9	0.5% (1.16x)	3.9% (1.17x)			2.5% (1.04x)
2.0	1.0% (1.11x)	3.3% (1.13x)	20.4% (1.09x)	10% (1.03x)	
2.1	0.7% (1.06x)	3.2% (1.10x)			0.0% (1.00x)
2.2	0.9% (1.02x)	3.0% (1.07x)	19.1% (1.04x)		
2.3	0.6% (1.01x)	1.9% (1.05x)	17.7% (1.02x)	5% (1.00x)	-3.9% (0.97x)

Table 2: Energy savings obtained for each of the benchmark runs carried out at various reduced CPU clock frequencies. Corresponding benchmark run-time penalties are shown in parentheses.




Figure 3a: Energy savings measured for running the benchmarks at reduced CPU clock frequencies.




Figure 3b: Run-time penalties seen for the benchmarks executed at reduced CPU clock frequencies.

Compared to the reference case where a benchmark is executed without limiting CPU clock frequencies (turbo mode activated), between 22% and 25% of energy was saved in our OpenFOAM and PALM benchmark runs at a CPU clock frequency of 1.6 GHz. With GROMACS approximately 11% of energy was saved at 1.6 GHz. For Siesta and DGEMM it was only 4% and 0.7%, respectively. One explanation why OpenFOAM and PALM consume less energy at reduced CPU clock frequency compared to, for example, DGEMM is as follows. The investigated cases of OpenFOAM and PALM are mainly bound by communication between processes, while for the execution of DGEMM the bottleneck is defined only by the computational power of the CPU. This is in line with the general rule of thumb that communication or memory bound programs are less affected by clocking down than compute bound programs [14,15]. Running at lower CPU clock frequency reduces power consumption and also affects benchmark run times, but the latter does not have a too strong impact in the case of a communication bound code which results in energy savings for applications of that particular characteristic.

Our main findings can be summarized as follows: On the one hand, limiting the total power consumption of a compute node turned out to trigger unpredictable discrepancies between clock frequency settings which can result in even higher energy consumptions due to load imbalances and, therefore, extended run times. On the other hand, ex-

plicit specification of reduced CPU clock frequency values valid for all CPU cores can lead to substantial energy savings. This, in particular, is true for applications characterized by communication type bottlenecks; however compute bound applications do not virtually yield to energy savings with reduced CPU clock frequencies.

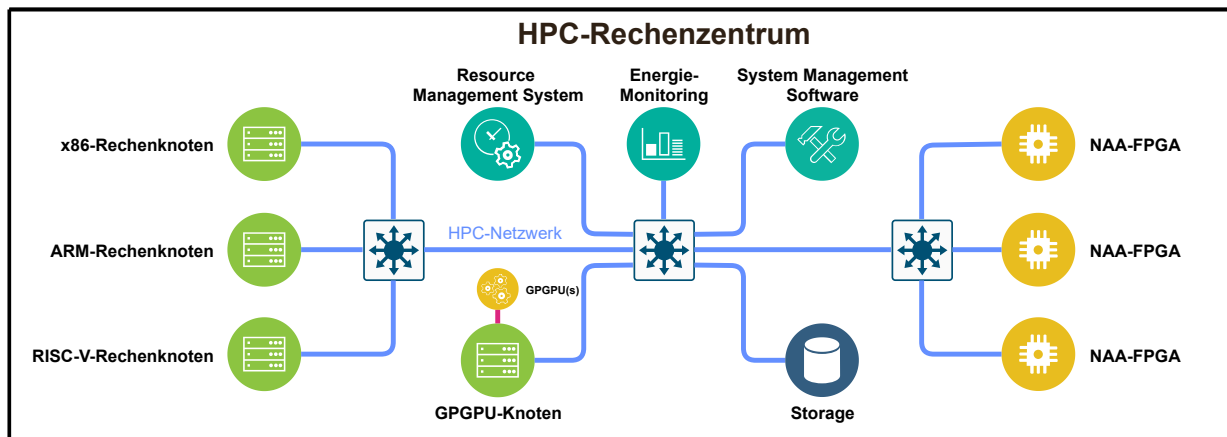


Figure 4: Conceptual overview of an HPC data center that employs FPGAs as network-attached accelerators (NAA).

Outlook

While the previously discussed actions address energy savings for ZIB's current HPC system, research has been focusing on the energy-efficient design of new systems [16,17,18,19]. One key aspect is the increasing usage of heterogeneous systems as discussed in the previous annual report. ZIB has followed this trend by using energy-efficient GPUs ([4], see above) as compute accelerators. Forty-two compute nodes, each with four such devices, were installed as an extension to "Lise" as of the end of 2022 and will enable key application codes to consume less energy, starting from early 2023.

Besides these already common GPUs, even more specialized accelerators, like Field Programmable Gate Arrays (FPGA), will be used in future systems. The Supercomputing research department at ZIB focuses on this particular technology due to its significantly larger energy efficiency and algorithmical flexibility [20,21,22] compared to GPUs. Although they are architecturally very different from CPUs and GPUs, FPGAs can be used with established programming models as was demonstrated by research activities of the HPCLab of the MODAL

research campus and within the ORKA-HPC project [23,24].

By using the internal expertise as well as by emphasizing existing and establishing new collaborations with recognized partners, the Algorithms for Innovative Architectures research group was able to acquire two projects funded by the BMBF within the "GreenHPC" program (see info box). Along with only seven other funded projects, ZIB and its partners (such as FU Berlin, Fraunhofer Heinrich Hertz Institute, and Intel) will contribute a significant share to the program results with a focus on the FPGA technology. In addition, ZIB became an "Intel oneAPI Center of Excellence" and will continue its research on the usage of new GPU programming models.

These combined efforts of system operations, understanding of application behavior, and research on new technologies emphasize the wide range of current and future activities of the Supercomputing Department to both reduce energy consumption and improve the energy efficiency for HPC in Berlin.

BMBF-funded GreenHPC projects

SeqAn@FPGA aims to demonstrate the usability of FPGAs for energy-efficient genome sequencing in collaboration with FU Berlin and Intel. This includes the development and evaluation of improved data structures, like distributed and interleaved Bloom filters, and new algorithms that can benefit from custom data types or pipelined execution.

NAAICE focuses on extending the usage of energy-efficient FPGAs in HPC data centers. The project builds around the idea to disaggregate the FPGAs from energy-wasting host systems and to attach them directly to an HPC interconnect (see Fig. 4), making them network-attached accelerators (NAAs) [25]. The project covers the design of efficient middleware to communicate with such NAAs, the offloading of performance-critical code paths of a relevant application, as well as the integration of NAAs into HPC workflows. The latter addresses challenges like remote configuration as well as the integration in resource management and monitoring infrastructure of the HPC data centers.

REFERENCES

Taking Math to the Heart

- [1] W. Abplanalp, N. Tucker, and S. Dimmeler. Single-cell technologies to decipher cardiovascular diseases. *European Heart Journal*, 43 (43):4536–4547, 2022.
- [2] R. Byrd, J. Nocedal, and R. Schnabel. Representations of quasi-Newton matrices and their use in limited memory methods. *Mathematical Programming* 63 (1–3):129–156, 1994.
- [3] F. Chegini. Multilevel Optimization Algorithm for Inverse Problem in Electrocardiography. *PhD thesis*, Università della Svizzera italiana, 2022.
- [4] F. Chegini, A. Kopanicakova, R. Krause, and M. Weiser. Efficient Identification of Scars using Heterogeneous Model Hierarchies. *EP Europace* 23:i113–i122, 2021.
- [5] F. Chegini, A. Kopanicakova, M. Weiser, and R. Krause. Quantitative Analysis of Nonlinear Multifidelity Optimization for Inverse Electrophysiology. In S.C. Brenner, E.T. Chung, A. Klawonn, F. Kwok, J. Xu, and J. Zou. Domain Decomposition Methods in Science and Engineering XXVI, pp. 65–76, *Springer*, 2022.
- [6] F. Chegini, T. Steinke, and M. Weiser. Higher-order time integration with algebraic adaptivity in discretization of cardiac models. *Submitted 2022*.
- [7] S. Gratton, A. Sartenaer, and Ph. Toint. Recursive trust-region methods for multiscale nonlinear optimization. *SIAM J. Optimization* 19(1):414–444, 2008.
- [8] C. Groß and R. Krause. On the Convergence of Recursive Trust-Region Methods for Multiscale Nonlinear Optimization and Applications to Nonlinear Mechanics. *SIAM J. Numerical Analysis* 47(4):3044–3069, 2009.
- [9] N. M. M. Huynh, F. Chegini, L. F. Pavarino, M. Weiser, and S. Scacchi. Convergence analysis of BDDC preconditioners for hybrid DG discretizations of the cardiac cell-by-cell model. *Submitted 2022*.
- [10] M. Rams and T. O. F. Conrad. Sparse dictionary learning allows model-free pseudotime estimation of transcriptomics data. *BMC Genomics*, 23:56, 2022.
- [11] K. Weiman and T. O. F. Conrad. Transfer Learning for ECG Classification. *Scientific Reports*, 11:5251, 2021.
- [12] M. Weiser and F. Chegini. Adaptive multirate integration of cardiac electrophysiology with spectral deferred correction methods. In P. Nithiarasu and C. Vergara. *CMBE22 – 7th International Conference on Computational & Mathematical Biomedical Engineering*. pp. 528–531, 2022.
- [13] M. Weiser and S. Ghosh. Theoretically optimal inexact spectral deferred correction methods. *Comm. Appl. Math. Comp. Sci.* 13(1):53–86, 2018.
- [14] M. Weiser and S. Scacchi. Spectral Deferred Correction methods for adaptive electro-mechanical coupling in cardiac simulation. In G. Russo, V. Capasso, G. Nicosia, and V. Romano. *Progress in Industrial Mathematics at ECMI 2014*. pp. 321–328, Springer, 2017.

