# ANNUAL REPORT 2019

ZUSE INSTITUTE BERLIN

# CONTENTS


ZUSE INSTIT


TUTE BERLIN

# EXECUTIVE SUMMARY



The year 2019 was a very intense year for ZIB, in which a series of long-planned and relentlessly pursued goals were achieved and a plethora of new things were started. Selected successes in random order:

1. We were able to welcome a new Vice President to ZIB and establish through him and around him a new department AI in Science, Society, and Technology ($AIS^2T$).
2. We created the new department in relation to a new Federal Competence Center for Big Data and Machine Learning in Berlin.
3. We started the activity of the new Cluster of Excellence MATH+ at ZIB.
4. We inaugurated our newest supercomputer.
5. We were part of the Berlin University Alliance's successful proposal in the context of the German excellence strategy.
6. We could celebrate a significant increase in our core budget through the state of Berlin.

In the following, we will elaborate on some of the above to outline their respective importance for ZIB.

**New Vice President.** In September 2019, Professor Sebastian Pokutta started as ZIB's new Vice President (VP). This ended a period of almost four years in which ZIB had to operate with a vacant VP position. We are therefore particularly happy that we were able to successfully fill this important position. Prior to joining ZIB and TU Berlin, Professor Pokutta, a mathematician by training, was the David M. McKenney Family associate professor at the School of Industrial and Systems Engineering and an associate director of the Machine Learning @ GT Center at the Georgia Institute of Technology. Professor Pokutta received the David M. McKenney Family Early Career Professorship in 2016, an NSF CAREER Award in 2015, the Coca-Cola Early Career Professorship in 2014, the outstanding thesis award of the University of Duisburg-Essen in 2006, as well as various best paper awards. His research is situated at the intersection of artificial intelligence and optimization.

**BIFOLD.** Machine Learning (ML) (together with decision-making) and Big Data (BD) are the key pillars of Artificial Intelligence (AI). Research in ML and BD has been an integral part of research activities at ZIB for ten years and is becoming a cornerstone of our research strategy. To galvanize our activities in this area, we decided to create a new research department "AI in Society, Science, and Technology ($AIS^2T$)" at ZIB in fall 2019. The department is strongly integrated with already-existing departments. Additionally, a new Berlin-based Federal Competence Center, the "Berlin Institute for the Foundations of Learning and Data" (BIFOLD), was founded that integrates the former Berlin Center for Machine Learning (BZML) and the Berlin Big Data Center (BBDC) into a new structure that is funded by the German Federal Ministry of Education and Research (BMBF). ZIB, having been part of both BZML and BBDC, is part of the consortium supporting BIFOLD and will establish new research groups in this context. We will further discuss our current research activities as well as our strategy to strengthen and intensify our presence in this area in the feature article "Machine Learning and Big Data."

**MATH+.** The Berlin-based Cluster of Excellence MATH+ succeeded in the Excellence Strategy research competition run by the German federal and state governments. MATH+ started in January 2019 as a joint project between the universities FU Berlin, HU Berlin, and TU Berlin as well as the two research centers WIAS and ZIB. The cluster is a cross-institutional and transdisciplinary Cluster of Excellence where researchers explore and further develop new approaches in application-oriented mathematics. A particular focus is merging mathematical modeling with mathematical techniques for analyzing ever-growing amounts of data in life and material sciences in energy and network research, and in the humanities and social sciences. The aim is to boost not only scientific progress, but also technological innovation and a comprehensive understanding of social processes. ZIB's research strengths in application-oriented mathematics and data-driven research in a transdisciplinary context including translation to industry and society at-large is recognized as one of the cornerstones of MATH+. At the institute, MATH+ started with ten new research projects funded through the Cluster of Excellence. These projects include about 40 researchers at ZIB alone and we provide an inside perspective into the research conducted in two of those projects in the feature articles "Riemannian Analysis of Time-Varying Shape Data – Understanding Geometric Evolution" and "Green Energy."

**HLRN.** On December 6, 2019, the newest supercomputer at ZIB was officially inaugurated. This fourth generation of supercomputer systems of the Northern German Supercomputing Alliance (HLRN) holds rank 40 in the latest TOP500 list of the fastest supercomputers in the world. It provides about six times more compute power and twice the online storage space than its predecessor, while consuming only 50% more electric power and occupying the same floor space. With its aalmost quarter of a million compute cores, this massively parallel system architecture provides a new level of parallelism to support computationally demanding scientific workloads from a broad spectrum of use cases. More details can be found in the article "The 40th-Fastest Supercomputer in the World."

**Berlin University Alliance.** In July 2019, the Berlin University Alliance (BUA) of FU Berlin, HU Berlin, and TU Berlin, together with their university hospital Charité, were admitted into the Universities of Excellence funding line of the German government's Excellence Strategy. The Berlin University Alliance's long-term goal is to turn Berlin into an integrated research environment. ZIB is one of the BUA's partner institutes for application-oriented mathematics and research infrastructure-related community building. The aim is to provide better access to and joint usage of top-level research infrastructure and community-centric services in order to foster interdisciplinary research communities around cutting-edge research infrastructure such as the HPC and large data facilities at ZIB.

**Increase of budget.** In mid-December 2019, as a much-celebrated Christmas gift, Berlin's parliament approved an increase of almost 15% in the basic budget of ZIB for the years 2020 and 2021 in order to support ZIB's current research activities, allowing for additional investments in restructuring and strengthening research, and improving research infrastructure. This highly necessary increase of our core budget and the complementing record high level of third-party funds bring ZIB into an economically strong position to tackle major research and innovation challenges in the future.

**More insights.** In addition to the topics already mentioned, this annual report provides insights into a variety of other success stories and gives a general overview of ZIB's organization and key factors for its successful development. For example, the feature article "On the Road to Autonomous Operating Room Scheduling" provides insight into how algorithmic intelligence and optimal decision-making can be utilized in the operating room through a joint project between ZIB and several partners including Berlin's university hospital Charité. This overview is complemented by the feature article "Solving Real-World Optimization Problems," that reports on ZIB's renowned in-house optimization suite SCIP and its success stories, outlining the possibilities of one of the fastest academic mathematical optimization software packages worldwide. In another feature article, "Reconfigurable Computing Today," it is explained why there is a new wave of interest in Field-Programmable Gate Arrays (FPGA). FPGAs can be used for implementing highly data-parallel architectures to target more power-efficient solutions as well as accelerators for HPC.

Last year was a fantastic year for ZIB with key successes that will positively impact on ZIB's development for a decade and beyond. Many other very positive developments make us confident that our institute has a bright future. Put differently, ZIB continues to be a place for excellent research and first-rate scientific services and infrastructure.

# Zuse Institute Berlin –
The Movie

www.zib.de/movie

Click on the image to start the movie.

# ORGANIZATION

## The Statutes

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



### Administrative Bodies

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

President of ZIB
**Prof. Dr. Christof Schütte**

Vice President
until February 28, 2019
**Prof. Dr. Martin Skutella**
since September 1, 2019
**Prof. Dr. Sebastian Pokutta**

The Board of Directors was composed in 2019 as follows:

**Prof. Dr. Peter Frensch**
Vice President, Humboldt-Universität zu Berlin (Chairman)

**Prof. Dr. Christian Thomsen**
President, Technische Universität Berlin (Vice Chairman)

**Dr.-Ing. Andrea Bör**
Provost, Freie Universität Berlin

**Prof. Dr. Günther Ziegler**
President, Freie Universität Berlin

**Frau Ellen Fröhlich**
Der Regierende Bürgermeister von Berlin, Senatskanzlei – Wissenschaft und Forschung

**Dr. Jürgen Varnhorn**
Senatsverwaltung für Wirtschaft, Energie und Betriebe

**Prof. Dr. Manfred Hennecke**
Bundesanstalt für Materialforschung und -prüfung (BAM)

**Thomas Frederking**
Helmholtz-Zentrum Berlin für Materialien und Energie (HZB)

**Prof. Dr. Heike Graßmann**
Max-Delbrück-Centrum für Molekulare Medizin (MDC)

The Board of Directors met on June 7, 2019, and November 27, 2019.

# Scientific Advisory Board

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

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

**Prof. Dr. Jörg-Rüdiger Sack**
Carleton University, Ottawa, Canada

**Prof. Dr. Cecilia Clementi**
Rice University, Houston, Texas, USA

**Prof. Dr. Michael Dellnitz**
Universität Paderborn, Germany

**Prof. Dr. Rolf Krause**
Université della svizzera italiana, Lugano, Switzerland

**Ludger D. Sax**
Grid Optimization Europe GmbH

**Prof. Dr. Reinhard Schneider**
Université du Luxembourg, Luxembourg

**Prof. Dr. Dorothea Wagner**
Karlsruher Institut für Technologie (KIT), Karlsruhe, Germany

The Scientific Advisory Board met on July 8 and 9, 2019, at ZIB.

**SCIENTIFIC ADVISORY BOARD**

**BOARD OF DIRECTORS**
Chairman: Prof. Dr. Peter Frensch
Humboldt-Universität zu Berlin (HUB)

**PRESIDENT**
Prof. Dr. Christof Schütte

**VICE PRESIDENT**
Prof. Dr. Martin Skutella until February 28, 2019 | Prof. Dr. Sebastian Pokutta since September 1, 2019

**MATHEMATICS FOR LIFE AND MATERIALS SCIENCES**
Prof. Dr. Christof Schütte

**PARALLEL AND DISTRIBUTED COMPUTING**
Prof. Dr. Alexander Reinefeld

**MATHEMATICAL OPTIMIZATION AND SCIENTIFIC INFORMATION**
Prof. Dr. Martin Skutella
until February 28, 2019

Prof. Dr. Christof Schütte
intermediate acting

Prof. Dr. Sebastian Pokutta
since September 1, 2019

**ADMINISTRATION**
Annerose Steinke

Dr. Kathrin Rost-Drese
(acting) since April 1, 2019

# ZIB STRUCTURE

- **MATHEMATICS FOR LIFE AND MATERIALS SCIENCES** – C. Schütte
  - **NUMERICAL MATHEMATICS** – M. Weiser
    - COMPUTATIONAL MEDICINE – M. Weiser, S. Zachow
    - COMPUTATIONAL MOLECULAR DESIGN – M. Weber
    - COMPUTATIONAL NANO OPTICS – S. Burger
    - COMPUTATIONAL SYSTEMS BIOLOGY – S. Winkelmann
    - UNCERTAINTY QUANTIFICATION – T. Sullivan
    - COMPUTATIONAL HUMANITIES – N. Conrad
  - **VISUAL DATA ANALYSIS** – H. Hege
    - VISUAL DATA ANALYSIS IN SCIENCE AND ENGINEERING – H. Hege
    - IMAGE ANALYSIS IN BIOLOGY AND MATERIALS SCIENCE – S. Prohaska, D. Baum
    - THERAPY PLANNING – S. Zachow
    - BIOINFORMATICS IN MEDICINE – T. Conrad

- **MATHEMATICAL OPTIMIZATION AND SCIENTIFIC INFORMATION** – S. Pokutta
  - **MATHEMATICAL OPTIMIZATION** – R. Borndörfer, T. Koch
    - MATHEMATICS OF TRANSPORTATION AND LOGISTICS – R. Borndörfer
    - MATHEMATICAL OPTIMIZATION METHODS – A. Gleixner
    - ENERGY NETWORK OPTIMIZATION – J. Zittel
    - MATHEMATICS OF HEALTH CARE – G. Sagnol
  - **SCIENTIFIC INFORMATION** – T. Koch
    - WEB TECHNOLOGY AND MULTIMEDIA – W. Dalitz
    - DIGITAL PRESERVATION – W. Peters-Kottig
    - RESEARCH AND COMPETENCE CENTER DIGITALIZATION BERLIN (DIGIS) – A. Müller
    - KOBV LIBRARY NETWORK – RESEARCH AND DEVELOPMENT – B. Rusch
    - KOBV LIBRARY NETWORK – OPERATING – S. Lohrum
    - FRIEDRICH-ALTHOFF-KONSORTIUM – U. Kaminsky
    - ZIB LIBRARY – K. Lachmann
  - **AI IN SOCIETY, SCIENCE, AND TECHNOLOGY (AIS²T)** – S. Pokutta

- **CORE FACILITY IT AND DATA SERVICES** – C. Schäuble
  - BRAIN BERLIN RESEARCH AREA INFORMATION NETWORK – C. Schäuble
  - PROCESS MODELING AND MANAGEMENT – C. Schäuble, K. Rost-Drese

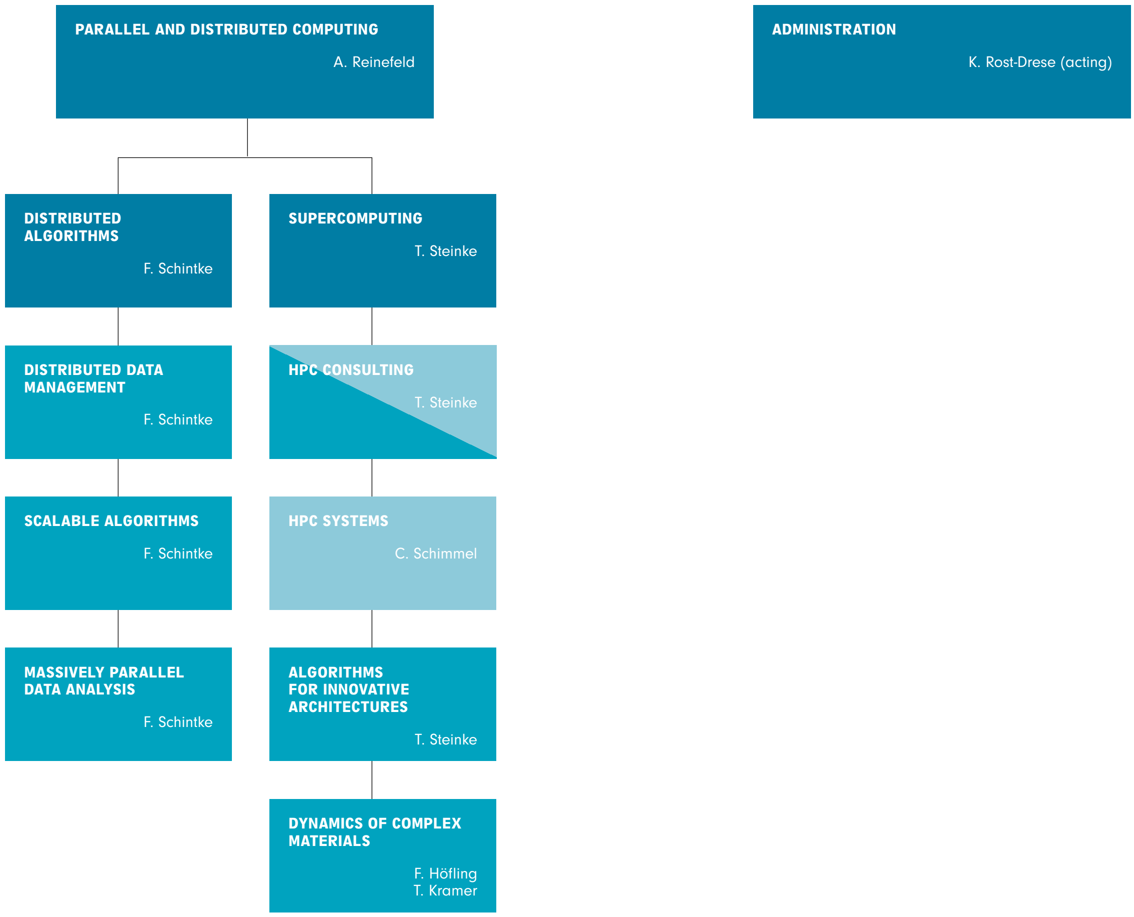


**ZIB is structured into four divisions: three scientific divisions and ZIB's administration.**

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

## LEGEND

- Scientific divisions and departments
- Research groups
- Research Service groups
- Core Facility

# ZIB MEMBERS

Michael Winkler, Helena Müller, Tamaz Amiranashvili,
Isabel Beckenbach, Tobias Achterberg, Ksenia Bestuzheva,
Marcel Drachmann, Karl-Heinz Haag, Alejandro Carderera,
Lovis Anderson, Sylke Arencibia, Ralf Borndörfer, Tim Sullivan,
Bahareh Banyassady, Josephine Brummer, Matteo Francobaldi,
Felix Ambellan, Ingo Meise, Ferdinand Bleschke, Daniel Rehfeldt,
Ulrike Homberg, Heike Balluneit, Luzie Helfmann, Matthias Noack,
Benjamin Hiller, Kathrin Rost-Drese, Pedro Maristany de las Casas,
Masoud Gholami Estahbanati,
Philipp-Immanuel Schneider,
The Anh Pham, Felix Thoma,
Tom Streubel, Oasama Bach,
Tim Conrad, Wolfgang Dalitz,
Felix Kalg, Renata Kussack,
Sven Helemann, Sourav Ray,
Jana Spiller, Marcus Weber,
Huy Le Duc, Felix Baumann,
Tobias Baumann, Leon Weis,
Mascha Berg, Beate Rusch,
Kai-Helge Becker, Jan Lutz,
Patrick Werk, Arthur Straube,
Gioni Mexi, Felix Binkowski,
Gregor Hendel, Max Zimmer,
Raphael Badel, Anja Müller,
Ridvan Uzun, Timo Berthold,
Sven Burger, Ralph Böhmert,
Jonas Patzelt, Alexander Tack,
Ursula Droebes, Wei Zhang,
Fridtjof Betz, Marco Blanco,
Julia Boltze, Julia Eichhorn,
Rémi Colom, Elmar Swarat,
Gábor Braun, Ricardo Euler,
Boris Grimm, Thomas Dierkes,
Ariane Ernst, Thomas Steinke, Robert Brandt, Felix Hennings, Kilian Amrhein,
Felix Herter, Katharina Lachmann, Nando Farchmin, Gerald Gamrath,
Marco Lübbecke, Alexander Sikorski, Fatemeh Chegini, Nils-Christian Kempke,
Natasa Conrad, Andreas Löbel, Franziska Egbers, Margarita Kostre,
Lu-Xi Feng, Konstantin Fackeldey, Andrea Kratz, Franziska Erlekam,
Armin Fügenschuh, Alexey Sharenkov, Anna-Lena Nowicki, Tobias Kramer,
Fabian Danecker, Birzhan Ayanbayev, Luitgard Kraus, Hagen Chrapary,
Steffen Christgau, Robert Clausecker, Cyrille Combettes, Leona Gottwald

Jan Szelag, Olaf Paetsch,
Martin Hanik, Stefan Düring,
Marc Hartung, Tim Hasler,
Silke Haase, Ramona Edler,
Stefan Heinz, Daniel Baum,
Fabian Löbel, Felix Höfling,
Hanna Jansen, Julian Bushe,
Dennis Jentsch, Klaus Jäger,
Benjamin Kaiser, Leon Eifler,
Uta Kaminsky, Oliver Kant,
Rainer Klaus, Ilja Klebanov,
Marco Klindt, Marius Knaust,
David Knötel, Thorsten Koch,
Gabriel Kressin, Dirk Krickel,
Karlotta Kruschke, Eric Kunze,
Vivian Köneke, Jonas Köppl,
Doreen Kühne, Arno Kühner,
Morgan Leborgne, Ralf Lenz,
Niels Lindner, Stefan Lohrum,
Christoph Kahl, Moritz Ehlke,
Natalia Ernst, Katrin Fedtke,
Guvenc Sahin, Guvenc Sahin,
Boro Šofranac, Signe Weihe,
Stefan Vigerske, Kati Wolter,
Yuji Shinano, Jesco Humpola,
Ursula Stanek, Norma Schüler,
Robin Chemnitz, Rick Grap,
Pooja Gupta, Felipe Serrano,
Phillip Semler, René Skillen,
Marek Fröhlich, Yildiz Şahin,
Denise Scholz, Renata Sechi,
Paul Seydel, Janina Zittel,
Felix Prause, Sascha Witte,
William Surau, Julia Plöntzke,
Tobias Weiß, Florian Wende,
Ying Wang, Sofi Tiwari,
Jens Schwidder, Stefan Wollny,
Denise Scholz, Oliver Meyer,
Katrin Woelk, Florian Willich,
Nils Gauglitz, Manish Sahu,
Felix Peppert, Niklas Wulkow,
Justus Vogel, Frank Schmidt,
Jakob Witzig, Stefan Zachow,

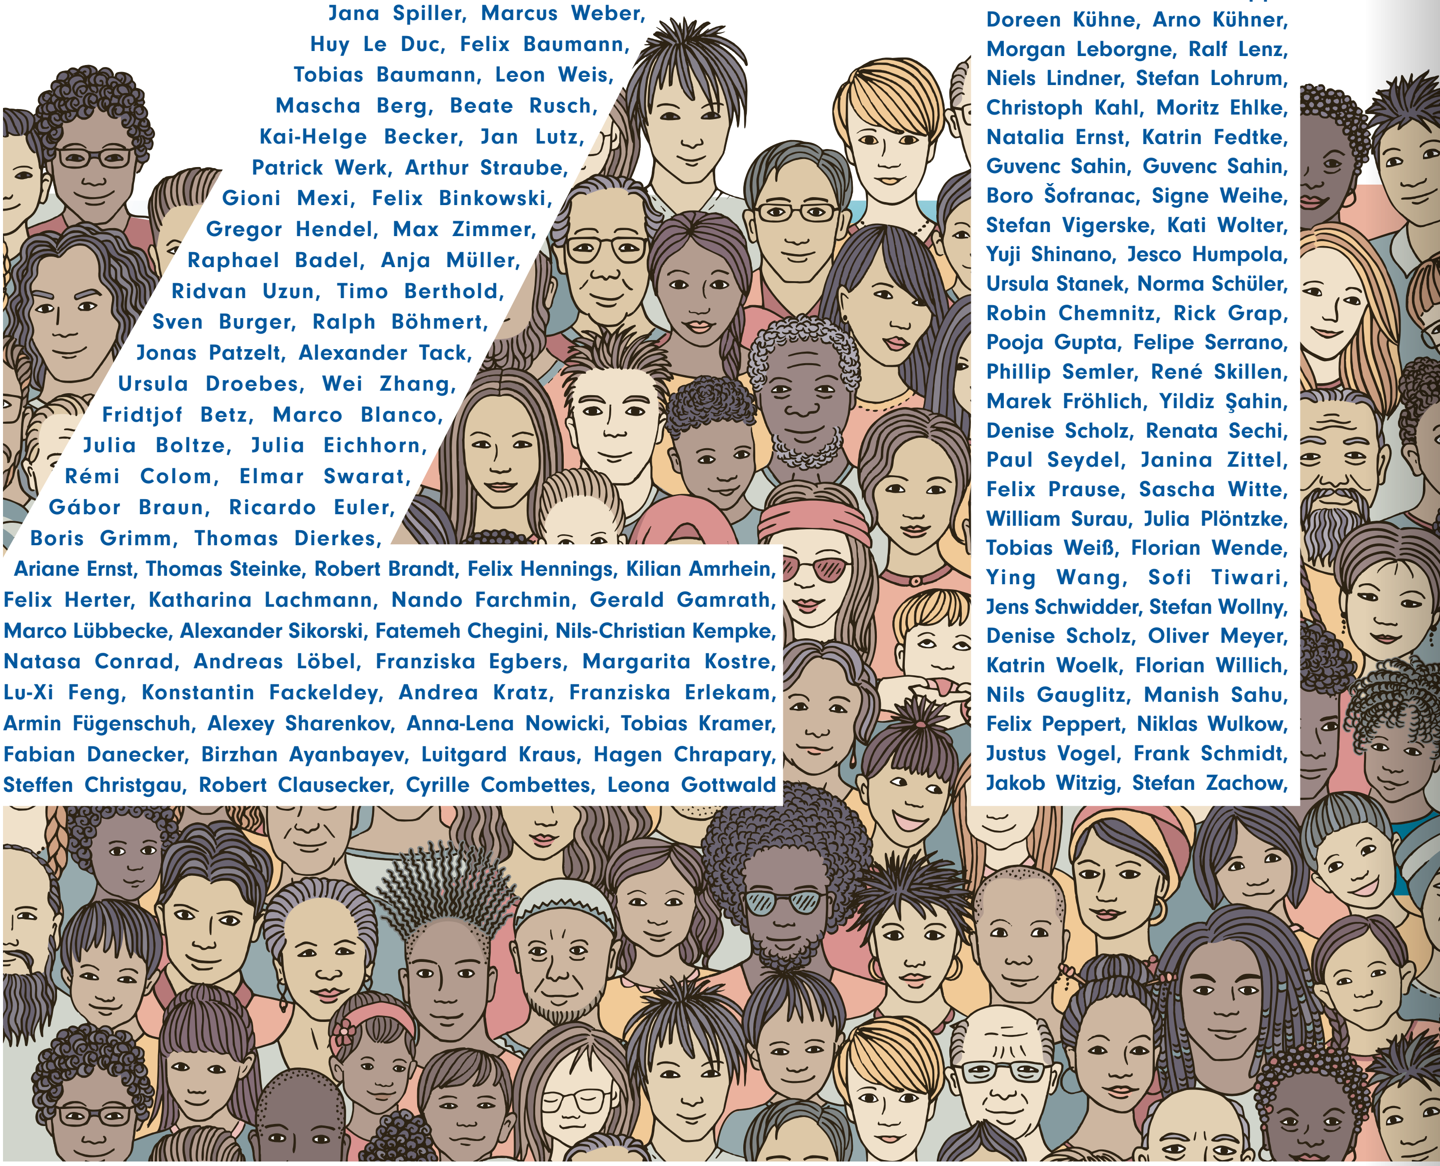

Hans-Christian Hege, Petra Fehlhauer, Han Lie, Florian Schintke, Neveen Eshtewy, Sascha Szott, Maria Mathew, Pierre Le Bodic, Martin Grötschel, Julien Roussel, Carl Martin Grewe, Kai Hoppmann, Jan Luca Naumann, Philipp Gutsche, Carsten Dreßke, Rainald Ehrig, Berenike Masing, Sahar Iravani, Christian Tuma, Matthias Läuter, Christine Tawfik, Sybille Mattrisch, Stephen Maher, Jan Skrzypczak, Hans Lamecker, Robert Julian Rabben, Jannis Polojannis, Susanna Röblitz, Lucas Siqueira Rodrigues, Alexander Reinefeld, Sandra Patzelt-Schütte, Jan-Patrick Clarner, Jorit Kleine-Möllhoff, Matthias Wittenberg, Ronja Strömsdörfer, Tobias Watermann, Wolfgang Peters-Kottig, Jan-Hendrik Niemann, Esfandiar Navayazdani, Hans-Hermann Frese, Xavier Garcia Santiago, Guillaume Sagnol, Hanh Dung Nguyen, N. Alexia Raharinirina, Stefanie Winkelmann, Lin Werner Zschiedrich, Felix Wiederschein, Marlies Engelke, Marwan Muhammad, Roxanne Vierhaus, Michael Wulkow, Marc Osterland, Anton Pakhomov, Ingmar Schuster, Viktoria Gerlach, Milena Petkovic, Matthias Plock, Uwe Uhlmann, Steffen Prohaska, Uwe Neumann, Dev Punjabi, Bernhard Reuter, Hanna Wulkow, Enric Ribera Borrell, Phillip Manley, Theresa Höhne, Elias Oltmanns, Sumit Kumar Vohra, Adam Schienle, Kateryna Melnyk, Mirta Rodríguez, Christof Schulz, Stanley Schade, Benjamin Müller, Philipp Heinrich, Julia Alexandra Goltz-Fellgiebel, Darshika Sharma, Enrico Bortoletto, Francesco D'Amato, Nicole Heidingsfelder, Lars Zimmermann, Christian Salzmann-Jäckel, Fabian Wegscheider, Heinz-Günter Kuper, Hildegard Franck, Steffi Conrad-Rempel, Martin Hammerschmidt, Matthias Miltenberger, Mattes Mollenhauer, Marian Moldenhauer, Sebastian Götschel, Thomas Schlechte, Erlinda C. Körnig, Franziska Schlösser, Christian Schimmel, Merlind Schotte, Stephan Schwartz, Horst-Holger Boltz, Fleur Schweigart, Aleksandr Gonopolskiy, Thomas Wornien, Jan Giesebrecht, Mohamed Omari, Farouk Salem, Carsten Schäuble, Jonas Schweiger, Ramin Khorsandi, Alexey Shestakov, Wolfram Sperber, Andreas Staude, Vikram Sunkara, Alexander Tesch, Stefanie Kuwatsch, Peter Tillmann, Mark Ruben Turner, Norbert Lindow, José Villatoro, Paul Wancura, Antonia Chmiela, Martin Weigelt, Johannes Zonker, Inci Yüksel-Ergün, Kuba Weimann, Martin Weiser, Christof Schütte, Sebastian Pokutta

In fall 2019, Professor Sebastian Pokutta from TU Berlin was appointed Vice President of the Zuse Institute Berlin. Having received both his diploma and Ph.D. in mathematics from the University of Duisburg-Essen in Germany, Professor Pokutta was a postdoctoral researcher and visiting lecturer at MIT, worked for IBM ILOG, and Krall Demmel Baumgarten. Prior to joining ZIB and TU Berlin, he was the David M. McKenney Family associate professor in the School of Industrial and Systems Engineering and an associate director of the Machine Learning @ GT Center at the Georgia Institute of Technology as well as a professor at the University of Erlangen-Nürnberg. Professor Pokutta received the David M. McKenney Family Early Career Professorship in 2016, an NSF CAREER Award in 2015, the Coca-Cola Early Career Professorship in 2014, the outstanding thesis award of the University of Duisburg-Essen in 2006, as well as various best paper awards. His research is situated at the intersection of artificial intelligence and optimization. A particular focus is on combining machine learning with optimization techniques, both discrete and continuous.

With the appointment of the new Vice President, the Zuse Institute Berlin is further expanding its foray into artificial intelligence, machine learning, and optimization. Apart from the integration with Berlin's AI and ML efforts, such as BIFOLD, a particular focus is on international collaborations and partnerships. To date, research collaborations and memorandums of understanding (MoU) have been signed with RIKEN-AIP in Tokyo, which is comparable to a Max Planck Institute, the McMaster University in Toronto, as well as the Fraunhofer IIS in Erlangen.

# A NEW AI DEPARTMENT

Tightly integrated with the activities of the VP is the newly founded department with the name AI in Society, Science, and Technology (AIS$^2$T). "AI is not just a methodology but a paradigm shift that permeates all disciplines. ZIB is embracing the opportunities and possibilities of this technology. At the same time, we do understand that this technology is touching human lives and society on so many levels that an inquiry disregarding societal aspects would be irresponsible," says Professor Pokutta. The department will focus on the development of new AI methodologies as well the interface and translation of these methods into an industry context. Initially, the department will be concerned with the following three major directions:

**1. Integrating Learning and Decision-Making.** We are at a point in time where we understand how to turn data into insights via machine learning very well. We also very much know how to compute optimal decisions using prior insights obtained, e.g., from data. However, the holistic integration of learning and decisions is still a key open problem and one major challenge that we aim to tackle.

**2. Specialized Hardware.** Today's SoCs and FPGAs are powerful customizable hardware that allows complex algorithms to be accelerated and implemented for specialized loads. While specialized/customized hardware is quite common in the context of deep learning (e.g. by means of GPUs or Google's TPUs, which one might think of as custom ASICs), no such custom implementations exist for discrete methods (such as mixed-integer programming algorithms or constraint programming algorithms) or heuristics commonly used in discrete methods. We are interested in designing custom hardware accelerators (via FPGAs) to significantly speed things up, i.e. discrete optimization algorithms by implementing specialized operations/functions. We are also interested in the edge computing regime, where we want to deploy complex algorithms in situ. An example in this context is, for example, the deployment of SCIP on a Raspberry Pi.

**3. Artificial Intelligence in Society.** A particular focus is on the interaction of AI methods with societal challenges as it has become clear that questions of fairness, accountability, transparency, and explainability (FATE) have to be not only addressed through policies but also directly within the algorithmic design.


# LIS



The new HLRN-IV supercomputer at ZIB is a massively parallel system for a broad spectrum of applications.

**On December 6th, 2019,** the next-generation supercomputer at ZIB was officially inaugurated. This fourth generation of supercomputer systems of the North German Supercomputing Alliance (HLRN) holds the 40th rank in the latest TOP500 list of the fastest supercomputers of the world.

With its name “Lise”, we honor the theoretical work of Lise Meitner to explain the experiments of Otto Hahn on nuclear fission of uranium in the late 1930s. Lise Meitner spent most of her scientific career in Berlin, where she was a physics professor and a department head at the Kaiser Wilhelm Institute (today Hahn-Meitner Institute), just a few hundred meters away from ZIB. Notable, she was the first woman to become a full professor of physics in Germany.

The HLRN-IV supercomputer “Lise” provides about six times more compute capacity and twice of the online storage space than its predecessor, the HLRN-III system “Konrad”, while consuming only 50% more electric power and occupying the same floor space. With its almost a quarter of million compute cores, this massively parallel system architecture provides a new level of parallelism to support computationally demanding scientific workloads out of a broad spectrum of usage scenarios.

## The HLRN-IV System “Lise” at a Glance

**1,270** nodes

in total, each with two Intel Xeon Cascade Lake Advanced Processors with 48 cores, i.e. 96 Intel Xeon cores per node:

- 1,236 standard nodes with 384 GByte main memory,
- 32 large memory nodes with 768 GByte memory, and
- 2 huge memory nodes with 1,536 GByte memory per node.

**14** compute racks

with a mixed warm-water and air cooling infrastructure.

**121,920** cores

is the total number of physical cores; as Hyperthreading is enabled twice as much logical cores are available to support highly parallel workloads.

**8** PFlop/s

is the theoretical peak performance of the “Lise” system which demonstrated a LINPACK performance of 5.4 PFlop/s.

**502** TByte

distributed main memory across all compute nodes is installed.

**2,540** CPUs

Intel Xeon Platinum 9242 (2.3 GHz) with 48 cores each.

**Fat Tree**

is the interconnect topology of the Intel Omni-Path network providing a low latency of 1.65 µs maximum and a high bandwidth, i.e., 7.8 TB/s bisection bandwidth, for the communication across the entire system. This interconnect is realized with two 1152-port OPA100 director switches and 54× 48-port edge switches.

**8.4** PByte

online storage capacity is provided by global parallel file systems (DDN Lustre and GRIDScaler).

E

## The 40$^{th}$ Fastest Supercomputer of the World



# The HLRN-IV Opening Event

**The official inauguration of the HLRN-IV supercomputer on December 6$^{th}$** was held at ZIB and was accompanied by high-profile guests and prominent speakers. Witnessed by state ministers and state secretaries of the seven HLRN federal states, Christof Schütte (ZIB), Wolf-Dieter Lukas (BMBF), Steffen Krach (Berlin), Björn Thümler (Lower Saxony), Ursula Morgenstern (Atos) and Hannes Schwaderer (Intel) delivered welcome greetings before pushing the red button to start the official operation of the new HLRN-IV system in Berlin. The highlight of the opening ceremony was the inspiring and entertaining keynote speech of Horst Simon, deputy director of the Berkeley Lab and co-author of the TOP500 list. In his keynote "Supercomputers and Superintelligence" he emphasized the tight relation of HPC and AI by illustrating the enormous potential and challenges of artificial intelligence solutions based on large-scale computations. Finally, in the afternoon session the impact of supercomputer resources for the scientific progress was emphasized. Joachim Sauer, Bettina Keller, and Siegfried Raasch presented their latest results on such diverse topics like heterogeneous catalysis, the dynamics of bio-molecular complexes, and urban climate simulations, and they explained why increasingly large supercomputer capacities are needed to advance their science domains.

## Storage Infrastructure of the HLRN-IV "Lise"

**8.4** PByte

online storage capacity is available in a globally accessible parallel file systems (Lustre) offering a high bandwidth access to application data in batch jobs.

**340** TByte

online storage capacity is provided in a globally accessible NAS appliance for permanent project data and program code development.

**110** PByte

**Peta-Scale**

tape library for long-term archiving of large data sets; operated independently by ZIB.

# ECONOMIC SITUATION IN 2019

**In 2019, the total income of ZIB comprised 25.1 million euros.** The main part of this was made available by the Federal State of Berlin as the basic financial stock of ZIB (10.3 million euros) including investments and Berlin's part of the budget of HLRN at ZIB. The second largest part of the budget resulted from third-party funds (8.6 million euros) acquired by ZIB from public funding agencies (mainly DFG and BMBF) and via industrial research projects. This was complemented by a variety of further grants, such as the HLRN budget made available by other German states or the research service budget of KOBV, summing up to almost 6.2 million euros in total.