Numerics Meets Quantum Technology

- [1] T. Gao, L. Rickert, F. Urban, J. Große, N. Srocka, S. Rodt, A. Musial, K. Zolnatz, P. Mergo, K. Dybka, W. Urbanczyk, G. Sek, S. Burger, S. Reitzenstein, and T. Heindel. A quantum key distribution test bed using a plug & play telecom-wavelength single-photon source. *Appl. Phys. Rev.* 9, 011412, 2022.
- [2] P. Gelß, S. Klus, S. Knebel, Z. Shakibaie, and S. Pokutta. Low-rank tensor decompositions of quantum circuits, *arXiv*:2205.09882, 2022.
- [3] F. Binkowski, F. Betz, M. Hammerschmidt, P.-I. Schneider, L. Zschiedrich, and S. Burger. Computation of eigenfrequency sensitivities using Riesz projections for efficient optimization of nanophotonic resonators. *Commun. Phys.* 5, 202, 2022.
- [4] F. Betz, F. Binkowski, and S. Burger. RPEExpand: Software for Riesz projection expansion of resonance phenomena. *SoftwareX* 15, 100763, 2021.
- [5] X. Garcia Santiago, S. Burger, C. Rockstuhl, and P.-I. Schneider. Bayesian optimization with improved scalability and derivative information for efficient design of nanophotonic structures. *J. Light. Technol.* 39, 167, 2021.
- [6] M. Plock, K. Andrie, S. Burger, and P.-I. Schneider. Bayesian Target-Vector Optimization for Efficient Parameter Reconstruction. *Adv. Theory Simul.* 5, 2200112, 2022.
- [7] R. Colom, F. Binkowski, F. Betz, Y. Kivshar, and S. Burger. Enhanced Purcell factor for nano-antennas supporting interfering resonances. *Phys. Rev. Research* 4, 023189, 2022.
- [8] Y. Wang, L. Vannucci, S. Burger, and N. Gregersen. Near-unity efficiency in ridge waveguide-based, on-chip single-photon sources. *Mater. Quantum Technol.* 2, 045004, 2022.
- [9] L. Bremer, C. Jimenez, S. Thiele, K. Weber, T. Huber, S. Rodt, A. Herkommer, S. Burger, S. Höfling, H. Giessen, and S. Reitzenstein. Numerical optimization of single-mode fiber-coupled single-photon sources based on semiconductor quantum dots. *Opt. Express* 30, 15913, 2022.
- [10] K. Volkova, J. Heupel, S. Trofimov, F. Betz, R. Colom, R. W. MacQueen, S. Akhundzada, M. Reginka, A. Ehresmann, J. P. Reithmaier, S. Burger, C. Popov, and B. Naydenov. Optical and spin properties of NV center ensembles in diamond nano-pillars. *Nanomaterials* 12, 1516, 2022.
- [11] F. Binkowski, T. Wu, P. Lalanne, S. Burger, and A. O. Govorov. Hot Electron Generation through Near-Field Excitation of Plasmonic Nanoresonators. *ACS Photonics* 8, 1243, 2021.
- [12] O. Avalos-Ovando, E. Y. Santiago, A. Movsesyan, X.-T. Kong, P. Yu, L. V. Besteiro, L. Khosravi Khoshnashad, H. Okamoto, J. M. Slocik, M. A. Correa-Duarte, M. Comesaña-Hermo, T. Liedl, Z. Wang, G. Markovich, S. Burger, and A. O. Govorov. Chiral Bioinspired Plasmonics: A Paradigm Shift for Optical Activity and Photochemistry. *ACS Photon.* 9, 2219, 2022.
- [13] A. Movsesyan, E. Y. Santiago, S. Burger, M. A. Correa-Duarte, L. V. Besteiro, Z. Wang, and A. O. Govorov. Plasmonic Nanocrystals with Complex Shapes for Photocatalysis and Growth: Contrasting Anisotropic Hot-Electron Generation with the Photothermal Effect. *Adv. Optical Mater.* 10, 2102663, 2022.

[14] S. Rej, E. Y. Santiago, O. Baturina, Y. Zhang, S. Burger, S. Kment, A. O. Govorov, and A. Naldoni. Colloidal Titanium Nitride Nanobars for Broadband Inexpensive Plasmonics and Photochemistry from Visible to Mid-IR Wavelengths. *Nano Energy* 104, 107989, 2022.

[15] K. Martens, T. Funck, E. Y. Santiago, A. O. Govorov, S. Burger, and T. Liedl. Onset of Chirality in Plasmonic Meta-molecules and Dielectric Coupling. *ACS Nano* 16, 16143, 2022.

[16] M. Dass, L. Kuen, G. Posnjak, S. Burger, and T. Liedl. Visible wavelength spectral tuning of absorption and circular dichroism of DNA-assembled Au/Ag core-shell nanorod assemblies. *Mat. Adv.* 3, 3438, 2022.

[17] K. Martens, F. Binkowski, L. N. Nguyen, L. Hu, A. O. Govorov, S. Burger, and T. Liedl. Long- and Short-Ranged Chiral Interactions in DNA Assembled Plasmonic Chains. *Nat. Commun.* 12, 2025, 2021.

[18] Y. Negrin-Montecelo, A. Movsesyan, J. Gao, S. Burger, Z. M. Wang, S. Nlate, E. Pouget, R. Oda, M. Comesana-Hermo, A. O. Govorov, and M. A. Correa-Duarte. Chiral Generation of Hot Carriers for Polarization-Sensitive Plasmonic Photocatalysis. *J. Am. Chem. Soc.* 144, 1663, 2022.

Figure 1 designed by L. Kulot (F4STUDIO).

Hydrogen Blending

[1] Bundesministerium für Wirtschaft und Energie (BMWi) *Die Nationale Wasserstoffstrategie*. 2020 (<https://www.bmwk.de/Redaktion/DE/Publikationen/Energie/die-nationale-wasserstoff-strategie.html>).

[2] <https://ehb.eu/>.

[3] J. Pedersen, K. Spreckelsen, U. Gotzes, J. Zittel, and T. Koch. Beimischung von Wasserstoff zum Erdgas: Eine Kapazitätsstudie des deutschen Gasnetzes, *gwf Gas + Energie*, 2023 (accepted for publication)

[4] FNB Gas, Netzenwicklungsplan Gas 2020–2030 (https://fnb-gas.de/wp-content/uploads/2021/09/fnb_gas_nep_gas_2020_de-1.pdf).

[5] J. Pedersen, K. Hoppmann-Baum, J. Zittel, and T. Koch. Blending hydrogen into natural gas: An assessment of the capacity of the German gas grid. *Operations Research Proceedings 2021*. OR 2021, 2022.

[6] K. Hoppmann-Baum. Mathematical programming for stable control and safe operation of gas transport networks. *Doctoral thesis, Technische Universität Berlin*, 2022.

[7] T. Koch, B. Hiller, M. E. Pfetsch, and L. Schewe, Eds. Evaluating Gas Network Capacities. *SI-AM-MOS Series on Optimization*, 2015.

[8] K. Hoppmann-Baum, F. Hennings, R. Lenz, U. Gotzes, N. Heinecke, K. Spreckelsen, and T. Koch. Optimal Operation of Transient Gas Transport Networks. *Optimization and Engineering*, Vol.22, pp. 735–781, 2021.

[9] F. Hennings, L. Anderson, K. Hoppmann-Baum, M. Turner, and T. Koch. Controlling transient gas flow in real-world pipeline intersection areas. *Optimization and Engineering*, Vol.22, pp. 687–734, 2021.

[10] M. Petkovic, Y. Chen, I. Gamrath, U. Gotzes, N. S. Hadjidimitrou, J. Zittel, X. Xu, and T. Koch. A hybrid approach for high precision prediction of gas flows. *Energy Systems*, Vol. 13, pp. 383–408, 2022.

Computer-Assisted Proofs in Extremal Combinatorics

[1] A. Razborov. Flag algebras. *The Journal of Symbolic Logic* 72.4:1239–1282, 2007.

[2] R. Graham, B. Rothschild, and J. Spencer. *Ramsey theory*. Vol. 20. John Wiley & Sons, 1991.

[3] A. Goodman. On sets of acquaintances and strangers at any party. *The American Mathematical Monthly* 66.9:778–783, 1959.

[4] P. Erdős. On the number of complete subgraphs contained in certain graphs. *Magyar Tud. Akad. Mat. Kutató Int. Közl.* 7(3):459–464, 1962.

[5] A. Thomason. A disproof of a conjecture of Erdős in Ramsey theory. *Journal of the London Mathematical Society* 2.2:246–255, 1989.

[6] J. Cummings et al. Monochromatic triangles in three-colored graphs. *Journal of Combinatorial Theory, Series B* 103.4:489–503, 2013.

[7] S. Kirkpatrick, D. Gelatt, and M. Vecchi. Optimization by Simulated Annealing. *Science*, 220(4598):671–680, 1983.

[8] F. Glover. Tabu search – part I. *ORSA Journal on computing*, 1(3):190–206, 1989.

[9] O. Parczyk, S. Pokutta, C. Spiegel, and T. Szabó. Fully Computer-Assisted Proofs in Extremal Combinatorics. *Proceedings of the 37th AAAI Conference*, 2023.

[10] L. Lovász and B. Szegedy. Limits of dense graph sequences. *Journal of Combinatorial Theory, Series B*, 2006.

[11] A. Kiem, S. Pokutta, and C. Spiegel. The 4-color Ramsey Multiplicity of Triangles. In *preparation*, 2023.

[12] K. Gatermann and P. Parrilo. Symmetry groups, semidefinite programs, and sums of squares. *Journal of Pure and Applied Algebra* 192.1–3: 95–128, 2004.

The Tropical Geometry of Periodic Timetabling

- [1] R. Borndörfer, H. Hoppmann, M. Karbstein, and N. Lindner. Separation of cycle inequalities in periodic timetabling. *Discrete Optimization*, vol. 35, p. 100552, Feb. 2020, doi: 10.1016/j.disopt.2019.100552.
- [2] R. Borndörfer, N. Lindner, and S. Roth. A concurrent approach to the Periodic Event Scheduling problem. *Journal of Rail Transport Planning & Management*, vol. 15, p. 100175, Sep. 2020, doi: 10.1016/j.jrtpm.2019.100175.
- [3] E. Bortoletto, N. Lindner, and B. Masing. Tropical Neighbourhood Search: A New Heuristic for Periodic Timetabling. In *22nd Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2022)*, Dagstuhl, Germany, 2022. doi:10.4230/OASlcs.ATMOS.2022.3.
- [4] E. Bortoletto, N. Lindner, and B. Masing. The tropical and zonotopal geometry of periodic timetables. *arXiv:2204.13501 [cs, math]*, Apr. 2022. Available: <http://arxiv.org/abs/2204.13501>.
- [5] P. Großmann, S. Hölldobler, N. Manthey, K. Nachtigall, J. Opitz, and P. Steinke. Solving Periodic Event Scheduling Problems with SAT. In *Advanced Research in Applied Artificial Intelligence*, Berlin, Heidelberg, 2012, pp. 166–175. doi: 10.1007/978-3-642-31087-4_18.
- [6] M. Joswig and K. Kulas. *Tropical and ordinary convexity combined*. Vol. 10, no. 2, pp. 333–352, Apr. 2010, doi: 10.1515/advegeom.2010.012.
- [7] M. Joswig and G. Loho. Weighted digraphs and tropical cones. *Linear Algebra and its Applications*, vol. 501, pp. 304–343, Jul. 2016, doi: 10.1016/j.laa.2016.02.027.
- [8] C. Liebchen and R. H. Möhring. The Modeling Power of the Periodic Event Scheduling Problem: Railway Timetables – and Beyond. In *Algorithmic Methods for Railway Optimization*, Berlin, Heidelberg, 2007, pp. 3–40. doi: 10.1007/978-3-540-74247-0_1.
- [9] N. Lindner and C. Liebchen. New Perspectives on PESP: T-Partitions and Separators. In *19th Symposium on Algorithmic Approaches for Transportation Modeling, Optimization, and Systems (ATMOS 2019)*, Dagstuhl, Germany, 2019, vol. 75, p. 2:1–2:18. doi: 10.4230/OASlcs.ATMOS.2019.2.
- [10] N. Lindner and C. Liebchen. Determining All Integer Vertices of the PESP Polytope by Flipping Arcs. In *20th Symposium on Algorithmic Approaches for Transportation Modeling, Optimization, and Systems (ATMOS 2020)*, Dagstuhl, Germany, 2020, vol. 85, p. 5:1–5:18. doi: 10.4230/OASlcs.ATMOS.2020.5.
- [11] N. Lindner and C. Liebchen. Timetable merging for the Periodic Event Scheduling Problem. *EURO Journal on Transportation and Logistics*, vol. 11, p. 100081, Jan. 2022, doi: 10.1016/j.ejtl.2022.100081.
- [12] N. Lindner, C. Liebchen, and B. Masing. Forward Cycle Bases and Periodic Timetabling. In *21st Symposium on Algorithmic Approaches for Transportation Modeling, Optimization, and Systems (ATMOS 2021)*, Dagstuhl, Germany, 2021, vol. 96, p. 2:1–2:14. doi: 10.4230/OASlcs.ATMOS.2021.2.
- [13] N. Lindner and J. Reisch. An analysis of the parameterized complexity of periodic timetabling. *J Sched*, vol. 25, no. 2, pp. 157–176, Apr. 2022, doi: 10.1007/s10951-021-00719-1.
- [14] K. Nachtigall and J. Opitz. Solving Periodic Timetable Optimisation Problems by Modulo Simplex Calculations. In *8th Workshop on Algorithmic Approaches for Transportation Modeling, Optimization, and Systems (ATMOS 2008)*, Dagstuhl, Germany, 2008, vol. 9. doi: 10.4230/OASlcs.ATMOS.2008.1588.
- [15] P. Serafini and W. Ukovich. A Mathematical Model for Periodic Scheduling Problems. *SIAM J. Discrete Math.*, vol. 2, no. 4, pp. 550–581, Nov. 1989, doi: 10.1137/0402049.

Programming on the Data Highway

- [1] A. Abel and J. Reineke. uops.info: characterizing latency, throughput, and port usage of instructions on Intel microarchitectures. *ASPLOS '19*, pp. 673–686, 2019.
- [2] A. Abel and J. Reineke. uiCA: accurate throughput prediction of basic blocks on recent Intel microarchitectures. *International Conference on Supercomputing*, pp. 33:1–33:14, 2022.
- [3] R. Clausecker and F. Schintke. A measure of quality for IDA* heuristics. *Proceedings of the International Symposium on Combinatorial Search*, 12(1):55–63, 2021.
- [4] A. Fog. Optimization Manuals. *Self-published 2022*, vol. 1–5. Published online at <https://www.agner.org/optimize/>.
- [5] H. Inoue et al. Accelerating UTF-8 decoding using SIMD instructions. *Information Processing Society of Japan, Transactions on Programming* 1(2):1–8, 2008.
- [6] M. Klarqvist et al. Efficient computation of positional population counts using SIMD instructions. *Preprint available at https://arxiv.org/abs/1911.02696*.
- [7] D. Lemire and W. Muta. Transcoding billions of Unicode characters per second with SIMD instructions. *Software: Practice and Experience*, 52(2):555–575, 2022.
- [8] W. Shen et al. KMCP: accurate metagenomic profiling of both prokaryotic and viral populations by pseudo-mapping. *Bioinformatics* 39(1):btac845, 2023.
- [9] H. Warren. Hacker's delight. Second Edition. *Pearson Education*, 2013.

Addressing Energy Challenges in HPC Data Centers

- [1] <https://www.top500.org/lists/green500/>.
- [2] D. D'Agostino, I. Merelli, M. Aldinucci, and D. Cesini. Hardware and software solutions for energy-efficient computing in scientific programming. *Scientific Programming*, 2021, doi: [10.1155/2021/5514284](https://doi.org/10.1155/2021/5514284).
- [3] R. Lucas et al. Doe advanced scientific computing advisory subcommittee (ascac) report: top ten exascale research challenges. *USDOE Office of Science* (United States), 2014.
- [4] J. Enos, C. P. Steffen, J. Fullop, M. T. Showerman, G. Shi, K. Esler, V. V. Kindratenko, J. E. Stone, and J. C. Phillips. Quantifying the impact of GPUs on performance and energy efficiency in HPC clusters. *Green Computing Conference*, 317–324, 2010.
- [5] R. Okazaki et al. Supercomputer Fugaku Cpu A64fx realizing high-performance, high-density packaging, and low power consumption. *Fujitsu Technical Review*, 2020.
- [6] E. Danovaro, A. Clematis, A. Galizia, G. Ripepi, A. Quarati, and D. D'Agostino. Heterogeneous architectures for computational intensive applications: a cost-effectiveness analysis. *Journal of Computational and Applied Mathematics*, 270, 63–77, 2014.
- [7] T. Kalaiselvi, P. Sriramakrishnan, and K. Somasundaram. Survey of using GPU CUDA programming model in medical image analysis. *Informatics in Medicine Unlocked*, 9, 133–144, 2017.
- [8] A. Marathe, P. E. Bailey, D. K. Lowenthal, B. Rountree, M. Schulz, and B. R. de Supinski. A run-time system for power-constrained HPC applications. In *International conference on high-performance computing*: pp. 394–408, 2015.
- [9] <https://www.gromacs.org/>.
- [10] https://gitlab.com/elsi_project/elsi_benchmark/-/tree/master/Periodic/Liquids/Water_64molecules_PEXSI.
- [11] G. Brook, D. Fuller, J. Swinburne, S. Christgau, M. Läuter, R. Pela, L. Stein, C. Tuma, and T. Steinke. An Early Scalability Study of Omni-Path Express. *ISC High Performance*, 2022.
- [12] H. David, C. Fallin, E. Gorbato, U. Hanebutte, and O. Mutlu. Memory power management via dynamic voltage/frequency scaling. *ICAC '11 Proceedings*, 2011.
- [13] V. Sundriyal and M. Sosonkina. Modeling of the CPU frequency to minimize energy consumption in parallel applications. *Sustainable Computing: Informatics and Systems*, 17, 2018.
- [14] R. Schöne, T. Ilsche, M. Bielert, A. Gocht, and D. Hackenberg. Energy Efficiency Features of the Intel Skylake-SP Processor and Their Impact on Performance. *International Conference on High Performance Computing & Simulation*, 399–406, 2019.
- [15] B. Wang, D. Schmidl, and M. S. Müller. Evaluating the Energy Consumption of OpenMP Applications on Haswell Processors. In: *Lecture Notes in Computer Science, Springer*, 2015.
- [16] T. Scogland, B. Subramaniam, and W. Feng. The Green500 list: escapades to exascale. *Comput. Sci. Res. Dev.* 28(2–3): 109–117, 2013.
- [17] Yu-Xiang Gao, Peng Zhang. A Survey of Homogeneous and Heterogeneous System Architectures in *High Performance Computing. SmartCloud*. 170–175, 2016.
- [18] S. Kozielski and D. Mrozek. Development of High-Performance Computing Systems. *Computer Network*: 52–63, 2020.
- [19] M. Sato, Y. Kodama, M. Tsuji, and T. Odajima. Co-Design and System for the Supercomputer "Fugaku". *IEEE Micro* 42(2): 26–34, 2022.
- [20] M. Qasaimeh, K. Denolf, J. Lo, K. A. Vissers, J. Zambreno, and P. H. Jones. Comparing Energy Efficiency of CPU, GPU, and FPGA Implementations for Vision Kernels. *ICESS*: 1–8, 2019.
- [21] J. Knapheide, B. Stabernack, and M. Kuhnke. A High Throughput MobileNetV2 FPGA Implementation Based on a Flexible Architecture for Depthwise Separable Convolution. *FPL*: 277–283, 2020.
- [22] M. Knaust, E. Seiler, K. Reinert, and T. Steinke. Co-Design for Energy-Efficient and Fast Genomic Search: Interleaved Bloom Filter on FPGA. *FPGA*: 180–189, 2022.
- [23] M. Knaust, F. Mayer, and T. Steinke. OpenMP to FPGA Offloading Prototype Using OpenCL SDK. *IPDPS Workshop*, 387–390, 2019.
- [24] S. Christgau, M. Knaust, and T. Steinke. A First Step Towards Support for MPI Partitioned Communication on SYCL-programmed FPGAs. *IEEE/ACM International Workshop on Heterogeneous High-performance Reconfigurable Computing*, H2RC@SC 2022, Dallas, TX, USA, November (2022, to appear).
- [25] F. Steinert, N. Schelten A. Schulte, and B. Stabernack. Hardware and Software Components towards the Integration of Network-Attached Accelerators into *Data Centers*. 149–153, 2020.
- [26] C. Kutzner: *A free GROMACS benchmark set*, <https://www.mpi-nat.mpg.de/grubmueller/bench>.
- [27] https://www.zib.de/de/research_services/supercomputing.