## ZIB INCOME

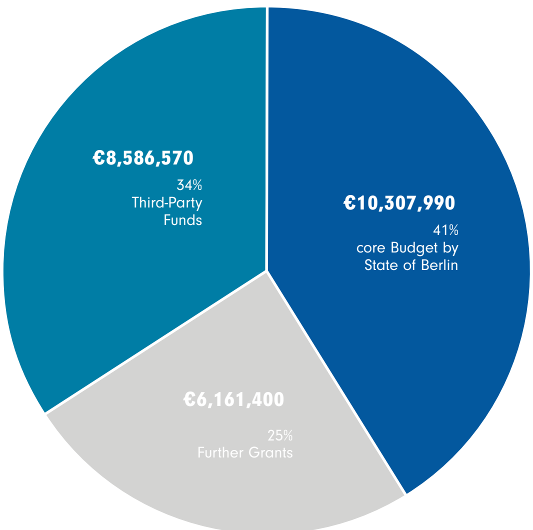

The Zuse Institute Berlin (ZIB) finances its scientific work via three main sources: the basic financial stock of the Federal State of Berlin and third-party funds from public sponsors and those of industrial cooperation contracts.

In 2019, ZIB raised third-party funding through a large number of projects. Project-related public third-party funds declined from 6.425 million euros in 2018 to 6.137 million euros in 2019, while industrial third-party projects rose from 1.707 million euros to 2.450 million euros. In total, 8.587 million euros in third-party funding marked again a new record in ZIB's history – an increase for the eighth year in a row.

## ZIB THIRD-PARTY FUNDS BY SOURCE

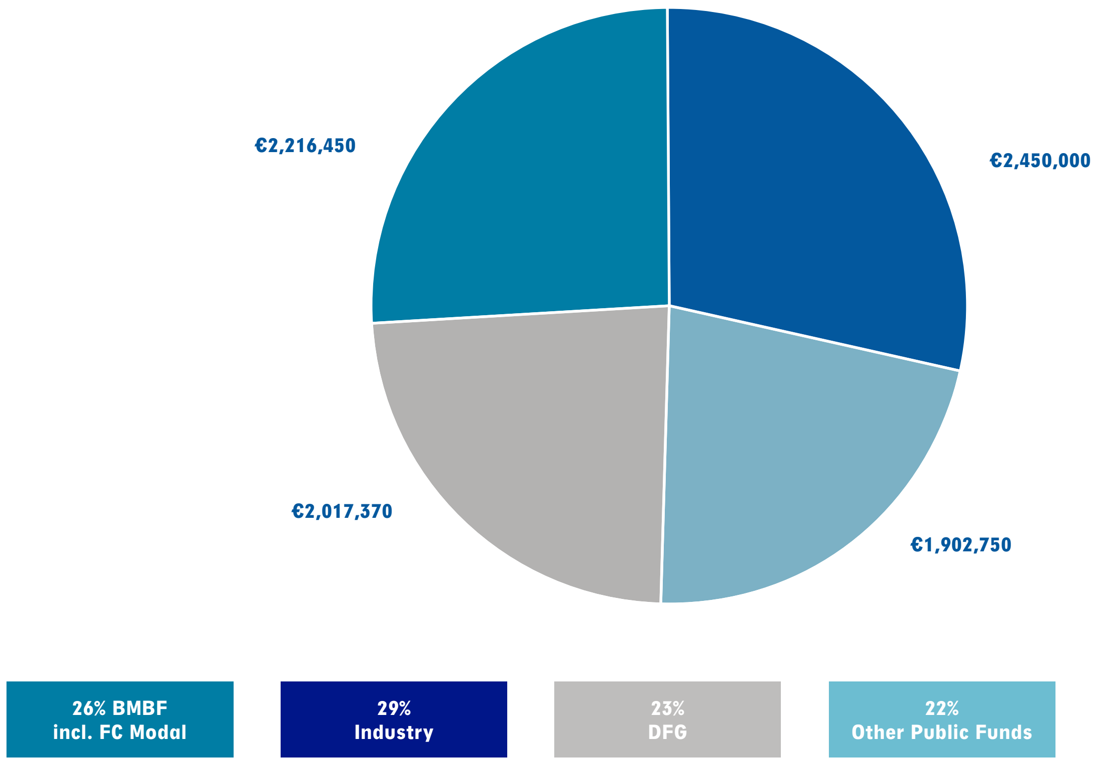

## ZIB THIRD-PARTY FUNDS IN EUROS

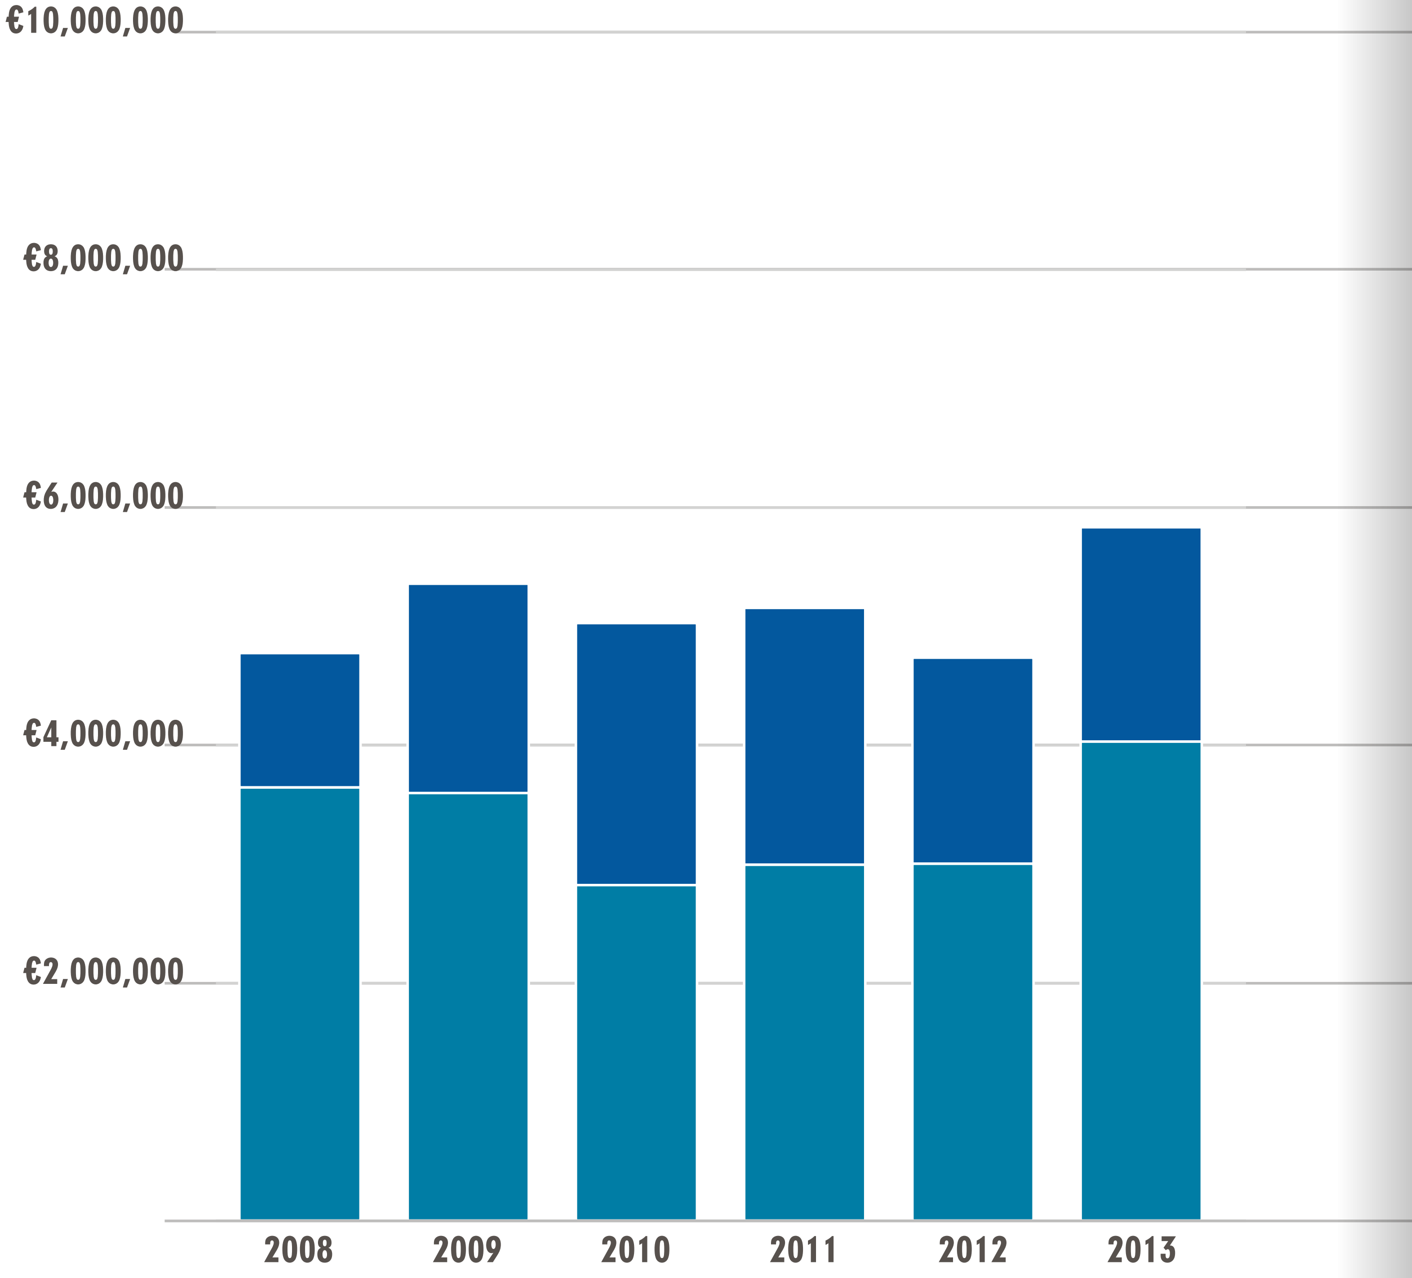

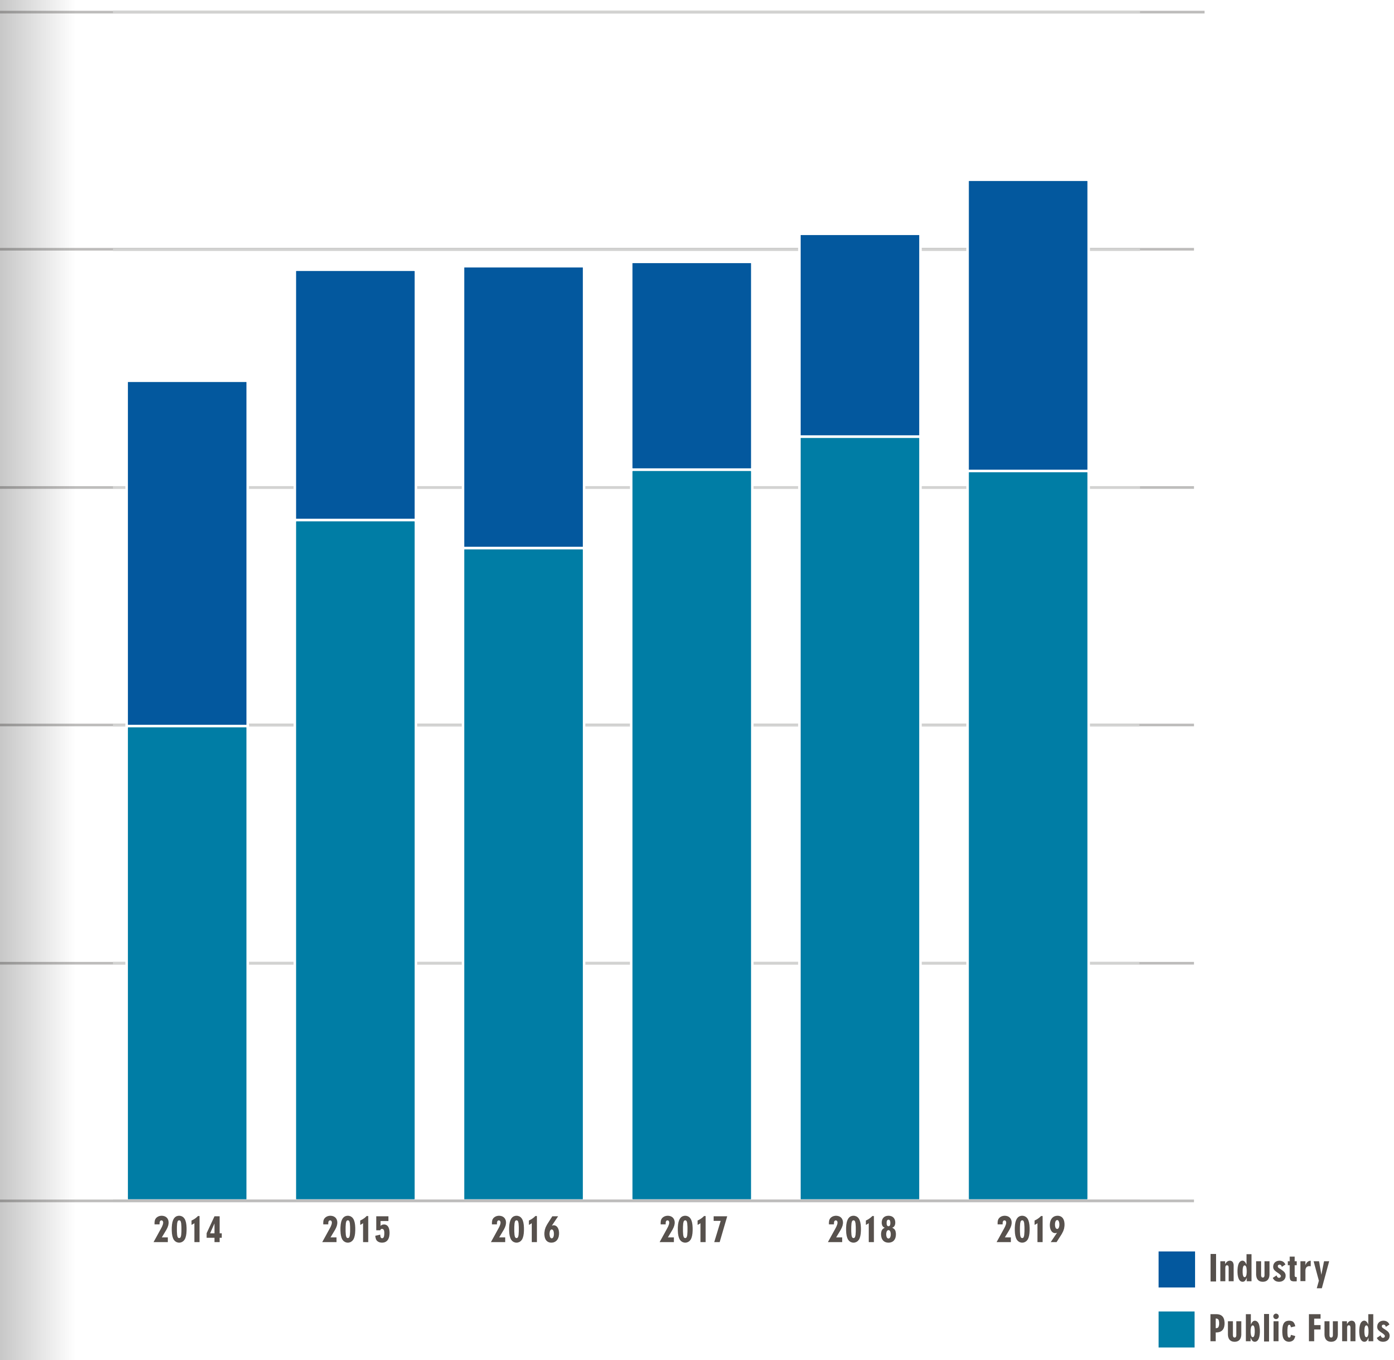
2014
2015
2016
2017
2018
2019
Industry
Public Funds

# 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

**AktuarData GmbH**
1998 | www.aktuardata.de
Development and distribution of risk-evaluation systems in health insurance

**Visage Imaging GmbH**
(Originating from the ZIB spin-off Visual Concepts GmbH)
1999 | www.visageimaging.com
Advanced visualization solutions for diagnositic 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
3-D landscape visualization, software development, and services

**JCMwave GmbH**
2005 | www.jcmwave.com
Simulation software for optical components

**onScale solutions GmbH**
2006 | www.onscale.de
Software development, consulting, and services for parallel and distributed storage and computing systems

**Laubwerk GmbH**
2009 | www.laubwerk.com
Construction of digital plant models

**1000shapes GmbH**
2010 | www.1000shapes.com
Statistical shape analysis

**TASK – Berthold Gleixner Heinz Koch GbR**
2010
Distribution, services, and consulting for ZIB's optimization suite

**Quobyte Inc.**
2013 I www.quobyte.com
Quobyte develops carrier-grade storage software that runs on off-the-shelf hardware

**Keylight GmbH**
2015 I www.keylight.de
Keylight develops scalable real-time Web services and intuitive apps. The focus is on proximity, marketing, iBeacon, and Eddystone for interactive business models

**DoloPharm Biosciences UG**
2017
A specialty pharmaceutical company focused on the clinical and commercial development of new products in pain management that meet the needs of acute and chronic care practitioners and their patients

# Number of Employees

In 2019, 199 people were employed at ZIB; of these, 119 positions were financed by third-party funds. The number of employees decreased in comparison to 2018, mainly becoause of a funding gap between the first and second funding periods of the Research Campus MODAL.

| | 1/1/2019 | | | 1/1/2020 | | |
|---|---|---|---|---|---|---|
| | Permanent | Temporary | **Total** | Permanent | Temporary | **Total** |
| MANAGEMENT<br>* without temporary management | 3 | 1 | 4 | 3 | 0 | 3* |
| SCIENTISTS | 22 | 96 | 118 | 21 | 78 | 99 |
| SERVICE PERSONNEL | 35 | 6 | 41 | 41 | 0 | 41 |
| KOBV HEADQUARTERS | 9 | 7 | 16 | 15 | 1 | 16 |
| STUDENTS | 0 | 42 | 42 | 0 | 40 | 40 |
| **Total** | 69 | 152 | **221** | 80 | 119 | **199** |