PUBLICATIONS

Peer-reviewed (154)

*Samer Alhaddad, Jens Förstner, Stefan Groth, Daniel Grünewald, Yevgen Grynko, Frank Hannig, Tobias Kenter, F. J. Pfreundt, Christian Plessl, Merlind Schotte, Thomas Steinke, J. Teich, Martin Weiser, Florian Wende (2022). The HighPerMeshes Framework for Numerical Algorithms on Unstructured Grids. *Concurrency and Computation: Practice and Experience*, 34(14). <https://doi.org/10.1002/cpe.6616> (Joint publication: Modeling and Simulation of Complex Processes, Distributed Algorithms, and Supercomputing)

Tamaz Amiranashvili, David Lüdke, Hongwei Li, Bjoern Menze, Stefan Zachow (2022). Learning Shape Reconstruction from Sparse Measurements with Neural Implicit Functions. In *Medical Imaging with Deep Learning*.

Lovis Anderson, Mark Turner, Thorsten Koch (2022). Generative deep learning for decision making in gas networks. *Mathematical Methods of Operations Research*, 95:503-532. <https://doi.org/10.1007/s00186-022-00777-x>

Oscar Avalos-Ovando, Eva Yazmin Santiago, Artur Movsesyan, Xiang-Tian Kong, Peng Yu, Lucas V. Besteiro, Larousse Khosravi Khoshashad, Hiromi Okamoto, Joseph M. Slocik, Miguel Correa-Duarte, Miguel Comesana-Hermo, Tim Liedl, Gil Markovich, Sven Burger, Alexander O. Govorov (2022). Chiral Bioinspired Plasmonics: A Paradigm Shift for Optical Activity and Photochemistry. *ACS Photon.*, 9:2219. <https://doi.org/10.1021/acsp Photonics.2c00445>

Birzhan Ayanbayev, Ilja Klebanov, Han Cheng Lie, T. J. Sullivan (2022). Γ -convergence of Onsager-Machlup functionals: I. With applications to maximum a posteriori estimation in Bayesian inverse

problems. *Inverse Problems*, 38(2). <https://doi.org/10.1088/1361-6420/ac3f81>

Birzhan Ayanbayev, Ilja Klebanov, Han Cheng Lie, T. J. Sullivan (2022). Γ -convergence of Onsager-Machlup functionals: II. Infinite product measures on Banach spaces. *Inverse Problems*, 38(2). <https://doi.org/10.1088/1361-6420/ac3f82>

Jonathan Bader, Joel Witzke, Soeren Becker, Ansgar Löber, Fabian Lehmann, Leon Doeher, Anh Duc Vu, Odej Kao (2022). Towards Advanced Monitoring for Scientific Workflows. In *2022 IEEE International Conference on Big Data, 11th Workshop on Scalable Cloud Data Management*.

Kaitlyn P Becker, Clark Teeple, Nicholas Charles, Yeonsu Jung, Daniel Baum, James C Weaver, L. Mahadevan, Robert J Wood (2022). Active entanglement enables stochastic, topological grasping. *PNAS*, 119(42):e2209819119. <https://doi.org/10.1073/pnas.2209819119>

Fidji Berio, Yann Bayle, Daniel Baum, Nicolas Goudemand, Mélanie Debais-Thibaud (2022). Hide and seek shark teeth in Random Forests: Machine learning applied to Scyliorhinus canicula populations. *PeerJ - Aquatic Biology*, 10:e13575. <https://doi.org/10.7717/peerj.13575>

Julius Berner, Lorenz Richter, Karen Ullrich (2022). An optimal control perspective on diffusion-based generative modeling. In *NeurIPS 2022 Workshop on Score-Based Methods*.

Ksenia Bestuzheva, Ambros Gleixner, Helena Völker (2022). Strengthening SONC Relaxations with Constraints Derived from

Variable Bounds. In *Proceedings of the Hungarian Global Optimization Workshop HUGO 2022*, 41-44.

Felix Binkowski, Fridtjof Betz, Martin Hammerschmidt, Philipp-Immanuel Schneider, Lin Zschiedrich, Sven Burger (2022). Computation of eigenfrequency sensitivities using Riesz projections for efficient optimization of nanophotonic resonators. *Commun. Phys.*, 5:202. <https://doi.org/10.1038/s42005-022-00977-1>

Felix Binkowski, Fridtjof Betz, Martin Hammerschmidt, Philipp-Immanuel Schneider, Lin Zschiedrich, Sven Burger (2022). Numerical methods for the investigation of resonances in nanophotonics. In S. Zouhdi, editor, *The 12th International Conference on Metamaterials, Photonic Crystals and Plasmonics (META 2022)*, 247.

Horst-Holger Boltz, Alexei Sirbu, Nina Stelzer, Primal de Lanerolle, Stefanie Winkelmann, Paolo Annibale (2022). The Impact of Membrane Protein Diffusion on GPCR Signaling. *Cells*, 11(10):1660. <https://doi.org/10.3390/cells11101660>

Julia Boltze, Annika Höllerl, Monika Kuberek, Stefan Lohrum, Heinz Pampel, Markus Putnings, Regina Retter, Beate Rusch, Hildegard Schäffler, Konstanze Söllner (2022). DeepGreen: Eine Infrastruktur für die Open-Access-Transformation. *O-Bib. Das Offene Bibliotheksjournal*, 9(1). <https://doi.org/10.5282/o-bib/5764>

Ralf Borndörfer, Stephan Schwartz, William Surau (2023). Vertex Covering with Capacitated Trees. *Networks*, 81(2):253-277. <https://doi.org/10.1002/net.22130> (epub ahead of print 2022-11-10)

Ralf Borndörfer, Stephan Schwartz,

William Surau (2022). Rooted Maximum Weight Connected Subgraphs with Balancing and Capacity Constraints. In *Proceedings of the 10th International Network Optimization Conference (INOC), Aachen, Germany, June 7-10, 2022*, 63-68. <https://doi.org/10.48786/inoc.2022.12>

*Ralf Borndörfer, Fabian Danecker, Martin Weiser (2022). A Discrete-Continuous Algorithm for Globally Optimal Free Flight Trajectory Optimization. In *22nd Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2022)*, volume 106 of Open Access Series in Informatics (OASlcs), pages 1-13. <https://doi.org/10.4230/OASlcs.ATMOS.2022.2> (Joint publication: Modeling and Simulation of Complex Processes, Network Optimization)

Enrico Bortoletto, Niels Lindner, Berenike Masing (2022). Tropical Neighbourhood Search: A New Heuristic for Periodic Timetabling. In *22nd Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2022)*, volume 106 of Open Access Series in Informatics (OASlcs), pages 3:1-3:19. <https://doi.org/10.4230/OASlcs.ATMOS.2022.3>

Lucas Bremer, Carlos Jimenez, Simon Thiele, Ksenia Weber, Tobias Huber, Sven Rodt, Alois Herkommer, Sven Burger, Sven Höfling, Harald Giessen, Stephan Reitzenstein (2022). Numerical optimization of single-mode fiber-coupled single-photon sources based on semiconductor quantum dots. *Opt. Express*, 30:15913. <https://doi.org/10.1364/OE.456777>

Glenn Brook, Douglas Fuller, John Swinburne, Steffen Christgau, Mat-

- thias Lauter, Ronaldo Rodrigues Pela, Stein Lewin, Tuma Christian, Thomas Steinke (2022). An Early Scalability Study of Omni-Path Express. In *ISC High Performance*. <https://doi.org/10.13140/RG.2.2.21353.57442>
- Ariel Caputo, Marco Emporio, Andrea Giachetti, Marco Cristani, Guido Borghi, Andrea D'Eusano, Minh-Quan Le, Hai-Dang Nguyen, Minh-Triet Tran, Felix Ambellan, Martin Hanik, Esfandiar Navayazdani, Christoph von Tycowicz (2022). SHREC 2022 Track on Online Detection of Heterogeneous Gestures. *Computers and Graphics*, 107:241-251. <https://doi.org/10.1016/j.cag.2022.07.015>
- Fatemeh Chegini, Alena Kopanickova, Martin Weiser, Rolf Krause (2022). Quantitative Analysis of Nonlinear Multifidelity Optimization for Inverse Electrophysiology. In *Domain Decomposition Methods in Science and Engineering XXVI*, 65-76.
- Steffen Christgau, Marius Knaust, Thomas Steinke (2022). A First Step towards Support for MPI Partitioned Communication on SYCL-programmed FPGAs. In *IEEE/ACM International Workshop on Heterogeneous High-performance Reconfigurable Computing, H2RC@SC 2022, Dallas, TX, USA, November 13-18, 2022*, 9-17. <https://doi.org/10.1109/H2RC56700.2022.00007>
- Remi Colom, Felix Binkowski, Fridtjof Betz, Yuri Kivshar, Sven Burger (2022). Enhanced Purcell factor for nanoantennas supporting interfering resonances. *Phys. Rev. Research*, 4:023189. <https://doi.org/10.1103/PhysRevResearch.4.023189>
- Natasa Djurdjevac Conrad, Jonas Koppl, Ana Djurdjevac (2022). Feedback Loops in Opinion Dynamics of Agent-Based Models with Multiplicative Noise. *Entropy*, 24(10). <https://doi.org/10.3390/e24101352>
- Francisco Criado, David Martınez-Rubio, Sebastian Pokutta (2022). Fast Algorithms for Packing Proportional Fairness and its Dual. In *Proceedings of the Conference on Neural Information Processing Systems*, 36.
- Mihir Dass, Lilli Kuen, Gregor Posnjak, Sven Burger, Tim Liedl (2022). Visible wavelength spectral tuning of absorption and circular dichroism of DNA-assembled Au/Ag core-shell nanorod assemblies. *Mater. Adv.*, 3:3438. <https://doi.org/10.1039/D1MA01211H>
- Siva Prasad Chakri Dhanakoti, John Maddocks, Martin Weiser (2022). Navigation of Concentric Tube Continuum Robots using Optimal Control. In *Proceedings of the 19th International Conference on Informatics in Control, Automation, and Robotics*, 146-154. <https://doi.org/10.5220/0011271000003271>
- Luca Donati, Marcus Weber, Bettina G. Keller (2022). A review of Girsanov Reweighting and of Square Root Approximation for building molecular Markov State Models. *Journal of Mathematical Physics*, 63(12):123306-1-123306-21. <https://doi.org/10.1063/5.0127227>
- Sarah Ehlers, Daniel Baum, Roland Muhlethaler, Hannelore Hoch, Peter Braunig (2022). Large abdominal mechanoreceptive sense organs in small plant-dwelling insects. *Biology Letters*, 18(4). <https://doi.org/10.1098/rsbl.2022.0078>
- Leon Eifler, Ambros Gleixner (2022). A computational status update for exact rational mixed integer programming. *Mathematical Programming*. <https://doi.org/10.1007/s10107-021-01749-5>
- Leon Eifler, Ambros Gleixner, Jonad Pulaj (2022). A Safe Computational Framework for Integer Programming applied to Chvatal's Conjecture. *ACM Transactions on Mathematical Software*, 48(2). <https://doi.org/10.1145/3485630>
- Lennart Eigen, Daniel Baum, Mason N. Dean, Daniel Werner, Jan Wolfer, John A. Nyakatura (2022). Ontogeny of a tessellated surface: carapace growth of the longhorn cowfish *Lactoria cornuta*. *Journal of Anatomy*, 241(3):565-580. <https://doi.org/10.1111/joa.13692>
- Franziska Erlekan, Maximilian Zumbansen, Marcus Weber (2022). Parameter estimation on multivalent ITC data sets. *Scientific Reports*, 12:13402. <https://doi.org/10.1038/s41598-022-17188-x>
- Ariane Ernst, Christof Schutte, Stephan Sigrist, Stefanie Winkelmann (2022). Variance of filtered signals: Characterization for linear reaction networks and application to neurotransmission dynamics. *Mathematical Biosciences*, 343. <https://doi.org/10.1016/j.mbs.2021.108760>
- Ricardo Euler, Ralf Borndorfer, Timo Strunk, Tuomo Takkula (2022). ULD Build-Up Scheduling with Dynamic Batching in an Air Freight Hub. In *Operations Research Proceedings 2021, Lecture Notes in Operations Research*, pages 254-260. https://doi.org/10.1007/978-3-031-08623-6_38
- Konstantin Fackeldey, Mathias Oster, Leon Sallandt, Reinhold Schneider (2022). Approximative Policy Iteration for Exit Time Feedback Control Problems driven by Stochastic Differential Equations using Tensor Train format. *SIAM Journal on Multiscale Modeling and Simulation*, 20(1):379-403. <https://doi.org/10.1137/20M1372500>
- Rita Fernandes, Suvrat Chakraborty, Noureldin Saleh, Natalia Mikula, Katerina Kanevche, Hans Berlepsch, Naoki Hosogi, Joachim Heberle, Marcus Weber, Christoph Bottcher, Beate Koksche (2022). Cyanine Dye Coupling Mediates Self-assembly of a pH Sensitive Peptide into Novel 3D Architectures. *Angewandte Chemie*, 61(48):e202208647.
- Sophie Fischer-Holzhausen, Susanna Roblitz (2022). Hormonal regulation of ovarian follicle growth in humans: Model-based exploration of cycle variability and parameter sensitivities. *Journal of Theoretical Biology*. <https://doi.org/https://doi.org/10.1016/j.jtbi.2022.111150>
- Sophie Fischer-Holzhausen, Susanna Roblitz (2022). Mathematical modelling of follicular growth and ovarian stimulation. *Current Opinion in Endocrine and Metabolic Research*. <https://doi.org/https://doi.org/10.1016/j.coemr.2022.100385>
- Steffen Furst, Christian Rakow (2022). Parallelization of large-scale agent-based epidemiological simulations. In *Proceedings of the 34th European Modeling & Simulation Symposium*. <https://doi.org/https://doi.org/10.46354/i3m.2022.emss.008>
- Timm Gao, Lucas Rickert, Felix Urban, Jan Groe, Nicole Srocka, Sven Rodt, Anna Musial, Kinga Zolnacz, Pawel Mergo, Kamil Dybka, Wacław Urbanczyk, Grzegorz Sek, Sven Burger, Stephan Reitzenstein, Tobias Heindel (2022). A quantum key distribution testbed