# GREEN ENERGY

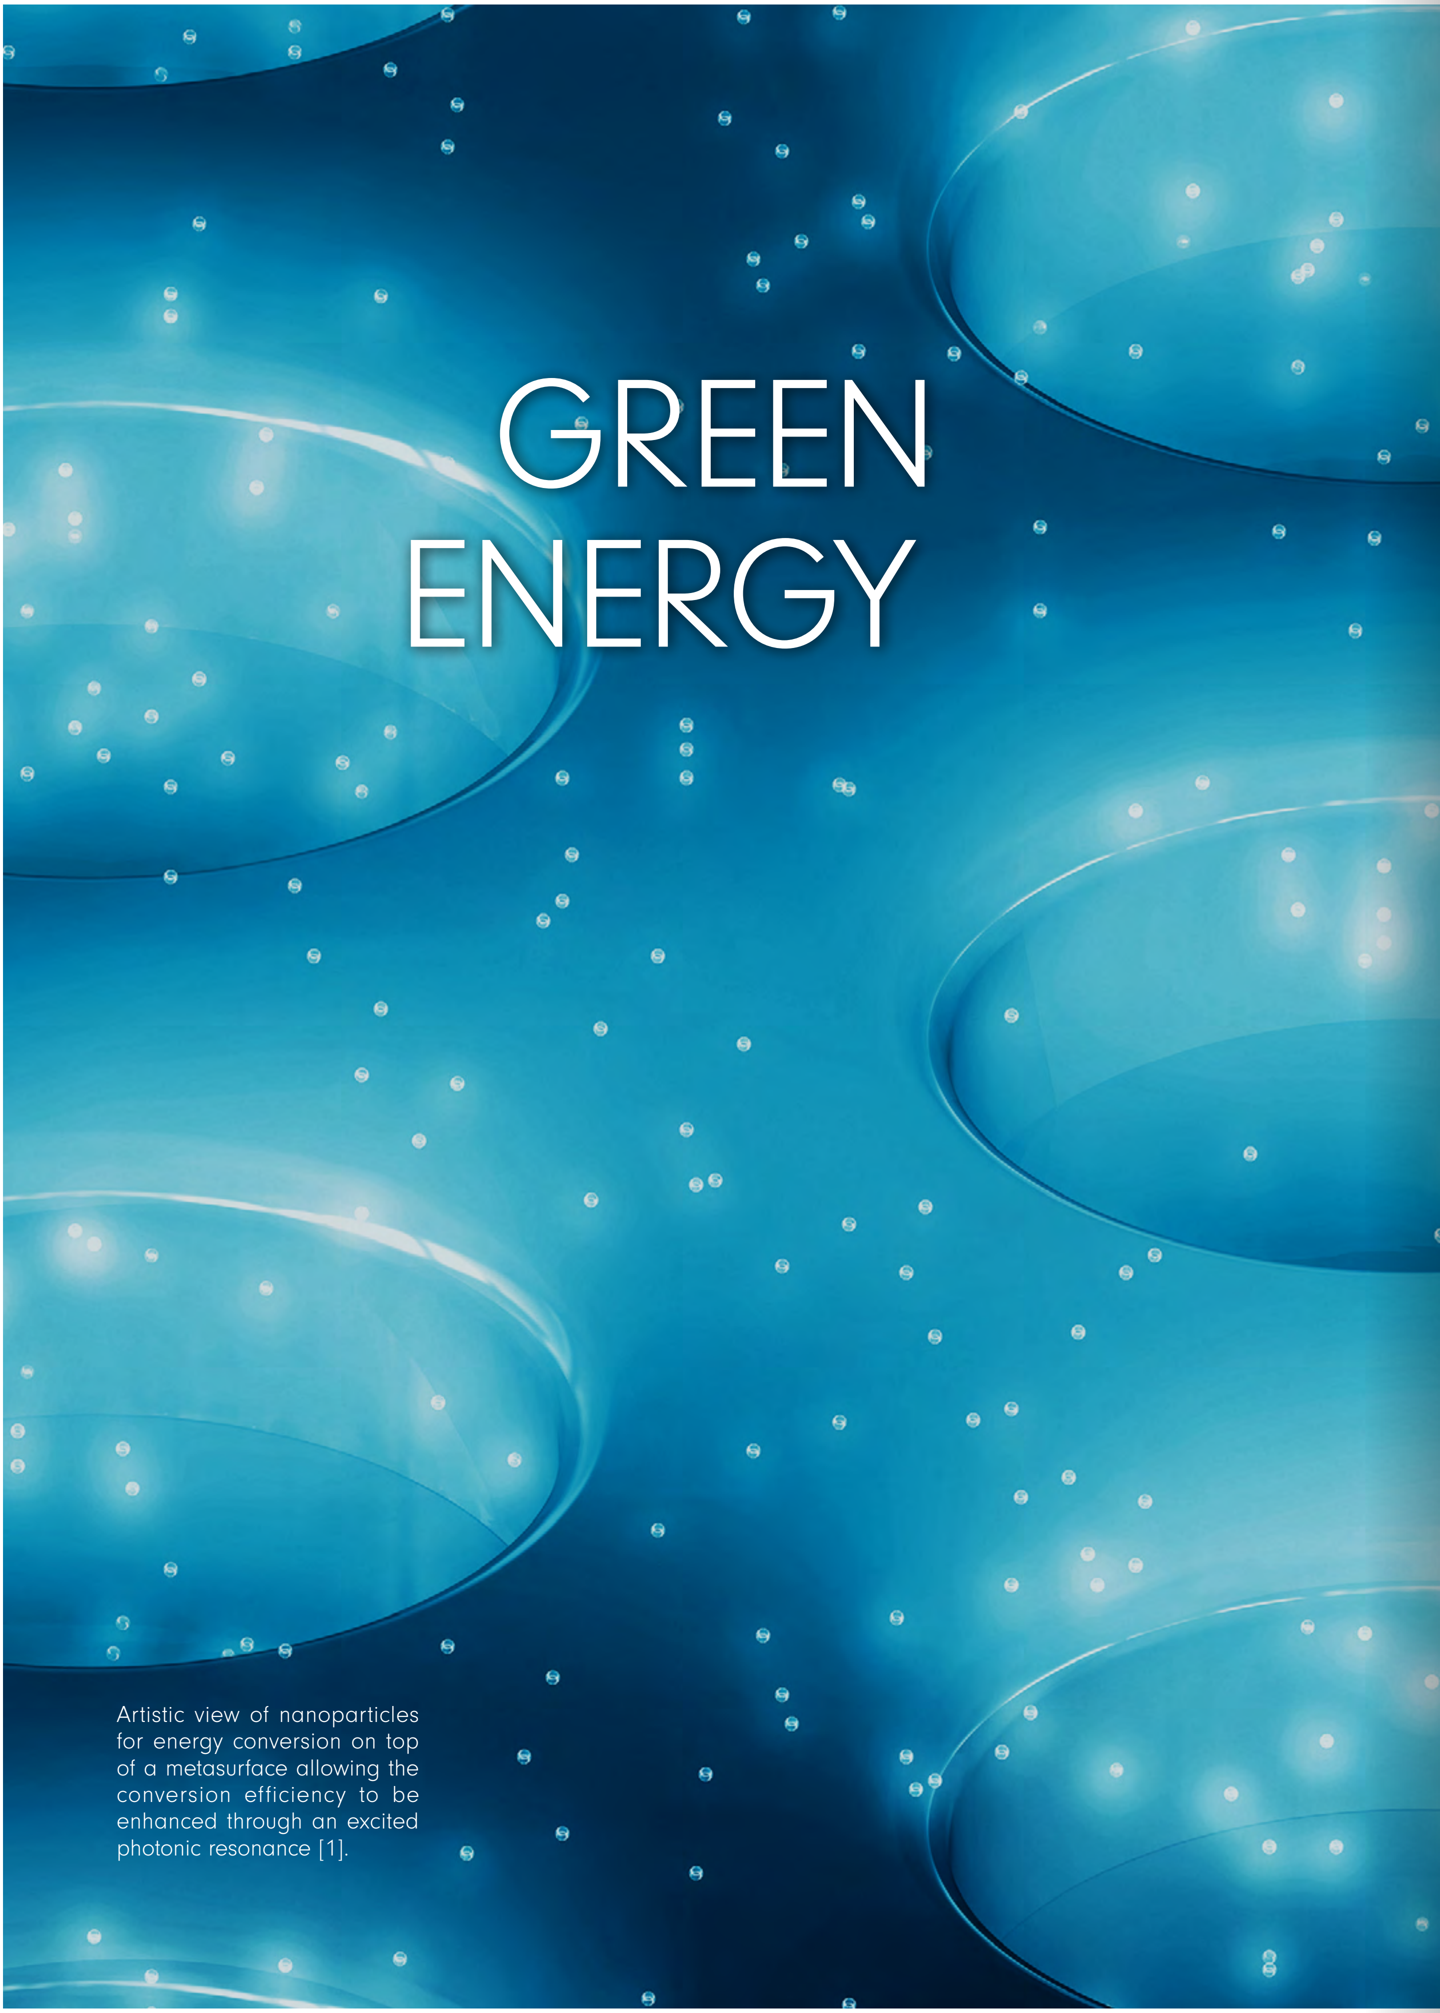

Artistic view of nanoparticles for energy conversion on top of a metasurface allowing the conversion efficiency to be enhanced through an excited photonic resonance [1].

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

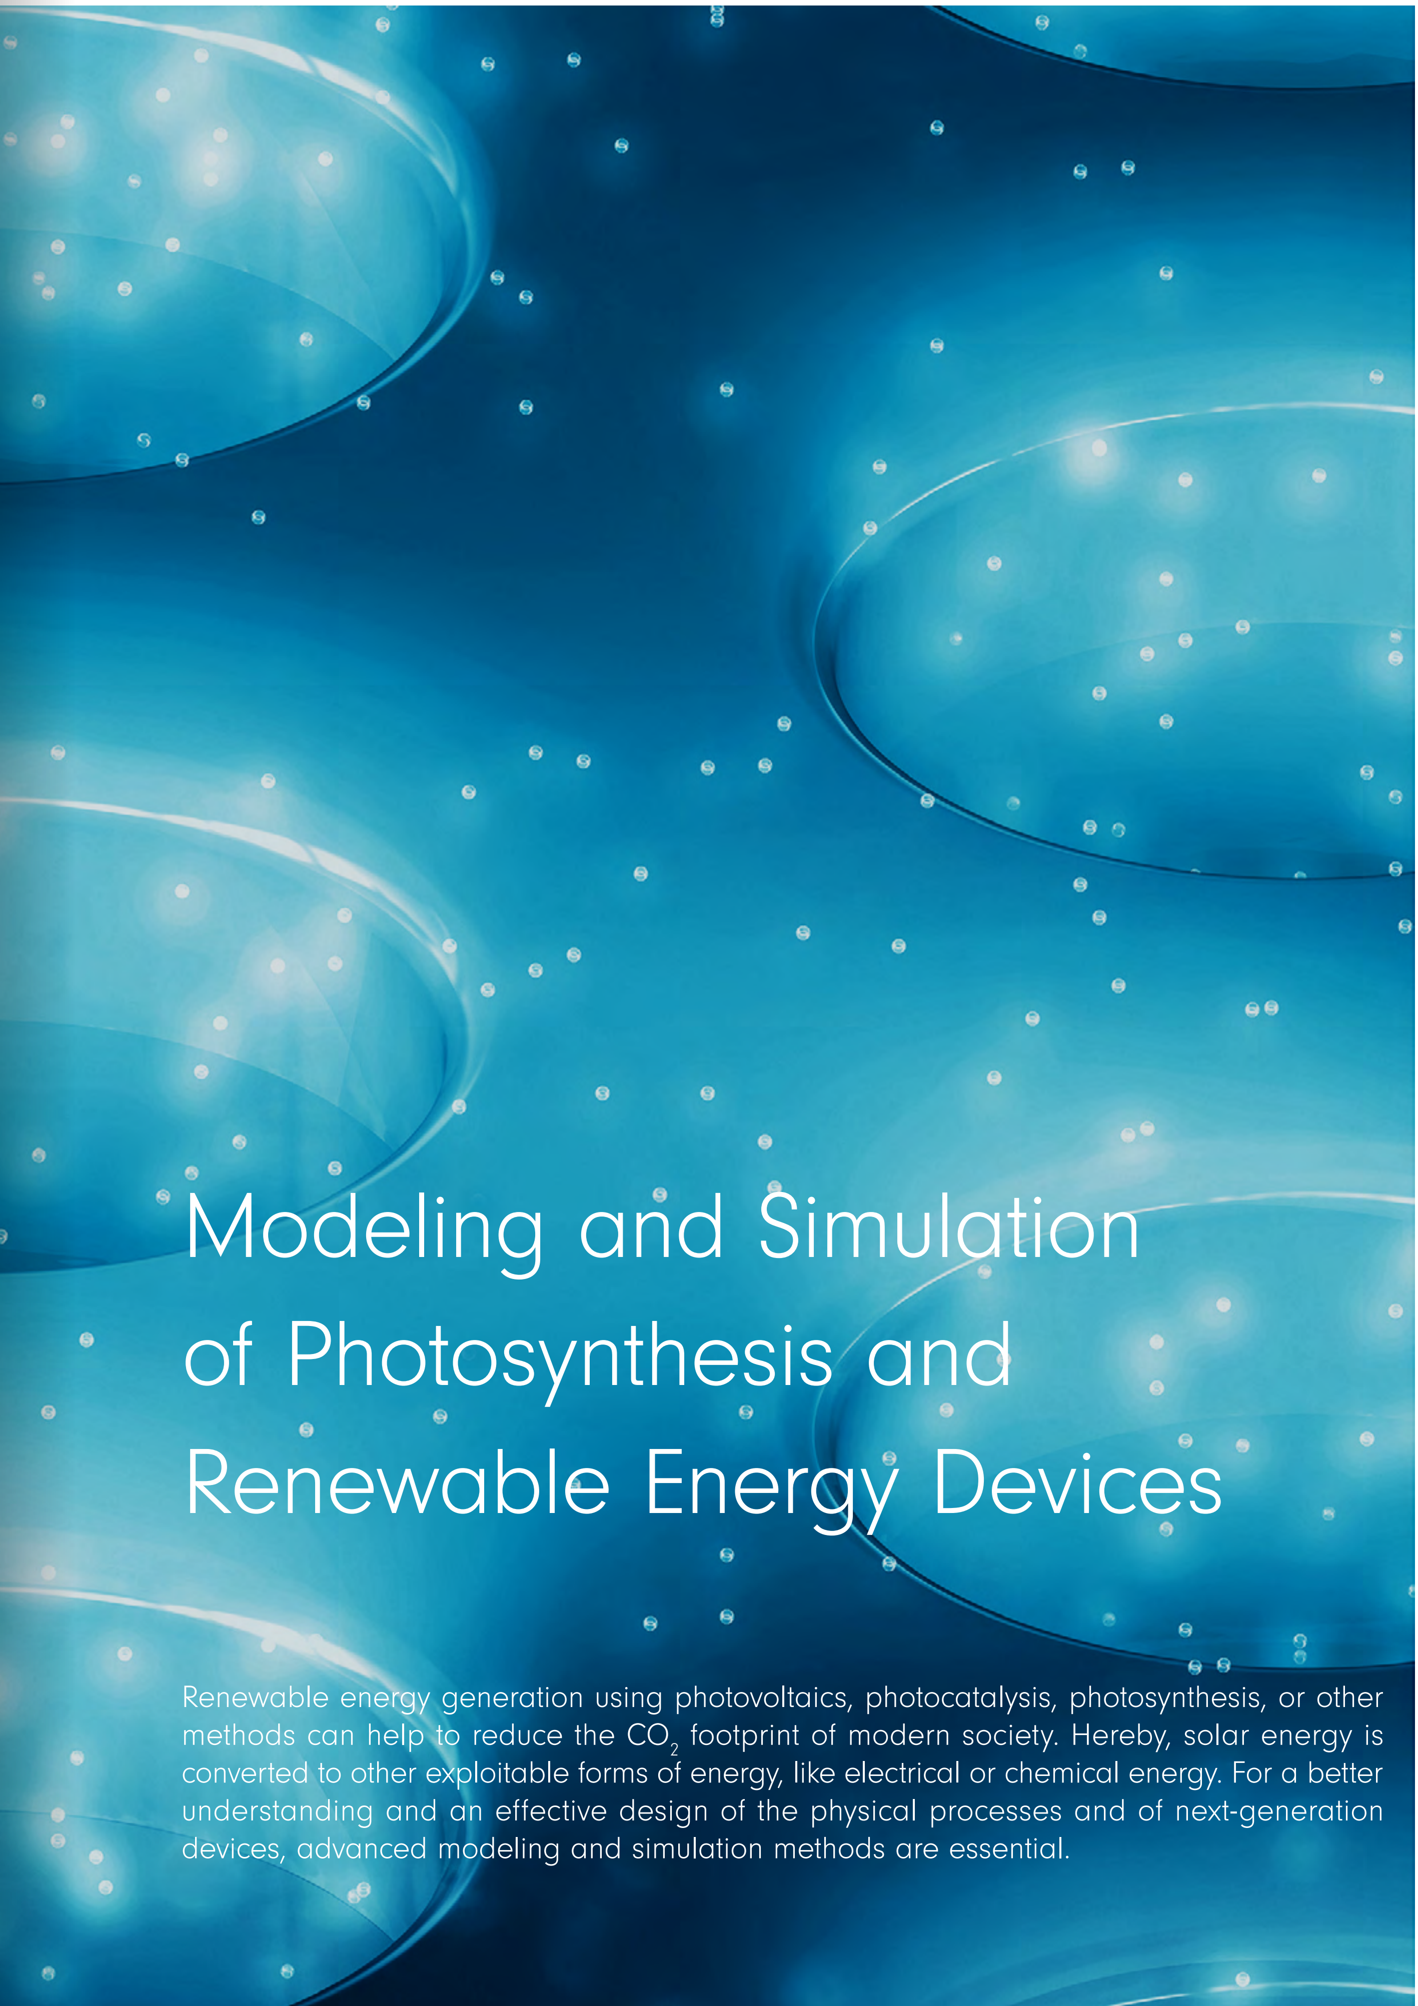

# Modeling and Simulation of Photosynthesis and Renewable Energy Devices

Renewable energy generation using photovoltaics, photocatalysis, photosynthesis, or other methods can help to reduce the $CO_2$ footprint of modern society. Hereby, solar energy is converted to other exploitable forms of energy, like electrical or chemical energy. For a better understanding and an effective design of the physical processes and of next-generation devices, advanced modeling and simulation methods are essential.

# Energy Research at ZIB

Green Energy

Several of ZIB's cooperations are addressing various topics in the field of green energy. Within the Cluster of Excellence MATH+, a cooperation between Helmholtz Zentrum Berlin für Materialien und Energie (HZB) and ZIB, investigates numerical methods needed to design optical energy conversion processes. Within the Helmholtz Excellence Network SolarMath, HZB and ZIB jointly investigate perovskite-silicon tandem solar cells and hybrid components for solar fuel manipulation. The training and networking of young scientists is supported within the Berlin Joint Lab for Optical Simulations in Energy Research (BerOSE) as well as in the Helmholtz Einstein International Berlin Research School in Data Science (HEIBRiDS).

**COLLABORATIONS**

Professor Christiane Becker
Helmholtz-Zentrum Berlin für Energie und Materialien

Professor Bernd Rech
Helmholtz-Zentrum Berlin für Energie und Materialien

Professor Roel van de Krol
Helmholtz-Zentrum Berlin für Energie und Materialien

Dr. Sebastian Matera
Freie Universität Berlin

# Solar Fuel Devices

## Solar Fuel Devices Can Convert Sunlight Directly into Long-Term Storable Fuels

The energy supply network needs to go through tremendous change in order to meet the demands to mitigate climate change. To achieve this, every aspect of the supply chain must be carefully evaluated and optimized. A cornerstone of any sustainable energy network will be solar energy. While conventional solar cells are rapidly increasing their penetration into energy markets, the requirement of electricity networks to provide on-demand power is at odds with the intermittent supply of solar cells. Solar fuel devices are able to convert solar energy directly into usable fuels, thus storing the energy and allowing distribution through existent networks. Of the various fuels that can be created using photochemistry, hydrogen, obtained from water splitting, is the most prominent candidate due to the abundant, low-cost supply of water. Therefore, a better understanding of solar-driven water splitting and developing the technology for water splitting on a large scale is a central challenge facing the transition to a green economy. This technology will likely first be based on the currently available photovoltaic (PV) powered electrolysis. However, it is expected that the next technology step will be to integrate light absorption and electrochemistry into a single device. Optical device design is an important problem for both of these technology steps: one of the limitations of both currently available PV technology and the developing solar fuel technologies is insufficient absorption of light due to reflective losses, transmissive losses, and parasitic absorption. In order to overcome these challenges, modeling, simulation, and optimization [2, 3] plays a crucial role. Proper light management in many photovoltaic devices has already lead to large increases in efficiency, typically through the use of nanostructures. In a collaboration between Helmholtz-Zentrum Berlin für Materialien und Energie and Zuse Institute Berlin, we have recently investigated concepts for absorption enhancement in integrated solar fuel devices. It could be shown that nanostructuring of the device can significantly increase efficiency, and that plasmonic resonances in the optical excitation of nanoparticles can play a major role in optimized devices [5].

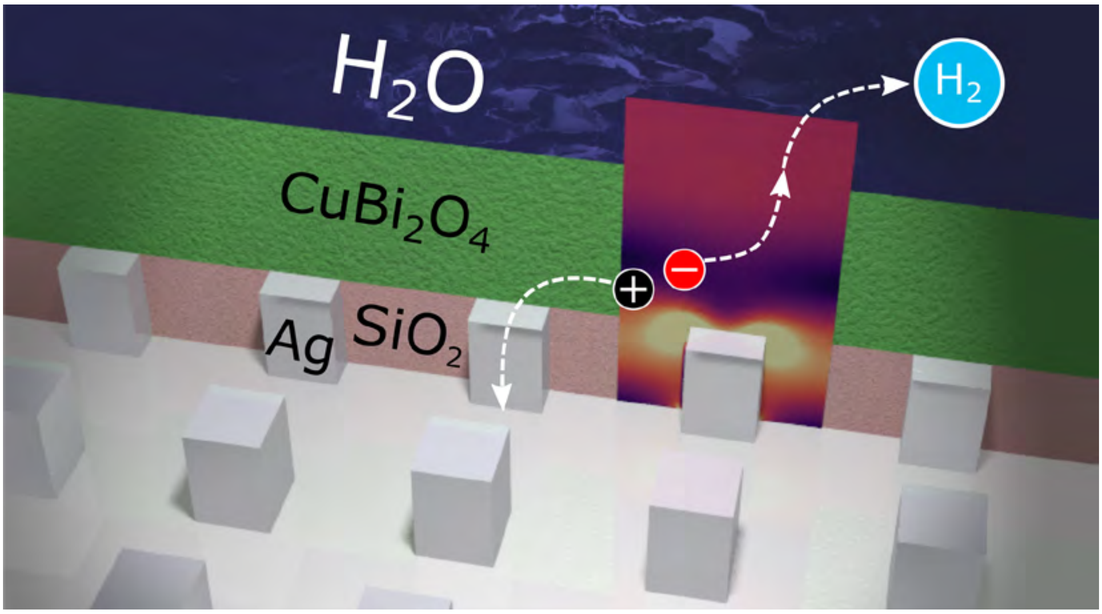


Fig. 1: Artistic view of a solar fuel device. Light is absorbed in a semiconductor material, $CuBi_2O_4$, generated electrons and holes drive the splitting of water to hydrogen ($H_2$) and oxygen ($O_2$, not shown). Metal (silver, Ag) nanostructures enhance the efficiency through providing plasmonic resonances (inset with electromagnetic field distribution in pseudocolor scale).

# Photosynthesis: Catch-As-Catch-Can for Photons

Photosynthesis fuels life on earth by converting incoming solar radiation into chemical energy. Some photosynthetic organisms have adapted particularly well to low light environments. For instance, green sulfur bacteria use "leftover" near-infrared light to thrive at places where other wavelengths have already been absorbed by other species. The sulfur-rich volcanic ponds in the Yellow Stone National Park provide a suitable natural habitat.

To further concentrate the available light, green sulfur bacteria developed a multi-stage photon trap: starting from a large antenna complex, grounded on a base-plate structure, the energy is directed toward the reaction center [5]. This link is established by the Fenna-Matthews-Olson complex, the first photosynthetic complex whose structure was determined by X-ray crystallography [6]. The bacteriochlorophyll-a complex embedded in the photosynthetic apparatus is tuned by surrounding proteins to establish an energetic funnel that guides the photoexcitation toward the reaction center. The directed energy transfer comes at a price: part of the light energy is converted to heat in order to guide the larger remaining part to the reaction center.

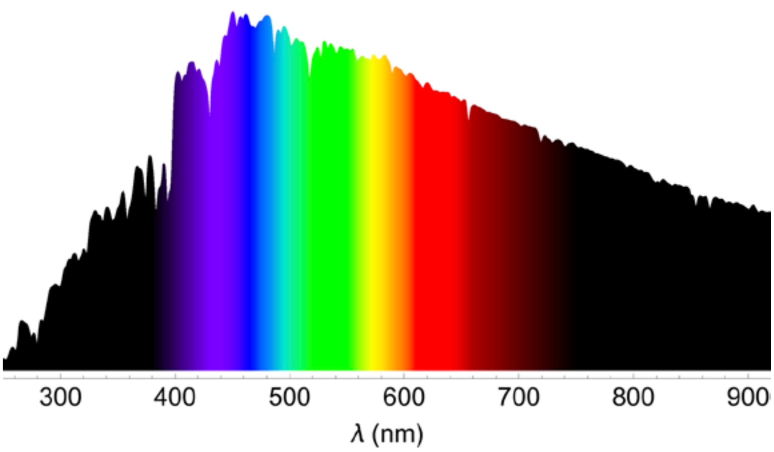


Sunlight is composed of light with different wavelengths. Green sulfur bacteria are harvesting near-infrared photons (750 nm), whereas plants absorb light in the blue and red colors, giving rise to the green color of leaves.

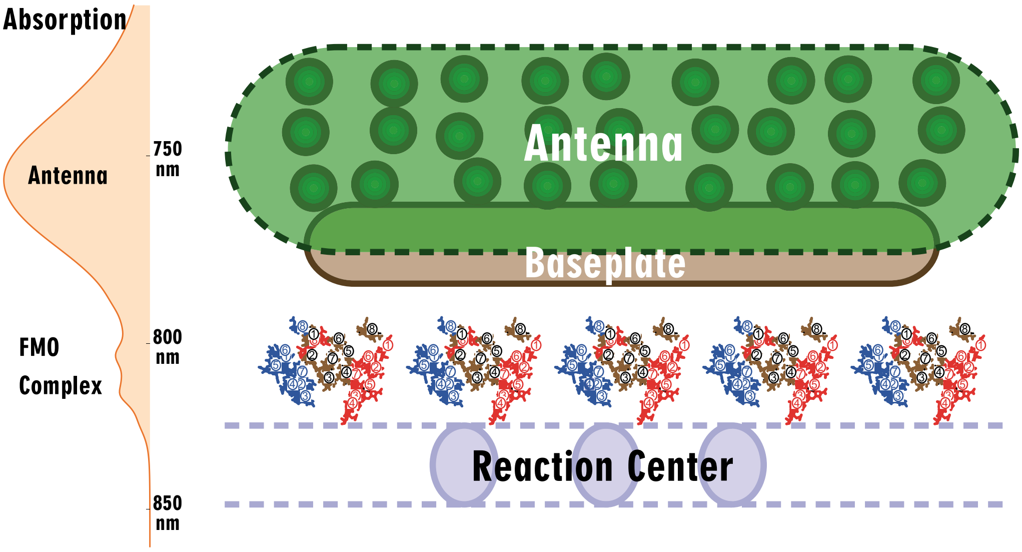


Schematics of the photosynthetic apparatus of green sulfur bacteria (C. tepidum), adapted from [5].

Since 2005, it has been possible to directly image the energy transfer using a sequence of laser pulses that first excite and later – at a specific time delay – probe the energetic state of the molecular complex [7]. The energy transfer is then directly monitored by the movement of emission peaks. At ZIB, we have developed methods to efficiently compute and predict the energetic pathways and to elucidate the intermediate steps. To capture the crucial role of the heat transferred to molecular vibrations, these computations make use of the quantum-mechanical density matrix description [8, 9]. The strong interaction of vibrational and electronic excitations forms an entangled state of the molecular complex.

Recent experiments employ fast-varying polarization changes of laser pulses to further probe the configuration and arrangement of the bacteriochlorophylls. The modeling of the signals requires considering large ensembles of molecular complexes, a task ideally suited for the HLRN supercomputer [12]. The resulting large data sets are efficiently encoded in neural networks using machine learning and are used to investigate the arrangement of different photosynthetic complexes.



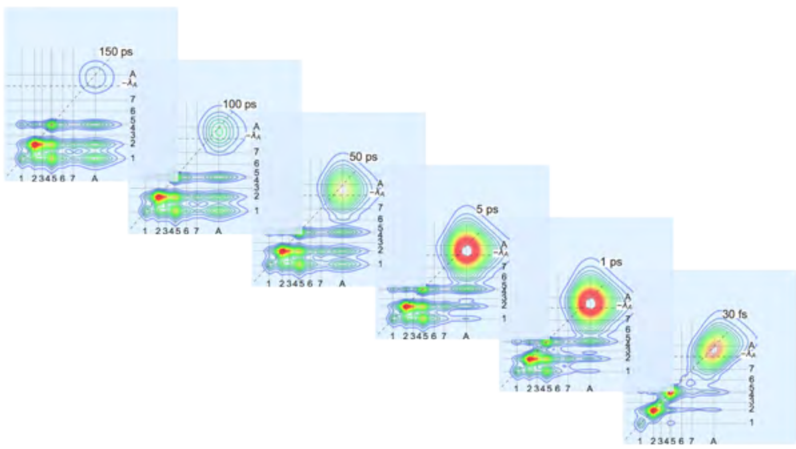


Time sequence of the energy transfer in green sulfur bacteria (C. tepidum). The horizontal axis denotes the excitation energy and the vertical axis the de-excitation energy. A decay of intensity to the lower diagonal tracks the energy transfer [10, 11].

# Light Management for Highly Efficient Photovoltaic Devices

## Nanotextured Perovskite Solar Cells



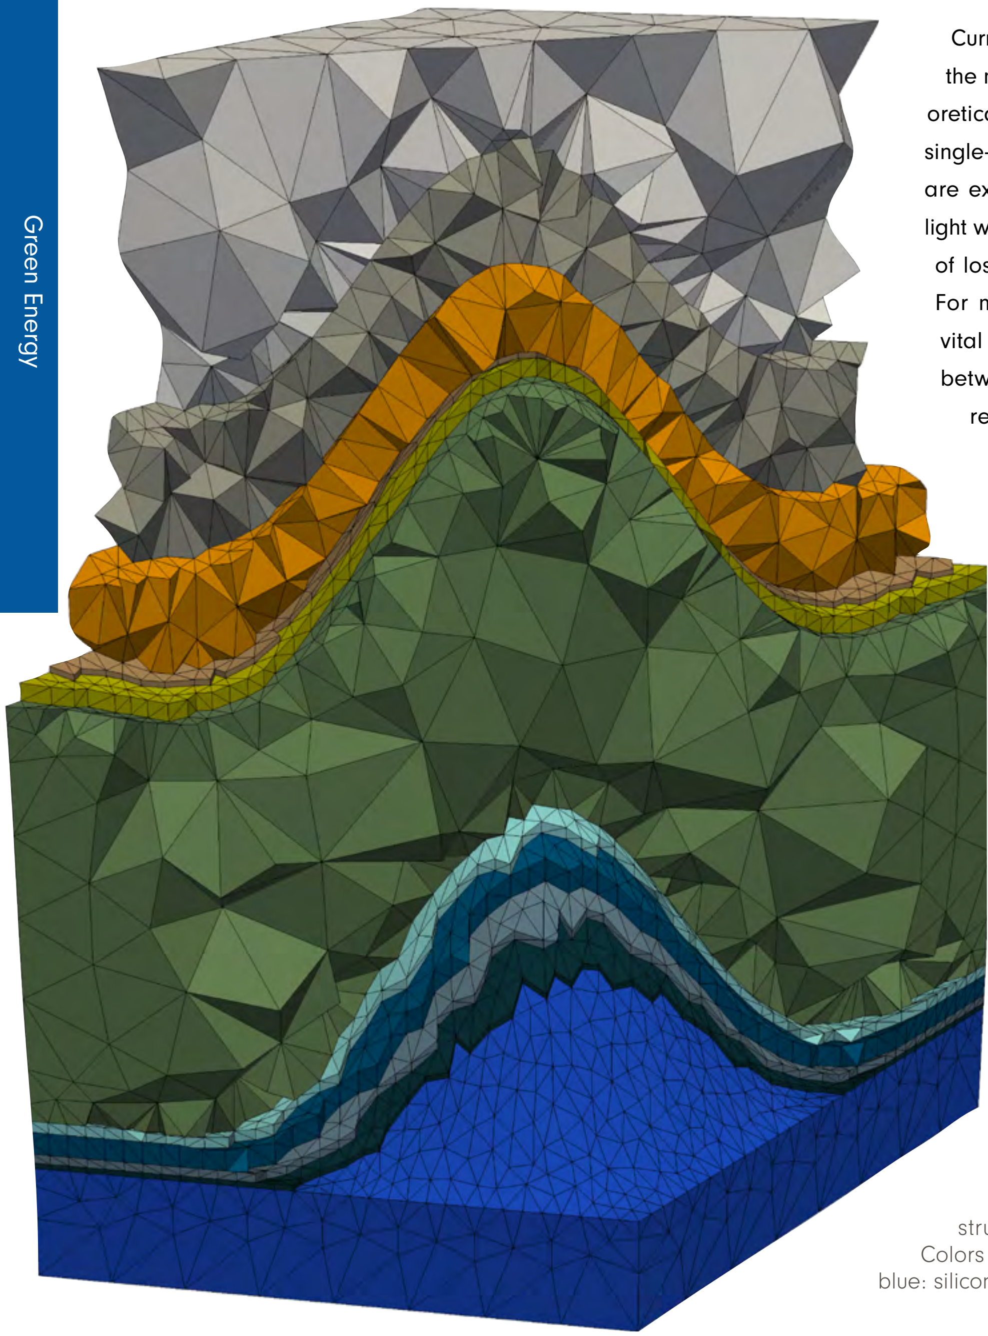

Currently, perovskite-silicon tandem solar cells are the most investigated concept to overcome the theoretical limit for the power conversion efficiency of single-junction silicon solar cells. Optical simulations are extremely valuable to study the distribution of light within the solar cells, and allow the minimization of losses from reflection and parasitic absorption. For monolithic perovskite-silicon solar cells, it is vital that the available light is equally distributed between the two subcells, which is known as current matching. At ZIB, we develop advanced finite-element-method-based (FEM) simulation tools for optimizing the solar-cell stacks and nanotextures. These tools are used to study, for example, how different light management approaches influence the sensitivity of the solar module to the illumination condition [13].

Tetrahedral mesh for a FEM simulation of electromagnetic field absorption in a nanostructured perovskite-silicon tandem solar cell. Colors indicate different materials (green: perovskite, blue: silicon).

# Realistic Weather Conditions Impact Bifacial Solar-Cell Installations

Bifacial solar cells are a promising technology with a quickly increasing market share. By converting light received at both sides of the module, the power output of a solar cell can be significantly increased. Modern cell designs are intrinsically suitable for bifacial modules and therefore the difference in production costs of bifacial compared to classic modules is relatively small. It is however challenging to estimate the energy output and to find optimal configurations of bifacial photovoltaic installations. Within the Joint Lab BerOSE, models for data-driven optimization of modules in field installations have been developed. This allows the amount of light that falls on the front- and backside of a module to be computed. The tilt angle of the modules and the distance between the module rows are important parameters that determine how much electricity can be generated. With increasing distance between the rows, the power per module will increase, however it also becomes more cost intensive due to the increasing ground consumption. In order to find the optimal configurations of tilt and row distance for realistic weather data, the combination of an economical model and a tool for optical simulation has been used. This allows the levelized cost of electricity (LCOE) that relates the total amount of energy produced by a power plant to the total cost over its lifetime to be calculated. The optical simulation can be performed with weather data that is readily available for large parts of the world. By using Bayesian optimization, we obtain parameters that allow LCOEs that are significantly reduced [14], compared to rules of thumb, as they are employed in current state-of-the-art installations.