- using a plug&play telecom-wave-length single-photon source. *Appl. Phys. Rev.*, 9:011412. <https://doi.org/10.1063/5.0070966>
- Patrick Gelß, Rupert Klein, Sebastian Matera, Burkhard Schmidt (2022). Solving the time-independent Schrödinger equation for chains of coupled excitons and phonons using tensor trains. *The Journal of Chemical Physics*, 156(2):024109. <https://doi.org/10.1063/5.0074948>
- Masoud Gholami, Florian Schintke (2022). IOSIG: Declarative I/O-Stream Properties Using Pragmas. *Datenbank-Spektrum*, 22(2):109-119. <https://doi.org/10.1007/s13222-022-00419-w>
- Korbinian Glatzeder, Igor Komnik, Felix Ambellan, Stefan Zachow, Wolfgang Potthast (2022). Dynamic pressure analysis of novel interpositional knee spacer implants in 3D-printed human knee models. *Scientific Reports*, 12. <https://doi.org/10.1038/s41598-022-20463-6>
- Christoph Gorgulla, Abhilash Jayaraj, Konstantin Fackeldey, Haribabu Arthanari (2022). Emerging frontiers in virtual drug discovery: From quantum mechanical methods to deep learning approaches. *Current Opinion in Chemical Biology*, 69:102156-102156-12. <https://doi.org/10.1016/j.cbpa.2022.102156>
- Carl Martin Grewe, Tuo Liu, Andrea Hildebrandt, Stefan Zachow (2022). The Open Virtual Mirror Framework for Enfacement Illusions – Enhancing the Sense of Agency With Avatars That Imitate Facial Expressions. *Behavior Research Methods*. <https://doi.org/10.3758/s13428-021-01761-9>
- Noushin Hajarolasvadi, Vikram Sunkara, Sagar Khavnekar, Florian Beck, Robert Brandt, Daniel Baum (2022). Volumetric macromolecule identification in cryo-electron tomograms using capsule networks. *BMC Bioinformatics*, 23(360). <https://doi.org/10.1186/s12859-022-04901-w>
- Martin Hanik, Hans-Christian Hege, Christoph von Tycowicz (2022). A Nonlinear Hierarchical Model for Longitudinal Data on Manifolds. In *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*, 1-5. <https://doi.org/10.1109/ISBI52829.2022.9761465>
- Martin Hanik, Hans-Christian Hege, Christoph von Tycowicz (2022). Bi-invariant Dissimilarity Measures for Sample Distributions in Lie Groups. *SIAM Journal on Mathematics of Data Science*, 4(4):1223-1249. <https://doi.org/10.1137/21M1410373>
- Martin Hanik, Mehmet Arif Demirtaş, Mohammed Amine Gharsallaoui, Islem Rekik (2022). Predicting cognitive scores with graph neural networks through sample selection learning. *Brain Imaging and Behavior*, 16:1123-1138. <https://doi.org/10.1007/s11682-021-00585-7>
- Philipp Harth, Sumit Vohra, Daniel Udvary, Marcel Oberlaender, Hans-Christian Hege, Daniel Baum (2022). A Stratification Matrix Viewer for Analysis of Neural Network Data. In *Eurographics Workshop on Visual Computing for Biology and Medicine (VCBM)*. <https://doi.org/10.2312/vcbm.20221194>
- Kai Hoppmann-Baum, Oleg Burdakov, Gioni Mexi, Carl Johan Casselgren, Thorsten Koch (2022). Length-Constrained Cycle Partition with an Application to UAV Routing. *Optimization Methods and Software*. Vol. 37, Issue 6 <https://doi.org/10.1080/10556788.2022.2053972>
- Junko Hosada, Stephen Maher, Yuji Shinano, Jonas Christoffer Villumsen (2022). Location, transshipment and routing: An adaptive transportation network integrating long-haul and local vehicle routing. *EURO Journal on Transportation and Logistics*, 11. <https://doi.org/10.1016/j.ejtl.2022.100091>
- Christoph Hunkenschröder, Sebastian Pokutta, Robert Weismantel (2022). Optimizing a low-dimensional convex function over a high-dimensional cub. *SIAM Journal on Optimization*.
- Christian Jung, Anna Nowacka, Tin Barisin, Dietmar Meinel, Olaf Paetsch, Szymon Grzesiak, Michael Salamon, Katja Schladitz, Claudia Redenbach, Matthias Pahn (2023). 3D imaging and analysis of cracks in loaded concrete samples. In *12th Conference on Industrial Computed Tomography, Fürth, Germany (iCT 2023)*. (accepted for publication 2022-10-04)
- Klaus Jäger, Johannes Sutter, Philipp Tockhorn, Martin Hammer-schmidt, Philipp-Immanuel Schneider, Steve Albrecht, Christiane Becker (2022). Improving perovskite/silicon tandem solar cells with optical simulations. In *Proc. SPIE*, PC12150:PC1215001. <https://doi.org/10.1117/12.2621255>
- Nina Kamčev, Christoph Spiegel (2022). Another note on intervals in the Hales-Jewett theorem. *Electronic Journal of Combinatorics*, 29(1):P1.62. <https://doi.org/10.37236/6730>
- Robert Kiewisz, Gunar Fabig, William Conway, Daniel Baum, Daniel Needleman, Thomas Müller-Reichert (2022). Three-dimensional structure of kinetochore-fibers in human mitotic spindles. *eLife*, 11:e75459. <https://doi.org/10.7554/eLife.75459>
- Torsten Klug, Markus Reuther, Thomas Schlechte (2022). Does Laziness Pay Off? – A Lazy-Constraint Approach to Timetabling. In: *22nd Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2022) of Open Access Series in Informatics (OASICS)*. <https://doi.org/10.4230/OASICS.ATMOS.2022.11>
- Stefan Klus, Natasa Djurdjevac Conrad (2022). Koopman-based spectral clustering of directed and time-evolving graphs. *Journal of Nonlinear Science*, 33. <https://doi.org/10.1007/s00332-022-09863-0>
- Marius Knaust, Enrico Seiler, Knut Reinert, Thomas Steinke (2022). Co-Design for Energy Efficient and Fast Genomic Search: Interleaved Bloom Filter on FPGA. In *FPGA '22: Proceedings of the 2022 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, 180-189. <https://doi.org/10.1145/3490422.3502366>
- Thorsten Koch, Timo Berthold, Jaap Pedersen, Charlie Vanaret (2022). Progress in mathematical programming solvers from 2001 to 2020. *EURO Journal on Computational Optimization*, 10:100031. <https://doi.org/10.1016/j.ejco.2022.100031>
- Tabea Kossen, Manuel A. Hirzel, Vince I. Madai, Franziska Boenisch, Anja Hennemuth, Kristian Hildebrand, Sebastian Pokutta, Kartikey Sharma, Adam Hilbert, Jan Sobesky, Ivana Galinovic, Ahmed A. Khalil, Jochen B. Fiebach, Dietmar Frey (2022). Towards Sharing Brain Images: Differentially Private TOF-MRA Images with Segmentation Labels Using Generative Adversarial Networks. *Frontiers in Artificial Intelligence*. <https://doi.org/https://doi.org/10.3389/frai.2022.813842>
- *Margarita Kostré, Vikram Sunkara, Christof Schütte, Natasa Djurdjevac Conrad (2022). Understanding the Romanization Spreading on Historical Interregional Networks in Northern Tunisia. *Applied*

Network Science, 7. <https://doi.org/10.1007/s41109-022-00492-w> (Joint publication: Modeling and Simulation of Complex Processes, Visual, and Data-centric Computing)

Jan Krüger, Phillip Manley, Lin Zschiedrich, Martin Hamerschmidt, Philipp-Immanuel Schneider, Rainer Köning, Bernd Bodermann (2022). Untersuchung des Schwingungseinflusses auf gemessene Punktspreizfunktionen in der optischen Mikroskopie. In *Proc. DGAO*, P5.