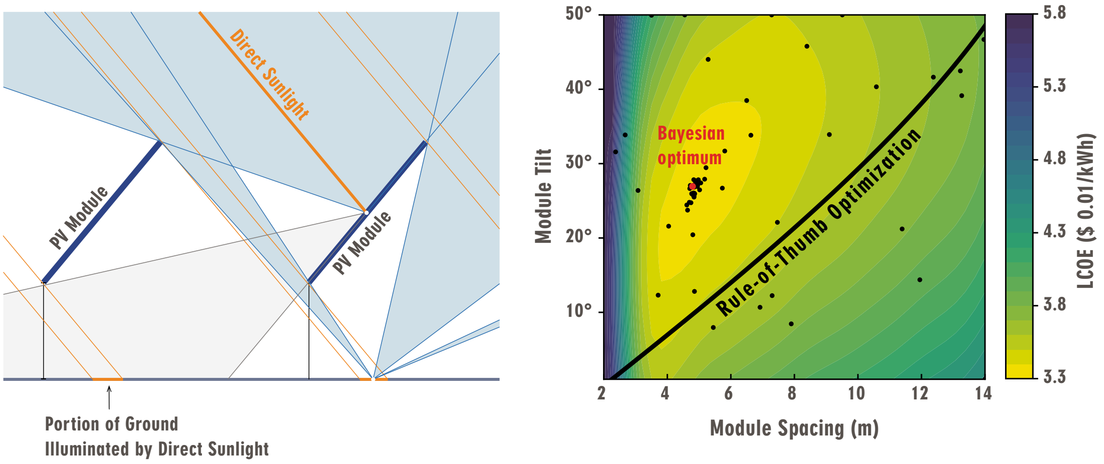


Model of an illuminated installation of bifacial solar-cell modules. Depending on the tilt angle of the modules and on the module spacing, different portions of direct sunlight and of sunlight scattered from the ground reach the front and back surfaces of the solar cell. Right: Levelized cost of electricity (LCOE) depending on module tilt and spacing, for realistic weather data and economic model at a specific location of the installation in the US. The optimum installation parameters allow for significantly reduced costs with respect to rule-of-thumb optimization.

Christoph von Tycowicz | +49 30 84185 - 350 | vontycowicz@zib.de

# RIEMANNIAN ANALYSIS OF TIME-VARYING SHAPE DATA – UNDERSTANDING GEOMETRIC EVOLUTIONS

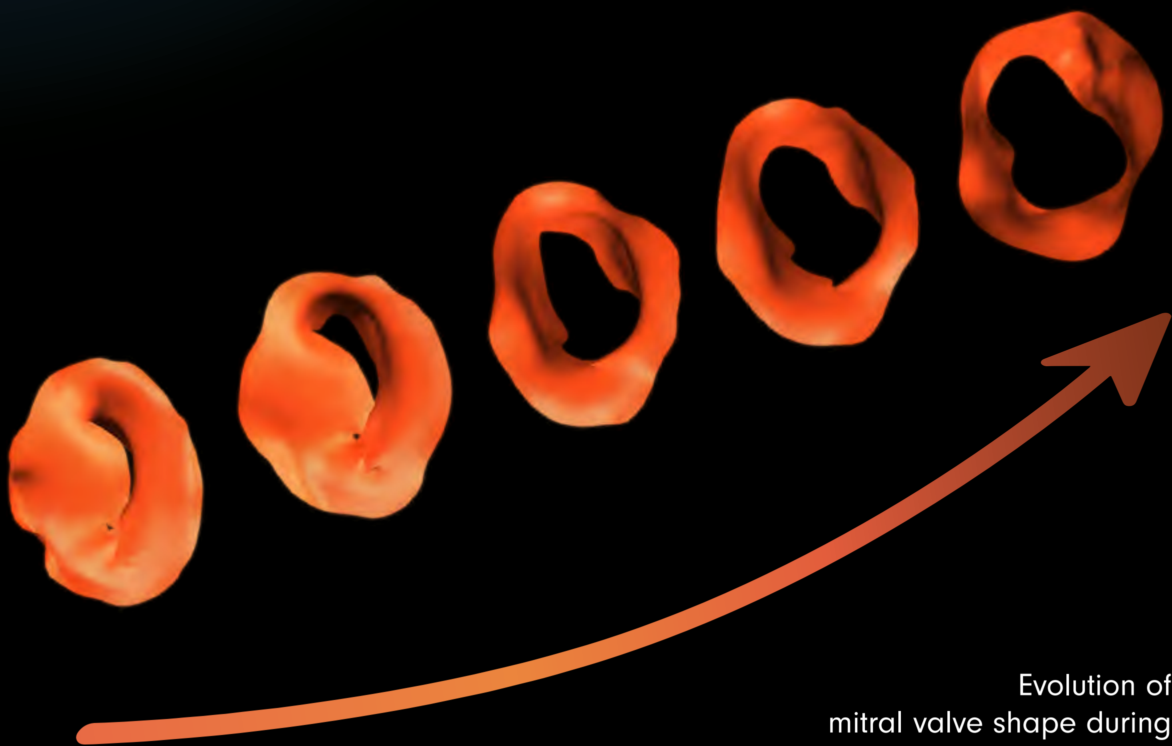

Evolution of
mitral valve shape during
distole phase of the cadiac cycle.

# Motivation

Material objects, both in the inanimate and the animate world, are characterized by their material and their shape. In our work we focus on the geometric aspect, i.e. the shape. Instead of characterizing shapes by a few meaningful parameters, we consider shapes in their entirety. This allows a full geometric characterization of empirically given shapes and more differentiated assertions about these. From a mathematical point of view, we consider individual shapes as single points in a high-dimensional space, the so-called “shape space”.

For many applications, e.g. in medicine or biology, the characterization of sets of shapes is of interest: What types of shapes do occur in a given ensemble? How frequently do the different types of shape show up? What is a typical and what is an atypical shape? These are questions where statistics comes into play, or mathematically speaking, statistics in high-dimensional shape spaces. In recent years, appropriate mathematical theories and computational tools have been developed. Statistical shape modeling makes it possible to consider all shape features and their correlations at once without having to pre-define discrete shape measurements. The complete coverage of the shape, including the correlations of the degrees of freedom, allows more differentiated conclusions and provides better quantitative measures that are statistically significant.

The focus of our project is the analysis of time-varying shapes. These exist in abundance, both in everyday life and in science – especially in life, environmental and earth sciences, but also in engineering and cultural sciences. A large class in life sciences are biological changes of shapes within and between individuals. Such changes can be tracked over time in longitudinal imaging studies, e.g. to gain insights into dynamic processes, such as ageing or disease progression. Another example is tracking the beating heart during the cardiac cycle, where atypical shape changes can indicate certain diseases and precise analysis of shape changes can potentially provide further diagnostic information. Also quite different time scales can be involved: for example in the analysis of the stylistic development of ornaments over cultural periods, or the analysis of the evolutionary development of bony structures in living beings.

We have extended the mathematical tools for the statistical analysis of (static) shapes to the analysis of time-varying shapes – or in mathematical terms, to shape trajectories, i.e. curves in the shape space. These trajectories are discretely sampled by the empirically given shapes. They are themselves geometrical objects and thus amenable to geometric techniques. Since each individual shape is defined by many degrees of freedom and for statistical reasons many objects have to be included in the analysis at many points in time, we are dealing with “big data”. A central concern is therefore the development of efficient algorithms.

# Shapes and Shape Spaces

For a mathematical concept of “shape”, we can start directly from the everyday concept in which we refer to the outer boundary (i.e. the surface) of an object. A possible framework for the formalization of the concept of shape is based on the comparison of related forms: An object class under study can be represented by a common deformable template that accounts for the typicality of the objects’ structure. The shape variability is represented by deformations that are applied to the template. Codifying shapes in such a way allows to interpret them as elements in a high-dimensional space of deformations. This so-called configuration space not only encodes the geometric form of objects but also their scale, position and orientation within the 3D space they are embedded in. By identifying shapes through similarity transformations we obtain the shape space [3]. This is suitable to statistical shape analysis. The last step (mathematically, a quotient taking through a group action), however, introduces curvature to shape space. Contrary to flat spaces, shortest connecting paths in shape space are not straight lines but curved trajectories referred to as geodesics (see Fig. 1).



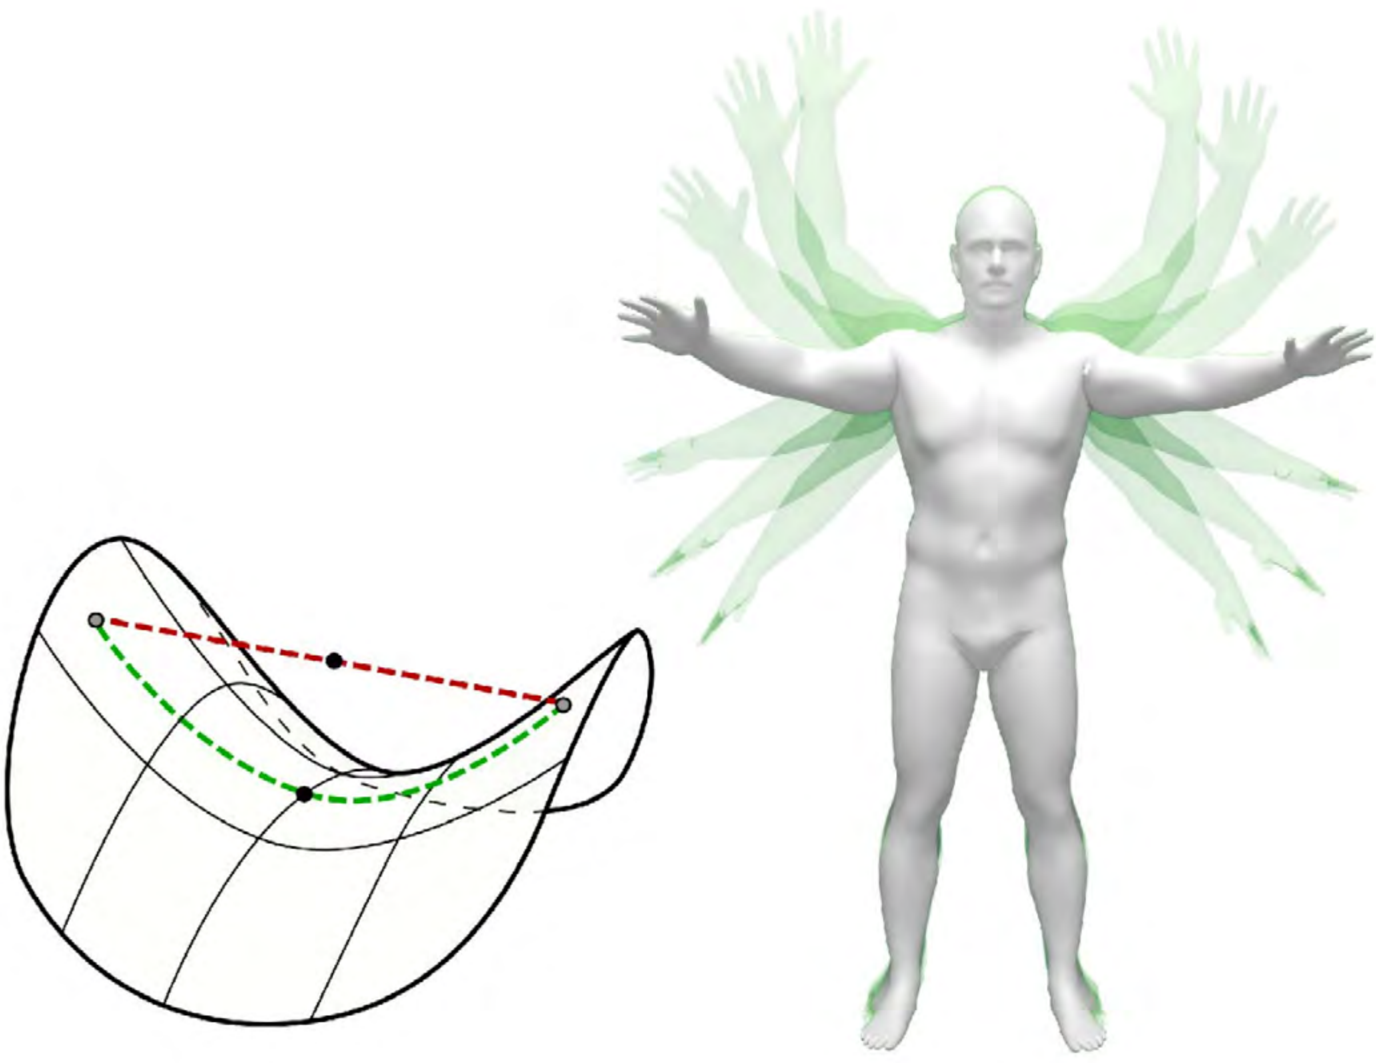

Fig. 1: Visualization of shortest paths, i.e. geodesics, connecting two body shapes w.r.t. the flat ambient space (red) and a curved shape space (green). The latter contains only valid shape instances whereas the former contains shapes with artifacts, like shrunken arms.

The nonlinearity of shape space further implies that there is no global system of coordinates. Consequently, there is no direct way to compare pairwise differences between shapes. To obtain a consistent description of structural changes at the population level (e.g. for group-wise analysis), the differences need to be transferred into a common reference frame. Among the different techniques proposed, constructions based on parallel transport provide the most natural approach. As parallel transport is rarely given in closed form, in general it has to be approximated numerically, e.g. employing Schild's ladder or fanning. In particular, except for the limited setting of planar shapes, this is the case for Kendall's shape space [5]. Utilizing closed form expressions of Kendall's pre-shape sphere, we reduce parallel transport to the solution of a homogeneous first-order differential equation that allows for highly efficient computations [7]. Moreover, we reduce the important case of parallel transport along a geodesic path to the solution of a low-dimensional equation that only depends on the dimension of the ambient space and not on the spatial resolution of the discrete representation.

Whereas the nonlinearity of the shape space ensures consistency, e.g. by preventing bias due to misalignment of shape observations, it also impedes the application of classical statistical tools. As a fully intrinsic treatment of the analysis problem can be computationally demanding, a common approach is to approximate it using extrinsic distances. For data with a large spread in shape space or within regions of high curvature, such linearization will introduce distortions that degrade the statistical power [9]. Therefore, we derive novel geometric structures that allow for efficient shape analysis and, thus, facilitate applications that require interactive response or involve large shape populations. To this end, we employ differential coordinates that are derived from the (deformation) gradient of the map that encodes the shape relative to the template and, hence, naturally belong to the group of orientation preserving linear transformations $GL^+(3)$ [2]. Performing intrinsic calculus on this representation allows for fast computations while, at the same time, accounts for the nonlinearity in shape variation. Following a surface-theoretic approach, we further extend the differential coordinates based on discrete fundamental forms to derive a shape representation that is invariant under Euclidean motion and thus alignment-free [1]. We endow this representation with a Lie group structure that admits bi-invariant metrics and therefore allows for consistent analysis using manifold-valued statistics based on the Riemannian framework. The rich structure of the derived shape space yields highly discriminative shape descriptors providing a compact representation that is amenable to learning algorithms. We evaluate the performance of our model w.r.t. shape-based classification of pathological malformations of the human knee and show that it outperforms state-of-the-art approaches especially in the presence of sparse training data.

# Regression

Given empirically-defined shape trajectories we would like to estimate the relationship between the observed variables, i.e. the shapes and their co-varying parameters. To this end, regression analysis is a reliable statistical approach to identify which variables have an impact on a quantity of interest. The process of performing regression allows one to confidently determine which parameters matter most, which parameters can be ignored, and how these parameters influence each other. Typically, the evolution of shapes is assumed to be smooth and related to a single explanatory variable (usually time), for example when studying growth patterns. In this context, spatiotemporal regression models allow to estimate continuous trajectories from sparse and potentially noisy samples. They also provide a way to describe the data at unobserved times (i.e. shape changes between observation times and – within certain limits – also at future times). Such a time interpolation of data further allows to compare shape evolutions between different subtypes that have been observed at unequal time points.

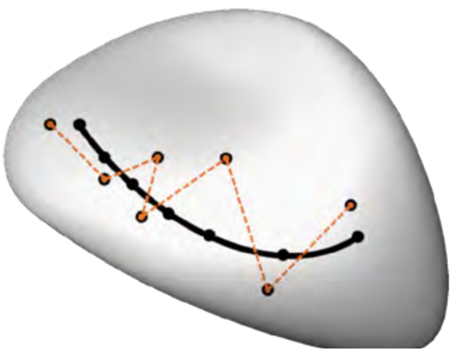

Fig. 2: Schematic depiction of best fitting geodesic (left) and a Jacobi field (right).

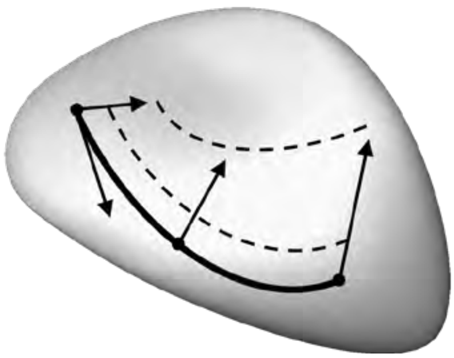

As shape spaces lack a global vector space structure, any linear combination of shapes may not lie on the manifold. Furthermore, embedding the manifold-valued variables in a Euclidean space might introduce distortions and, thus, results in a poor estimation of the model. These problems advocate the development of novel, intrinsic regression methods for non-Euclidean data. The most widely used approach is to approximate the observed temporal shape data by geodesics in shape space. Geodesic models are attractive as they feature a compact representation (similar to the slope and intercept term in linear regression) and therefore allow for computationally efficient inference. In particular, the goal of geodesic regression is to find a geodesic curve in shape space that best fits the data in a least-squares sense.