Tillmann Krüger, Brigitte Leeners, Enrico Tronci, Toni Mancini, Fabian Ille, Marcel Egli, Harald Engler, Susanna Röblitz, Helge Frieling, Christopher Sinke, Kirsten Jahn (2022). The androgen system across the menstrual cycle: Hormonal, (epi-)genetic, and psychometric alterations. *Physiology and Behavior*. <https://doi.org/https://doi.org/10.1016/j.physbeh.2022.114034>

Alejandra Laguillo-Diego, Robert Kiewisz, Carlos Martí-Gómez, Daniel Baum, Thomas Müller-Reichert, Isabelle Vernos (2022). MCRS1 modulates the heterogeneity of microtubule minus-end morphologies in mitotic spindles. *Molecular Biology of the Cell*, 34(1). <https://doi.org/10.1091/mbc.E22-08-0306-T>

Annemarie Lang, Antja Vester, M. Ferreira-Gomes, G.M. Guerra, Fredrich R. Heinrich, C.C. Götzke, Sebastian Kurmies, Vikram Sunkara, Pawel Durek, Joel D. Boerckel, Mir-Farzin Mashreghi (2022). Pipeline for single cell sequencing of human chondrocyte pellet cultures to delineate il-1 β modulated changes in cell heterogeneity. *Osteoarthritis and Cartilage*, 30(Supplement 1):S90. <https://doi.org/10.1016/j.joca.2022.02.112>

Thi Thai Le, Yasuhide Fukumoto (2022). Effect of depth discontinuity on interfacial stability of

tangential-velocity discontinuity in shallow-water flow. *Physics Letters A*, 436:128073. <https://doi.org/10.1016/j.physleta.2022.128073>

Thi Thai Le, Thorsten Koch (2022). Viscosity and porosity effects on tangential-discontinuity surface stability in 3D Compressible Media. *Physics of Fluids*, 34(7). <https://doi.org/10.1063/5.0095970>

Tiago Ramos Leite, Lin Zschiedrich, Orhan Kizilkaya, Kevin M. McPeak (2022). Resonant Plasmonic-Biomolecular Chiral Interactions in the Far-Ultraviolet: Enantiomeric Discrimination of sub-10 nm Amino Acid Films. *Nano Lett.*, 22:7343. <https://doi.org/10.1021/acs.nanolett.2c01724>

Tony Lelievre, Gabriel Stoltz, Wei Zhang (2022). Multiple projection MCMC algorithms on submanifolds. *IMA Journal of Numerical Analysis*. <https://doi.org/10.1093/imanum/drac006>

Ralf Lenz, Felipe Serrano (2022). Tight Convex Relaxations for the Expansion Planning Problem. *Journal of Optimization Theory and Applications*, 194:325-352. <https://doi.org/10.1007/s10957-022-02029-8>

Ralf Lenz, Kai-Helge Becker (2022). Optimization of Capacity Expansion in Potential-driven Networks including Multiple Looping – A comparison of modelling approaches. *OR Spectrum*, 44:179-224. <https://doi.org/https://doi.org/10.1007/s00291-021-00648-7>

Niels Lindner, Christian Liebchen (2022). Timetable merging for the Periodic Event Scheduling Problem. *EURO Journal on Transportation and Logistics*, 11:100081. <https://doi.org/10.1016/j.ejtl.2022.100081>

Niels Lindner, Rolf van Lieshout

(2022). Benders Decomposition for the Periodic Event Scheduling Problem. In *Operations Research Proceedings 2021*, 289-294. https://doi.org/10.1007/978-3-031-08623-6_43

Niels Lindner, Julian Reisch (2022). An analysis of the parameterized complexity of periodic timetabling. *Journal of Scheduling*, 25:157-176. <https://doi.org/10.1007/s10951-021-00719-1>

Eduardo Lopez-Fraguas, Felix Binkowski, Sven Burger, Sylvia Hagedorn, Braulio Garcia-Camara, Ricardo Vergaz, Christiane Becker, Phillip Manley (2022). Tripling the light extraction efficiency of a deep ultraviolet LED using a nano-structured p-contact. *Sci. Rep.*, 12:11480. <https://doi.org/10.1038/s41598-022-15499-7>

Matthias Läuter, Tobias Kramer, Martin Rubin, Kathrin Altwegg (2022). The ice composition close to the surface of comet 67P/Churyumov-Gerasimenko. *ACS Earth and Space Chemistry*, 6(5):1189-1203. <https://doi.org/10.1021/acsearthspacechem.1c00378>

Angar Löffler, Joel Witzke, Florian Schintke, Björn Scheuermann (2022). BottleMod: Modeling Data Flows and Tasks for Fast Bottleneck Analysis. In *2022 IEEE International Conference on Big Data, Posters*.

Marvin Lücke, Jingyi Lu, Daniel E. Quevedo (2022). Coding for secrecy in remote state estimation with an adversary. *IEEE Transactions on Automatic Control*, 67(9):4955-4962. <https://doi.org/10.1109/TAC.2022.3169839>

Marvin Lücke, Feliks Nüske (2022). tgEDMD: Approximation of the Kolmogorov Operator in Tensor Train Format. *Journal of Nonlinear Science*, 32. <https://doi.org/10.1007/s00332-022-09801-0>

David Lüdke, Tamaz Amiranashvili, Felix Ambellan, Ivan Ezhov, Bjoern Menze, Stefan Zachow (2022). Landmark-free Statistical Shape Modeling via Neural Flow Deformations. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, volume 13432 of *Lecture Notes in Computer Science*. https://doi.org/10.1007/978-3-031-16434-7_44

Jan MacDonald, Mathieu Besançon, Sebastian Pokutta (2022). Interpretable Neural Networks with Frank-Wolfe: Sparse Relevance Maps and Relevance Orderings. In *Proceedings of the International Conference on Machine Learning*.

Phillip Manley, Jan Krüger, Lin Zschiedrich, Martin Hamerschmidt, Bernd Bodermann, Rainer Köning, Philipp-Immanuel Schneider (2022). Elementary, my dear Zernike: model order reduction for accelerating optical dimensional microscopy. In *EPJ Web Conf.*, 266:10010. <https://doi.org/10.1051/epj-conf/202226610010>

Kevin Martens, Timon Funck, Eva Y. Santiago, Alexander O. Govorov, Sven Burger, Tim Liedl (2022). Onset of Chirality in Plasmonic Meta-Molecules and Dielectric Coupling. *ACS Nano*, 16:16143. <https://doi.org/10.1021/acsnano.2c04729>

David Martínez-Rubio, Sebastian Pokutta (2022). Accelerated Riemannian Optimization: Handling Constraints with a Prox to Bound Geometric Penalties. In *NeurIPS OPT*.

Julius Mayer, Daniel Baum, Felix Ambellan, Christoph von Tycowicz (2022). A Soft-Correspondence Approach to Shape-based Disease Grading with Graph Convolutional Networks. In *Proceedings of Machine Learning Research*.