In the absence of an analytic solution, the regression problem has to be solved numerically. Again, to obtain consistent and efficient computational schemes the non-Euclidean structure has to be taken into account. To this end, the gradient of the cost function can be computed using Jacobi fields (Fig. 2, right), since they express the derivatives of the exponential map. Based on an analytic derivation for these expressions [7], we can fully leverage the geometry of the shape space using Riemannian optimization procedures. Such an intrinsic approach leads to highly efficient algorithms reducing the computational expense by several orders of magnitude over general, nonlinear constrained optimization.

# Hierarchical Models

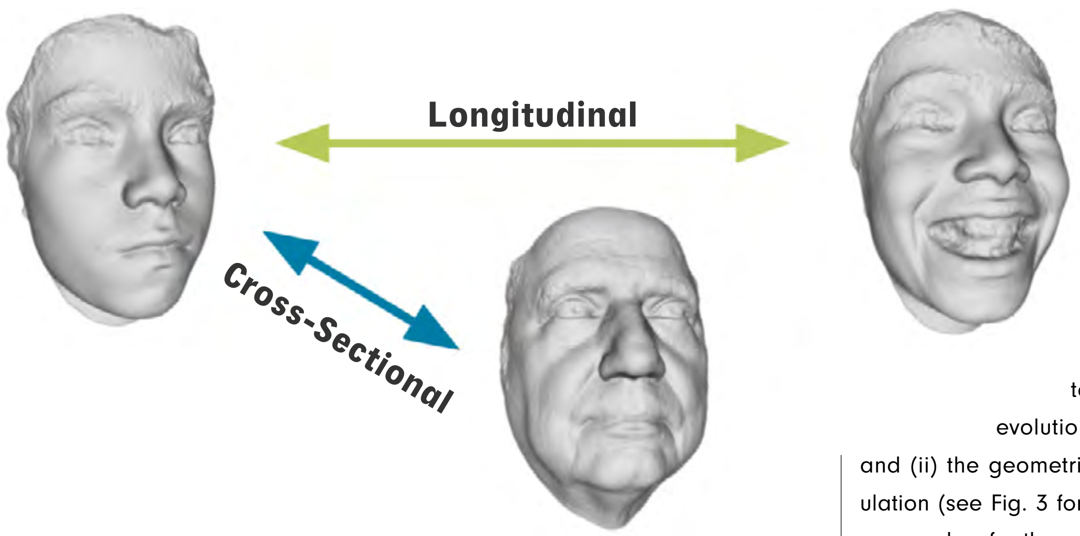


Fig. 3: Different types of shape changes.

Processes such as disease recovery, style progression of prehistoric remains, or growth are inherently time-dependent, requiring measurements at multiple time points to be sufficiently described. To gain insight into such dynamical processes, morphological studies rely on longitudinal datasets that capture shape changes within and across subjects over time. When analyzing such observations of shape trajectories we have to distinguish between morphological differences due to (i) temporal shape evolutions of a single subject and (ii) the geometric variability in a population (see Fig. 3 for an illustration). While approaches for the analysis of time series of scalar data are well understood and routinely employed in statistics and medical imaging communities, generalization to complex data such as shapes are at an early stage of research. Methods designed for cross-sectional data analysis, e.g. regression, do not consider the inherent correlation of repeated measurements of the same subject, nor do they inform how an instance relates to a population-average trend. Therefore, analysis methods for longitudinal datasets must capture and disentangle the cross-sectional variability in shape and the temporal variability due to underlying processes of change. Fig. 4 shows a synthetic example where cross-sectional regression fails to estimate a population-average trend correctly.

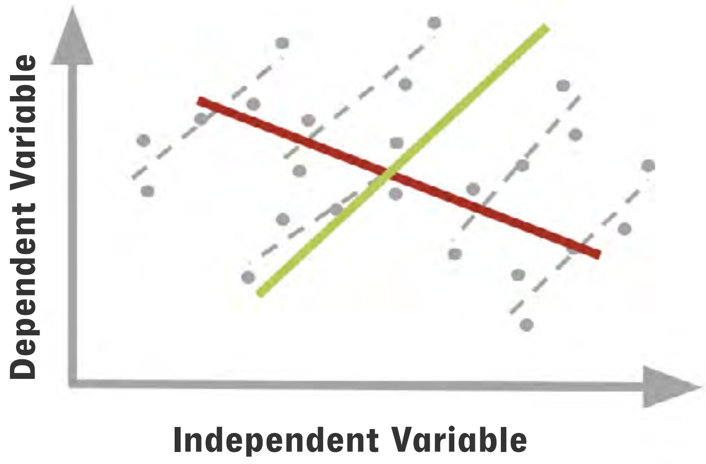


Fig. 4: Sketch of longitudinal dataset together with population-average trends estimated by cross-sectional (red) and a hierarchical (green) model.

These problems motivate the use of hierarchical models that include intra-individual changes in the response variable and thereby have the ability to differentiate between cohort and temporal effects. In the first stage of a hierarchical model, inner-individual changes are modeled as smooth, parametric curves determined via spatiotemporal regression. In the second stage, these subject-specific trends are considered as disturbances of a population-average trajectory. To this end, a principled way of comparing shape trends is needed. In particular, this requires a notion of distance for shape trajectories that is consistent with the Riemannian metric of the underlying shape space. As for manifold-valued regression, state-of-the-art approaches model shape trends as geodesic trajectories, which can be parametrized as points in the tangent bundle of shape space [6]. While the Sasaki metric is a natural metric on the tangent bundle, its geodesic computations require time-discrete approximation schemes involving the Riemannian curvature tensor. This not only incurs high computational costs but also impacts numerical stability. We consider a novel approach that overcomes these shortcomings [8]. To this end, we identify tangent vectors of the tangent bundle with vector fields along the geodesic trend. This provides a notion of a canonical metric that is motivated from a functional view of parameterized curves in the shape space. Considering the space of the geodesics as a submanifold in the space of shape trajectories, this allows in particular the use of a naturally induced distance. The corresponding shortest path, logarithmic map and average geodesic, can be computed by variational time-discretization. Remarkably, the underlying energy function allows for fast and simple evaluation, increasing computational efficiency. In particular, it neither requires curvature computation nor decomposition in horizontal and vertical components.

Based on the derived metric for geodesic trends, we obtain a notion of mean, covariance, and Mahalanobis distance. This allows us to perform a statistical hypothesis test for comparing group-wise mean trends. More precisely, we can test for significant differences in average shape trajectories using a manifold-valued Hotelling $t^2$ statistic in a non-parametric permutation test setup. This framework allowed us to reveal differences between physio- and pathological shape evolution in a long-term study on incident knee Osteoarthritis [8].

# Current and Future Work

The many advances we have made in recent years in the analysis of shapes and shape trajectories open up new avenues and allow us to address remaining challenges to further extend the scope of the derived methodologies.

## Higher-Order Models

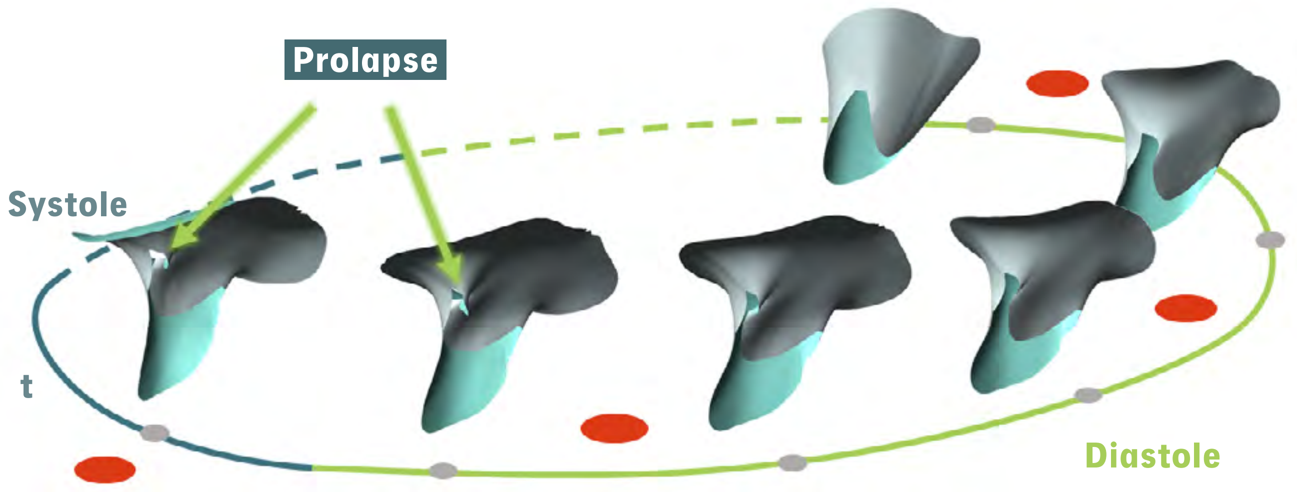


Fig. 5: Spline-based shape trajectory of mitral valve reconstructed from four ultrasound measurements.

Non-monotonous shape changes, e.g. present in time-series of cardiac shape motion or anatomical changes in the human brain over the course of decades, do generally not adhere to constraints of geodesicity. This necessitates the development of higher-order models. To this end, we are generalizing our approaches to Riemannian spline models based on constructive algorithms [4]. Fig. 5 shows a first preliminary result of a spline-based regression of a mitral valve shape trajectory.

# Geometric Deep Learning

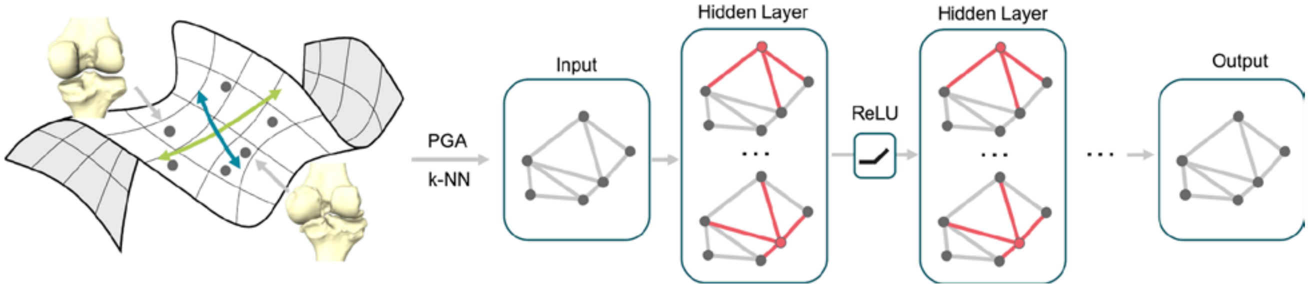


Fig. 6: Schematic depiction of a semi-supervised osteophyte classification approach applying a graph convolutional neural network on a sub-manifold embedded in shape space.

Recent advances in the field of machine learning have led to qualitative breakthroughs on a wide variety of tasks. Consequently there is a high interest of translating learning approaches to the high-dimensional and non-Euclidean setting of shapes and time-series thereof. In [10] we derived a transductive learning approach for morphometric osteophyte grading based on geometric deep learning (see Fig. 6). We formulate the grading task as semi-supervised node classification problem on a graph embedded in shape space. To account for the high-dimensionality and nonlinear structure of shape space we employ a combination of an intrinsic dimension reduction together with a graph convolutional neural network. Based on this, we plan to derive approaches for the analysis of cardiac shape trajectories to improve decision support and therapy planning.

## Multiple Co-Varying parameter

One line of future work will be to extend the Riemannian mixed-effects framework to include multiple explanatory variables. This will provide the means to study shape trends that correlate to multiple co-varying parameters (e.g. archaeological changes may also correlate with excavation site and used stone material).

## Man-Made Objects and Beyond

A largely unaddressed topic is the extension of shape analysis to other types of data. For example, the geometrical artifacts studied within archaeology typically belong to the class of man-made objects and, thus, feature a rich geometrical, structural, and topological variability. This violates the assumption that instances of an object class can be modeled as deformation of a template shape. A promising direction for the analysis of objects with such variability is to adapt concepts from co-analysis of shape collections and topology-varying correspondence estimation.