- Natalia Mikula, Tom Dörrfel, Daniel Baum, Hans-Christian Hege (2022). An Interactive Approach for Identifying Structure Definitions. *Computer Graphics Forum*, 41(3):321-332. <https://doi.org/10.1111/cgf.14543>
- Saeed Mohammadzadeh, Vitor H. Nascimento, Rodrigo C. de Lamare, Noushin Hajarolasvadi (2022). Robust Beamforming Based on Complex-Valued Convolutional Neural Networks for Sensor Arrays. *IEEE Signal Processing Letters*, 29:2018-2021. <https://doi.org/10.1109/LSP.2022.3212637>
- Zahra Mokhtari, Robert I. A. Paterson, Felix Höfling (2022). Spontaneous trail formation in populations of auto-chemotactic walkers. *New Journal of Physics*, 24:013012. <https://doi.org/10.1088/1367-2630/ac43ec>
- Mattes Mollenhauer, Stefan Klus, Christof Schütte, Péter Koltai (2022). Kernel Autocovariance Operators of Stationary Processes: Estimation and Convergence. *Journal of Machine Learning Research*, 23(327):1-34.
- Alberto Montefusco, Christof Schütte, Stefanie Winkelmann (2022). A route to the hydrodynamic limit of a reaction-diffusion master equation using gradient structures. *SIAM Journal on Applied Mathematics*.
- Artur Movsesyan, Eva Yazmin Santiago, Sven Burger, Miguel Correa-Duarte, Lucas V. Besteiro, Zhiming Wang, Alexander O. Govorov (2022). Plasmonic nanocrystals with complex shapes for photocatalysis and growth: Contrasting anisotropic hot-electron generation with the photothermal effect. *Adv. Opt. Mater.*, 10:2102663. <https://doi.org/10.1002/adom.202102663>
- Artur Movsesyan, Eva Yazmin Santiago, Sven Burger, Miguel Correa-Duarte, Lucas V. Besteiro, Zhiming Wang, Alexander O. Govorov (2022). The anisotropy of hot carriers' spatial distribution contrasts the isotropy of photothermal effects in complex and small plasmonic nanocrystals with complex shapes. In S. Zouhdi, editor, *The 12th International Conference on Metamaterials, Photonic Crystals, and Plasmonics* (META 2022), 285.
- Adele Myers, Saiteja Utpala, Shubham Talbar, Sophia Sanborn, Christian Shewmake, Claire Donnat, Johan Mathe, Umberto Lupo, Rishi Sonthalia, Xinyue Cui, Tom Szwagier, Arthur Pignet, Andri Bergsson, Søren Hauberg, Dmitriy Nielsen, Stefan Sommer, David Klindt, Erik Hermansen, Melvin Vaupel, Benjamin Dunn, Jeffrey Xiong, Noga Aharoni, Itsik Pe'er, Felix Ambellan, Martin Hanik, Esfandiar Navayazdani, Christoph von Tycowicz, Nina Miodlane (2022). ICLR 2022 Challenge for Computational Geometry & Topology: Design and Results. In *Proceedings of Topology, Algebra, and Geometry in Learning*, volume 196 of *Proceedings of Machine Learning Research*, pages 269-276.
- Benjamin Müller, Gonzalo Muñoz, Maxime Gasse, Ambros Gleixner, Andrea Lodi, Felipe Serrano (2022). On generalized surrogate duality in mixed-integer nonlinear programming. *Mathematical Programming*, 192(1):89-118. <https://doi.org/10.1007/s10107-021-01691-6>
- Esfandiar Nava-Yazdani, Hans-Christian Hege, Christoph von Tycowicz (2022). A Hierarchical Geodesic Model for Longitudinal Analysis on Manifolds. *Journal of Mathematical Imaging and Vision*, 64(4):395-407. <https://doi.org/10.1007/s10851-022-01079-x>
- Yoel Negrin-Montecelo, Artur Movsesyan, Jie Gao, Sven Burger, Zhiming Wang, Sylvain Nlate, Emilie Pouget, Reiko Oda, Miguel Comesana-Hermo, Alexander O. Govorov, Miguel Correa-Duarte (2022). Chiral Generation of Hot Carriers for Polarization-Sensitive Plasmonic Photocatalysis with Hybrid Nanostructures. *J. Am. Chem. Soc.*, 144:1663. <https://doi.org/10.1021/jacs.1c10526>
- Patrick E Obermeier, Albert Heim, Barbara Biere, Elias Hage, Maren Alchikh, Tim Conrad, Brunhilde Schweiger, Barbara A Rath (2022). Linking digital surveillance and in-depth virology to study clinical patterns of viral respiratory infections in vulnerable patient populations. *iScience*, 25(5). <https://doi.org/10.1016/j.isci.2022.104276>
- *Stephan Oehme, Philippe Moewis, Heide Boeth, Benjamin Bartek, Annika Lippert, Christoph von Tycowicz, Rainald Ehrig, Georg Duda, Tobias Jung (2022). PCL insufficient patients with increased translational and rotational passive knee joint laxity have no increased range of anterior-posterior and rotational tibiofemoral motion during level walking. *Scientific Reports*, 12(1):1-11. <https://doi.org/10.1038/s41598-022-17328-3> (Joint publication: Modeling and Simulation of Complex Processes, Visual, and Data-centric Computing)
- Olaf Parczyk, Sebastian Pokutta, Christoph Spiegel, Tibor Szabó (2023). Fully Computer-assisted Proofs in Extremal Combinatorics. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 37. (accepted for publication 2022-11-19)
- Olaf Parczyk, Sebastian Pokutta, Christoph Spiegel, Tibor Szabó (2022). New Ramsey Multiplicity Bounds and Search Heuristics. In *Discrete Mathematics Days*.
- Martha Paskin, Mason Dean, Daniel Baum, Christoph von Tycowicz (2022). A Kendall Shape Space Approach to 3D Shape Estimation from 2D Landmarks. In *Computer Vision - ECCV 2022*, 363-379. https://doi.org/10.1007/978-3-031-20086-1_21
- Jaap Pedersen, Kai Hopmann-Baum, Janina Zittel, Thorsten Koch (2022). Blending hydrogen into natural gas: An assessment of the capacity of the German gas grid. In *Operations Research Proceedings 2021*, 182-187. https://doi.org/https://doi.org/10.1007/978-3-031-08623-6_28
- Ronaldo Rodrigues Pela, Claudia Draxl (2022). Ehrenfest dynamics implemented in the all-electron package exciting. *Electron. Struct.*, 4:037001. <https://doi.org/10.1088/2516-1075/ac7afc>
- *Felix Peppert, Max von Kleist, Christof Schütte, Vikram Sunkara (2022). On the Sufficient Condition for Solving the Gap-Filling Problem Using Deep Convolutional Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems*, 33(11):6194-6205. <https://doi.org/10.1109/TNNLS.2021.3072746> (Joint publication: Modeling and Simulation of Complex Processes, Visual, and Data-centric Computing)
- Milena Petkovic, Ying Chen, Inken Gamrath, Uwe Gotzes, Natalia Selini Hadjidimitrou, Janina Zittel, Xiaofei Xu, Thorsten Koch (2022). A hybrid approach for high precision prediction of gas flows. *Energy Systems*, 13:383-408. <https://doi.org/10.1007/s12667-021-00466-4>
- Milena Petkovic, Nazgul Zakiyeva, Janina Zittel (2022). Statistical Analysis and Modeling for Detecting Regime Changes in Gas Nomination Time Series. In *Operations Research Proceedings 2021*. OR 2021, *Lecture Notes in Operations Research*, pages 188-193. https://doi.org/10.1007/978-3-031-08623-6_29

- Milena Petkovic, Nazgul Zakiyeva (2022). Mathematical Optimization for Analyzing and Forecasting Nonlinear Network Time Series. *Operations Research Proceedings* 2022.
- Matthias Plock, Kas Andrie, Sven Burger, Philipp-Immanuel Schneider (2022). Bayesian Target-Vector Optimization for Efficient Parameter Reconstruction. *Adv. Theory Simul.*, 5:2200112. <https://doi.org/10.1002/adts.202200112>
- Julia Plöntzke, Mascha Berg, Rainald Ehrig, Sabine Leonhard-Marek, Kerstin-Elisabeth Müller, Susanna Röblitz (2022). Model Based Exploration of Hypokalemia in Dairy Cows. *Scientific Reports*, 12, 19781. <https://doi.org/10.1038/s41598-022-22596-0>
- Anup K. Prasad, Chandni Tiwari, Sourav Ray, Stephanie Holden, David A. Armstrong, K. Johan Rosengren, Alison Rodger, Ajay S. Panwar, Lisandra L. Martin (2022). Secondary Structure Transitions for a Family of Amyloidogenic, Antimicrobial Uperin 3 Peptides in Contact with Sodium Dodecyl Sulfate. *ChemPlusChem*, 87:e202100408. <https://doi.org/10.1002/cplu.202100408>
- Dimitar Prodanov, Sumit Kumar Vohra (2022). Active Segmentation: Differential Geometry meets Machine Learning. In *Proceedings of the 23rd International Conference on Computer Systems and Technologies*, 1-6. <https://doi.org/10.1145/3546118.3546154>
- N. Alexia Raharinirina, Esteban Acevedo-Trejos, Agostino Merico (2022). Modelling the acclimation capacity of coral reefs to a warming ocean. *Plos computational biology*. <https://doi.org/10.1371/journal.pcbi.1010099>
- Mona Rams, Tim Conrad (2022). Dictionary learning allows model-free pseudotime estimation of transcriptomics data. *BMC Genomics*, 23. <https://doi.org/10.1186/s12864-021-08276-9>
- Mauricio del Razo, Daniela Frömberg, Arthur Straube, Christof Schütte, Felix Höfling, Stefanie Winkelmann (2022). A probabilistic framework for particle-based reaction-diffusion dynamics using classical Fock space representations. *Letters in Mathematical Physics*, 112(49). <https://doi.org/10.1007/s11005-022-01539-w>
- Daniel Rehfeldt, Henriette Franz, Thorsten Koch (2022). Optimal Connected Subgraphs: Integer Programming Formulations and Polyhedra. *Networks*, 80(3):314-332. <https://doi.org/10.1002/net.22101>
- Daniel Rehfeldt, Hannes Hobbie, David Schönheit, Thorsten Koch, Dominik Möst, Ambros Gleixner (2022). A massively parallel interior-point solver for LPs with generalized arrowhead structure, and applications to energy system models. *European Journal of Operational Research*, 296(1):60-71. <https://doi.org/10.1016/j.ejor.2021.06.063>
- Daniel Rehfeldt (2022). Faster Algorithms for Steiner Tree and related Problems: From Theory to Practice. In *Operations Research Proceedings* 2022.
- Marco Reidelbach, Mei Bai, Martin Sebastian Zöllner, Michaela Schneeberger, Katharina Kubicek, Henning Kirchberg, Christian Bressler, Michael Thorwart, Carmen Herrmann (2022). Solvent dynamics of aqueous halides before and after photoionization. *Journal of Physical Chemistry*.
- Sourav Rej, Eva Yazmin Santiago, Olga Baturina, Yu Zhang, Sven Burger, Stepan Kment, Alexander O. Govorov, Alberto Naldoni (2022). Colloidal Titanium Nitride Nanobars for Broadband Inexpensive Plasmonics and Photochemistry from Visible to Mid-IR Wavelengths. *Nano Energy*, 104:107989. <https://doi.org/10.1016/j.nanoen.2022.107989>
- Lorenz Richter, Julius Berner (2022). Robust SDE-Based Variational Formulations for Solving Linear PDEs via Deep Learning. In *Proceedings of the 39th International Conference on Machine Learning*, PMLR, 162:18649-18666.
- Lucas Rickert, Timm Gao, Felix Urban, Jan Große, Nicole Srocka, Sven Rodt, Anna Musial, Kinga Zolnacz, Pawel Mergo, Kamil Dybka, Wacław Urbanczyk, Grzegorz Sek, Sven Burger, Stephan Reitzenstein, Tobias Heindel (2022). A Plug&Play Telecom-Wavelength Single-Photon Source for Quantum Key Distribution. In *Quantum 2.0, Technical Digest Series* (Optica), page QTu4C.6. <https://doi.org/10.1364/QUANTUM.2022.QTu4C.6>
- Lucas Siqueira Rodrigues, John Nyakatura, Stefan Zachow, Johann Habakuk Israel (2022). An Immersive Virtual Paleontology Application. In 13th International Conference on Human Haptic Sensing and Touch Enabled Computer Applications, *EuroHaptics* 2022, 478-481. <https://doi.org/10.1007/978-3-031-06249-0>
- Xavier Garcia Santiago, Martin Hammerschmidt, Johannes Sachs, Sven Burger, Hyunah Kwon, Marvin Knöller, Tilo Arens, Peer Fischer, Ivan Fernandez-Corbaton, Carsten Rockstuhl (2022). Toward maximally electromagnetically chiral scatterers at optical frequencies. *ACS Photonics*, 9:1954. <https://doi.org/10.1021/acsphotonics.1c01887>
- Thomas Schlechte, Ralf Borndörfer, Jonas Denißen, Simon Heller, Torsten Klug, Michael Küpper, Niels Lindner, Markus Reuther, Andreas Söhlke, William Steadman (2022). Timetable Optimization for a Moving Block System. *Journal of Rail Transport Planning & Management*, 22:100315. <https://doi.org/10.1016/j.jrtpm.2022.100315>
- Benedikt Schmitt, Jürgen Titschack, Daniel Baum (2022). Polyp-Cavity Segmentation of Cold-Water Corals guided by Ambient Occlusion and Ambient Curvature. *Eurographics Workshop on Visual Computing for Biology and Medicine* (VCBM). <https://doi.org/10.2312/vcbm.20221189>
- Philipp-Immanuel Schneider, Phillip Manley, Jan Krüger, Lin Zschiedrich, Rainer Köning, Bernd Bodermann, Sven Burger (2022). Reconstructing phase aberrations for high-precision dimensional microscopy. *Proc. SPIE*, 12137:121370I. <https://doi.org/10.1117/12.2620762>
- Philipp-Immanuel Schneider, Lin Zschiedrich, Martin Hammerschmidt, Lilli Kuen, Ivan Sekulic, Julien Kluge, Bastian Leykauf, Markus Krutzik, Sven Burger (2022). Boosting quantum optics experiments with Bayesian optimization. In *Proc. SPIE*, 12227:122270F. <https://doi.org/10.1117/12.2632419>
- Stephan Schwartz (2022). An overview of graph covering and partitioning. *Discrete Mathematics*, 345(8). <https://doi.org/https://doi.org/10.1016/j.disc.2022.112884>
- Renata Sechi, Konstantin Fackeldey, Surahit Chewle, Marcus Weber (2022). SepFree NMF: A Toolbox for Analyzing the Kinetics of Sequential Spectroscopic Data. *Algorithms*, 15(9):297. <https://doi.org/10.3390/a15090297>

- Isam Ben Soltane, Rémi Colom, Brian Stout, Nicolas Bonod (2022). Derivation of the Transient and Steady Optical States from the Poles of the S-Matrix. *Laser Photon. Rev.* <https://doi.org/10.1002/lpor.202200141>
- Vikram Sunkara, Angus Lewis, Giang T. Nguyen, Malgorzata M. O'Reilly, Nigel Bean (2022). A discontinuous Galerkin method for approximating the stationary distribution of stochastic fluid-fluid processes. Methodology and Computing in *Applied Probability*. <https://doi.org/10.1007/s11009-022-09945-2>
- Peter Tillmann, Klaus Jäger, Asher Karsenti, Lev Kreinin, Christiane Becker (2022). Model-Chain Validation for Estimating the Energy Yield of Bifacial Perovskite/Silicon Tandem Solar Cells. *Sol. RRL*, 6:2200079. <https://doi.org/10.1002/solr.202200079>
- Philipp Tockhorn, Johannes Sutter, Alexandros Cruz, Philipp Wagner, Klaus Jäger, Danbi Yoo, Felix Lang, Max Grischek, Bor Li, Jinzhao Li, Oleksandra Shargaieva, Eva Unger, Amran Al-Ashouri, Eike Köhnen, Martin Stollerfoht, Dieter Neher, Rutger Schlatmann, Bernd Rech, Bernd Stannowski, Steve Albrecht, Christiane Becker (2022). Nano-optical designs for high-efficiency monolithic perovskite-silicon tandem solar cells. *Nat. Nanotechnol.*, 17:1214. <https://doi.org/10.1038/s41565-022-01228-8>
- Philipp Tockhorn, Johannes Sutter, Alexandros Cruz, Philipp Wagner, Klaus Jäger, Danbi Yoo, Felix Lang, Max Grischek, Bor Li, Amran Al-Ashouri, Eike Köhnen, Martin Stollerfoht, Dieter Neher, Rutger Schlatmann, Bernd Rech, Bernd Stannowski, Steve Albrecht, Christiane Becker (2022). Nanooptically Enhanced Perovskite/Silicon Tandem Solar Cells with 29.80% Power Conversion Efficiency. In *OSA Adv. Photonics, OSA Technical Digest*, page JT4H.2. <https://doi.org/10.1364/NOMA.2022.JT4H.2>
- Kazuma Tsuji, Ken'ichiro Tanaka, Sebastian Pokutta (2022). Pairwise Conditional Gradients without Swap Steps and Sparser Kernel Herding. In *Proceedings of the International Conference on Machine Learning*.
- Erhun Giray Tuncay, Riza Cenk Erdur, Tim Conrad (2022). Parallel Exchange of Randomized SubGraphs for Optimization of Network Alignment: PERSONA. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*. <https://doi.org/10.1109/TCBB.2022.3231489>
- Daniel Udvardy, Philipp Harth, Jakob H. Macke, Hans-Christian Hege, Christiaan P. J. de Kock, Bert Sakmann, Marcel Oberlaender (2022). The Impact of Neuron Morphology on Cortical Network Architecture. *Cell Reports*, 39(2). <https://doi.org/10.1016/j.celrep.2022.110677>
- Kseniia Volkova, Julia Heupel, Sergei Trofimov, Fridtjof Betz, Rémi Colom, Rowan W. MacQueen, Sapida Akhundzada, Meike Reginka, Arno Ehresmann, Johann Peter Reithmeier, Sven Burger, Cyril Popov, Boris Naydenov (2022). Optical and Spin Properties of NV Center Ensembles in Diamond Nano-Pillars. *Nanomaterials*, 12:1516. <https://doi.org/10.3390/nano12091516>
- Yujing Wang, Luca Vannucci, Sven Burger, Niels Gregersen (2022). Near-unity efficiency in ridge waveguide-based, on-chip single-photon sources. *Mater. Quantum Technol.*, 2:045004. <https://doi.org/10.1088/2633-4356/aca8e8>
- Martin Weiser, Fatemeh Chegini (2022). Adaptive multirate integration of cardiac electrophysiology with spectral deferred correction methods. In *CMBE22 - 7th International Conference on Computational & Mathematical Biomedical Engineering*, 528-531.
- Elias Wirth, Sebastian Pokutta (2022). Conditional Gradients for the Approximately Vanishing Ideal. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*.
- Stephan Wäldchen, Felix Huber, Sebastian Pokutta (2022). Training Characteristic Functions with Reinforcement Learning: XAI-methods Play Connect Four. In *Proceedings of the International Conference on Machine Learning*.
- Lina Maria Zangerl, Wolfgang Peters-Kottig, Achim Oßwald (2022). Personal Digital Archiving: Eine neue Aufgabe für Öffentliche und Wissenschaftliche Bibliotheken. *Bibliothek Forschung und Praxis*, 46(1):152-161. <https://doi.org/10.1515/bfp-2021-0058>
- Wei Zhang, Tiejun Li, Christof Schütte (2022). Solving eigenvalue PDEs of metastable diffusion processes using artificial neural networks. *Journal of Computational Physics*, 465. <https://doi.org/10.1016/j.jcp.2022.111377>
- Boro Šofranac, Ambros Gleixner, Sebastian Pokutta (2022). Accelerating domain propagation: An efficient GPU-parallel algorithm over sparse matrices. *Parallel Computing*, 109:102874. <https://doi.org/10.1016/j.parco.2021.102874>
- Boro Šofranac, Ambros Gleixner, Sebastian Pokutta (2022). An Algorithm-independent Measure of Progress for Linear Constraint Propagation. *Constraints*, 27:432-455. <https://doi.org/10.1007/s10601-022-09338-9>
- Torsten Klug, Markus Reuther, Thomas Schlechte (2022). Does Laziness Pay Off? – A Lazy-Constraint Approach to Timetabling. In: *22nd Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2022)*
- Berenike Masing (2022). Optimal Line Planning in the Parametric City. In: *Operations Research Proceedings 2021*
- Berenike Masing, Niels Lindner, Ralf Borndörfer (2022). Optimal Line Plans in the Parametric City and the Impact of In-Motion Costs. In: *Operations Research Proceedings 2021*

Books (2)

Mattia D'Emidio, Niels Lindner, editors (2022). 22nd Symposium on *Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2022)*, Volume 106

Katsuki Fujisawa, Shizuo Kaji, Toru Ishihara, Masaaki Kondo, Yuji Shinano, Takuji Tanigawa, Naoko Nakayama, editors (2022). *Construction of Mathematical Basis for Realizing Data Rating Service*. Institute of Mathematics for Industry, Lecture Notes, Kyushu University.

Dissertations (11)

Felix Ambellan (2022). Efficient Riemannian Statistical Shape Analysis with Applications in Disease Assessment. Freie Universität Berlin.

Fatemeh Chegini (2022). Multilevel optimization algorithm for Inverse Problem in Electrocardiography. Università della Svizzera italiana.

Luzie Helfmann (2022). Non-stationary Transition Path Theory with applications to tipping and agent-based models. Freie Universität Berlin.

Kai Hoppmann-Baum (2022). Mathematical programming for stable control and safe operation of gas transport networks. Technische Universität Berlin.

Sahar Iravani (2022). Interpretable Deep Learning Approaches for Biomarker Detection from High-Dimensional Biomedical Data. Freie Universität Berlin.

Jan-Hendrik Niemann (2022). Learning Reduced Models for Large-Scale Agent-Based Systems. Freie Universität Berlin.

Mona Milena Rams (2022). New approaches for unsupervised transcriptomic data analysis based on Dictionary learning. Freie Universität Berlin.

Manish Sahu (2022). Vision-based Context-awareness in Minimally Invasive Surgical Video Streams. Technische Universität Berlin.

Tom Streubel (2022). Simulation of Piecewise Smooth Differential Algebraic Equations with Application to Gas Networks. Humboldt-Universität zu Berlin.

Mattes Mollenhauer (2022). On the statistical approximation of conditional expectation operators. Freie Universität Berlin

Niklas Wulkow (2022). Modelling Observations of Dynamical Systems with Memory. Freie Universität Berlin

IMPRINT

Published by

Zuse-Institut Berlin
Takustrasse 7 | 14195 Berlin | Germany

Phone +49 30 84185-0
Fax +49 30 84185-125
Email info@zib.de
www.zib.de

Layout

Christiane Weimann Grafik + Produktion
www.schnittmarke.biz

Printing

Lochmann grafische Produktion GmbH

ISSN

0934-5892

©2022 Zuse-Institut Berlin

NOTES