Ralf Borndörfer +49 30 84185-243 | borndoerfer@zib.de

# ON THE ROAD TO AUTONOMOUS OPERATING ROOM SCHEDULING

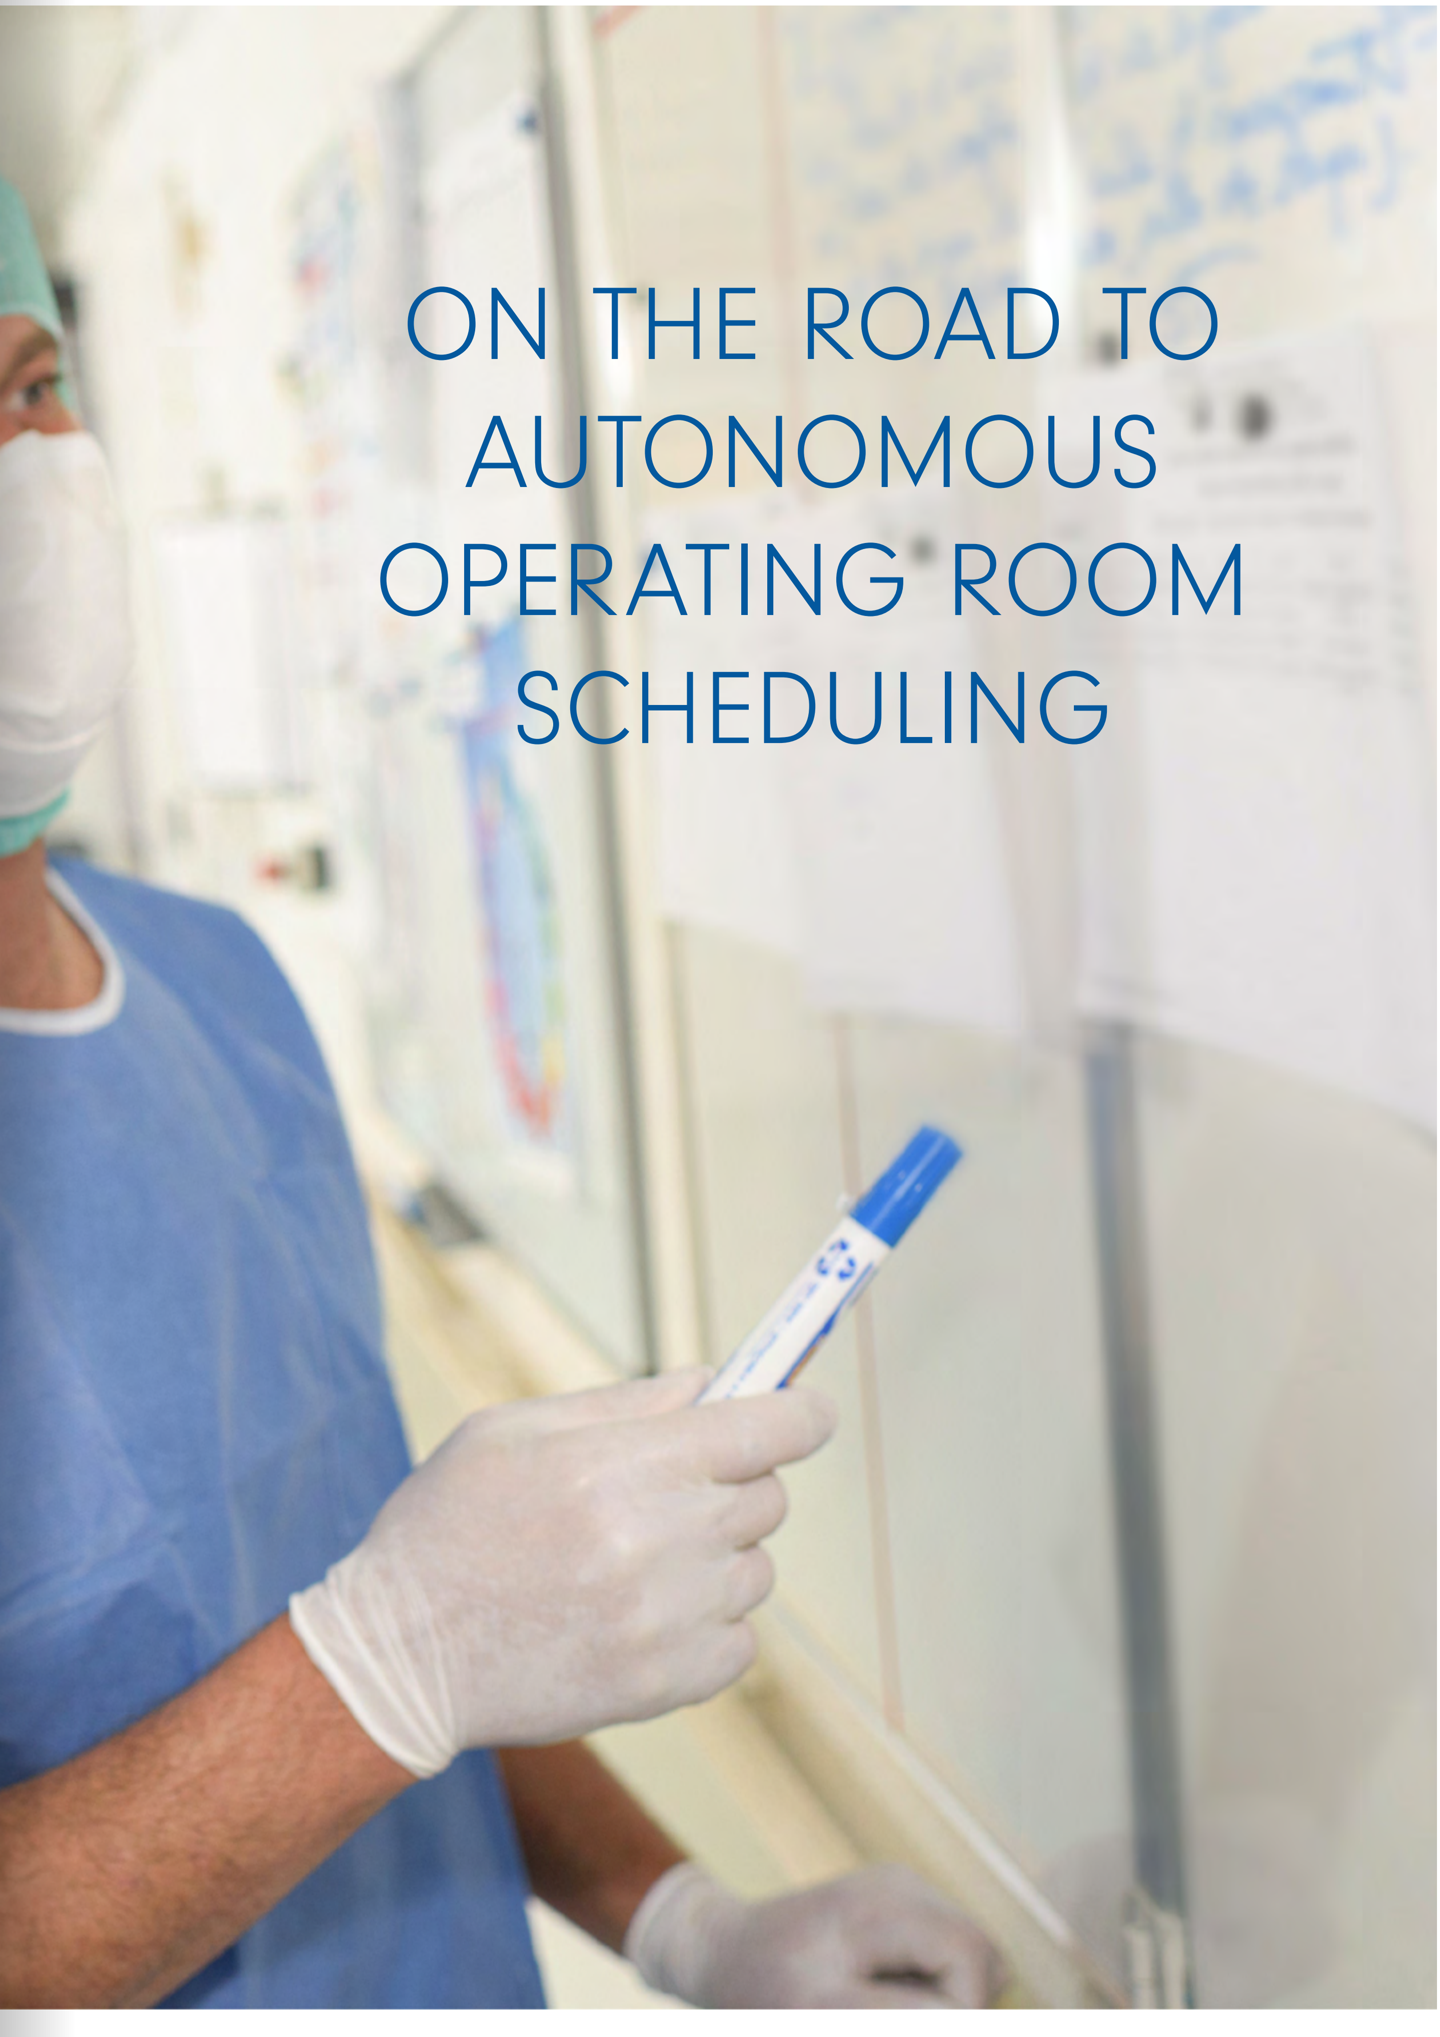

# Scarce Resources in Hospitals

**Since 2005, the number of annual operations in Germany has increased by approximately 25%.** However, the available resource capacities of the hospitals such as operating rooms and staff remained almost equal or even decreased. In 2019 alone, there were about 17,000 open nurse positions in Germany. Some hospitals even had to shut down some of their operating rooms because they could not find sufficient numbers of qualified personnel. Other hospitals suffered from a lack of operating rooms such that they had to implement middle and late shifts to cope with the demand. Overall, today's hospitals have to focus much more on efficient resource management to deal with potential bottlenecks. At the same time, they have to work more cost-efficiently because a large share of a hospital's profits and expenses come from surgeries. Since all surgeries share scarce resources such as operating rooms and staff, they become increasingly interdependent on each other. These dependencies must be planned giving them a high degree of attention. At the moment, the planning complexity can only be coped with by more human planners and a huge communication effort.

## IBOSS Project

In order to improve the quality of the daily operating room schedules, ZIB started a BMBF-funded project "Information-Based Optimization of Surgery Schedules" (IBOSS) together with Freie Universität Berlin and the University of Paderborn in cooperation with the Charité Berlin, which is one of the world's most prestigious hospitals. The goal of the project is to develop mathematical optimization tools that can compute daily operating room schedules of the highest quality.

There are many different indicators for the quality of an operating room schedule. The main objective for the hospital is to use the available resources as efficiently as possible. That means, in particular, minimizing overtime, idle time, delays/cancellations of surgeries, waiting time, and incorporating planning stability such that the initially planned schedule is robust enough to cope with unforeseeable events during the planning day, such as emergencies.

In particular, the different objectives are motivated by different parties such as the patient, staff, or management. Hence, it may be the case that objectives contradict each other, for example, if one has to decide whether a surgery is still performed one hour before the regular shift ends, in good knowledge that the surgery will presumably take two hours. Either the staff or the patient is going to be affected negatively. Thus, one has to weigh the different objectives to generate schedules that implement the best compromises between all the objectives.

The IBOSS Logo

## Daily Struggle with Constraints

Around afternoon, the baseline schedule for the next planning day has to be finished. Normally, minor changes will be applied until the next morning but the baseline remains fixed. Hence, the first planning step is to create a baseline schedule. The baseline schedule already has to satisfy many side constraints. It is dictated by the OR capacity plan, which is a week-based plan that shows for how long which department is booked into which operating room. Departments describe different medical disciplines such as general surgery, gynecology, urology, or orthopedics, each of which has its own qualified employees. The OR capacity plan is valid for one year, after which adjustments are made according to current demands. The baseline schedule consists of the assignment and sequencing of surgeries to the available resources. Resources are, for example, operating rooms, surgeons, anesthetists, nurses, and medical devices.

Every surgery generates its own supply chain in the hospital system, which consists of different subphases such as station pickup, OR preparation, induction, surgery, OR cleanup, recovery, and possibly ICU. Every phase has an individual resource demand and some phases run in parallel. For example, the operating room is prepared by the nurses while in parallel the patient receives anesthesia by an anesthetist and a supporting nurse. For the surgery phase, one surgeon must be present. Often, a second surgeon will attend the surgery only during the main operation so that sometimes he/she can switch smoothly between two operating rooms.

A surgery usually requires one to two surgeons, two nurses, one anesthetist, and one anesthetist nurse. For complicated surgeries, additional staff may be required.

One goal is to minimize the delays between the subphases since every minute that exceeds the OR opening hours is considered as overtime. This is because the staff rosters are built based on the planned OR capacity. For each staff type (surgeons, anesthetists, nurses), the schedule is covered by individual shifts of different quantities, e.g. 7 a.m.–3 p.m. (6×), 10 a.m.–6 p.m. (2×), and 12–8 p.m. (2×).

Additionally, there are also individual scheduling constraints that make the construction of a plan by hand really complex and hard to grasp. For example:

- Patient p is first available at 11 a.m.
- Surgeon s is not available from 10 a.m.–12 p.m.
- Sudden absence (illness etc.) of staff
- Integration of emergencies into the running process (immediately or within 6, 12, or 24 hours) with available resources.

Sophisticated surgeries additionally require at least one very experienced surgeon so that different experience levels may affect the schedule. For each department, this requires careful planning to guarantee the best treatment for each patient. The additional planning effort comes on top of the work of the department heads, who are often surgeons or physicians themselves.

Similarly, nurses are usually qualified for one specific department. For example, it is desired that a surgery is only attended by nurses of the associated department. However, if there is a deficiency of nurses in one department, it is sometimes allowed to borrow a nurse from another department in order to finally perform the surgery. Again, this induces dependencies across the departments that are difficult to manage manually, since the departments are mainly operated separately. Furthermore, in contrast to surgeons, a nurse can be replaced by another nurse during surgery to get a mandatory break or to finish their shift on time. However, even recognizing such potential resource availabilities is a hard task because most digital real-time information is poorly formatted so that the planner cannot use it efficiently. Hence, much of the information flow still relies on direct communication.

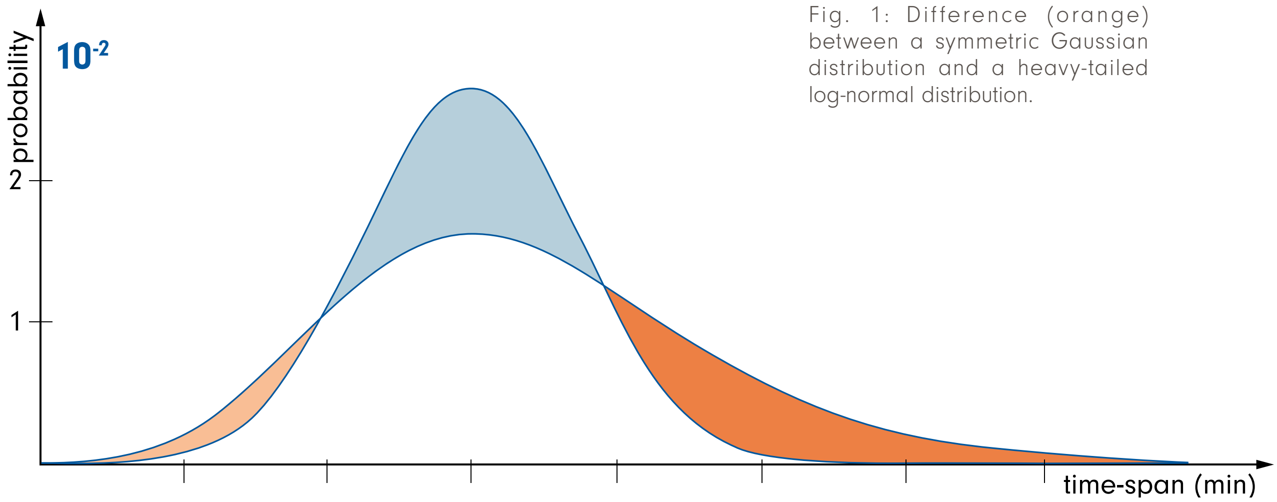


Fig. 1: Difference (orange) between a symmetric Gaussian distribution and a heavy-tailed log-normal distribution.

## Static vs. Dynamic Scheduling

In general, it is desired to perform the baseline schedule as planned. In practice, however, the baseline schedule will almost never be realized. This is due to the high uncertainty in the surgery durations that can sometimes vary by several hours. If a surgery is delayed or even cancelled, then replanning is necessary to keep the schedule cost-efficient. Hence, instead of considering the static cost of the baseline schedule, we want to optimize according to the dynamic cost of the baseline schedule that includes the estimated replanning decisions. Thus, we want to compute an optimal decision policy that minimizes the total expected cost with respect to the static baseline schedule as starting point.

In order to quantify the uncertainty in the system, we compute probability distributions for the duration of the subprocesses of each surgery. In general, a surgery consists of multiple smaller operations, each of which has a unique representation in form of a so-called OPS code that shows exactly which operation is being performed. For this, we developed a machine-learning model that predicts a probability distribution for surgery based on a given set of OPS codes. Our database consists of six years of historical data that contains about 140,000 surgeries. The model uses a cross-entropy objective to maximize the likelihood that the estimated distribution matches the observed durations. It turned out that most surgeries do not follow a Gaussian distribution but rather a long-tail distribution, such as a log-normal or log-logistic distribution (see Fig. 1 [3]). This seems plausible, since a surgery is more likely to take longer than usual.

# Algorithmic Approach

In order to compute a baseline schedule that satisfies all the different constraints, we formulate the underlying scheduling problem as a mathematical program that is solved by a branch-and-bound algorithm based on constraint-propagation and simulation techniques. It turns out to be extremely difficult to provably compute the optimal solution. In fact, the simultaneous assignment and sequencing character of the scheduling problem is solvable only for very small instances in the proposed time frame of five minutes, even in the deterministic setting. Therefore, we combine our branch-and-bound algorithm with a primal search algorithm that heuristically cuts off unpromising branches and diversifies the search regions. The built-in constraint propagators take care of the feasibility of the solutions. Stochastic cost bounds are determined by internal simulation. Our method is able to quickly generate schedules that can be used in practice. In addition, our algorithm works online, meaning that it suggests rescheduling decisions if the current schedule turns out to be inefficient due to delays, cancellations, or no-shows. The mere time effort for the generation of a schedule is estimated to drop by at least 80% if measured in working hours. Moreover, overtime is estimated to be lowered by 17% by a better a priori evaluation of surgery durations and scheduling decisions. Our next goal is to deploy our algorithm for practical usage. However, it is a long way from an algorithm prototype to an algorithm that is used in practice by default. Data interfaces, visualizations, and UIs must be provided to run the algorithm with the correct parameters (see Fig. 2). Moreover, it is important to give the practitioners a basic understanding of what the algorithms are doing. The acceptance of the practitioners is a factor that should not be underestimated, since nobody will use even the best algorithm if they do not comprehend its benefits.

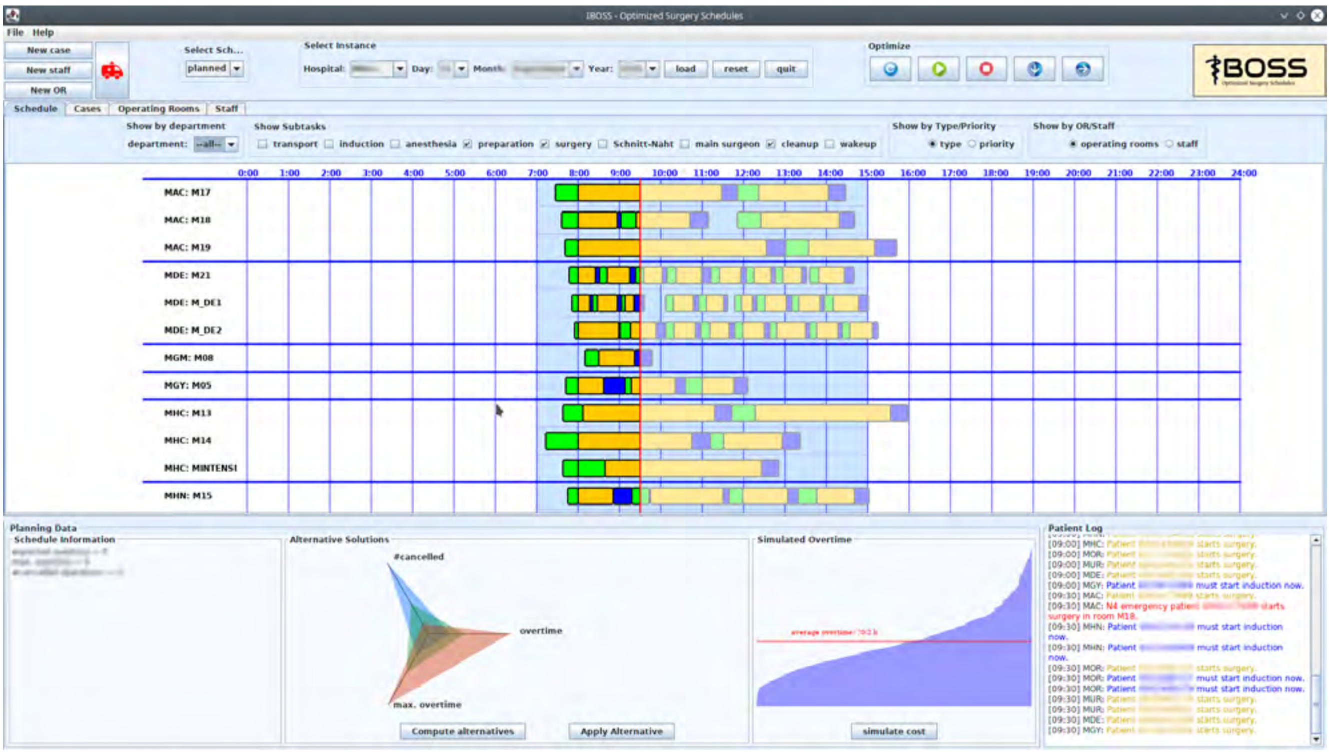


Fig. 2:Prototype of the IBOSS surgery schedule planner.

## Vision of Full Autonomy

The developed algorithm makes replanning suggestions automatically so that, in the best case, the planner only has to confirm the decision made by the algorithm. This almost gives the impression of a fully autonomous system that can handle the surgery scheduling by itself – like AI. Perspectively, how realistic is the fully autonomous artificial OR planner?

Well, currently there are still many obstacles between now and that vision. First of all, current surgery planning is still far from being completely digitized. It is still based on a lot of human interaction and communication because the planning is highly decentralized. The rostering and daily scheduling for physicians, surgeons, nurses and anesthetists is mostly done by separate organizational units, but which all belong to the same supply chain of a surgery. This is historically grown but there are also few possibilities to do it differently since the manual organizational effort is already very high for all the planners. Hence, moving from decentralized to centralized planning requires the will of the hospital staff and management to be implemented practically.

A second aspect is the planning responsibility, which is a bigger hurdle for full autonomy. Similar to autonomous driving, one always needs a person that is in charge of the current events. In particular in hospitals we need to be able to act fully autonomously as a human where bottleneck decision may be the difference between life and death. Hence, it is questionable if there will ever be a fully autonomous operating room scheduling system. It is much more likely that a semiautonomous system will be implemented that will support the planner with all the necessary information and tools to run the daily operative business as efficient as possible.

## (Re-)focus on the Patient

An algorithmic decision support tool yields obvious benefits to save time and money for the planners and the hospital management. However, it remains for how the patient would ultimately benefit from that change to be discussed, since all the processes should be centered around the best possible treatment for the patient. We believe the benefits for the patients are manifold.

An algorithmic tool is capable of making objective decisions based on the given real-time data. There is probably nothing more annoying for a patient than traveling to the hospital for surgery, waiting for several hours, and finally getting a notification that the surgery has been cancelled. Prescriptive analytical tools have all the information available: the status in the operating rooms, the expected waiting time of each patient, and the likelihood of cancellation, even hours before the actual appointment. If a critical value is exceeded, the patient may get an earlier notification of cancellation or the surgery can still be performed by an improved resource management.

Furthermore, a smaller number of planners lowers the danger of miscommunication and potential conflicts between staff in the operating rooms. The less personal communication is necessary, the more the staff can focus on their work. Current polls show that about 50% of the OR staff feels more stressed because of personal conflicts at their workplace, which they consider as a serious risk for the patient. Less necessary communication avoids unnecessary personal conflicts. This has immediate consequences on the treatment quality of the patient, since the whole working atmosphere in the operating room can be improved by using a computational support tool in the background. This will probably lead to best improvement in the treatment quality of a patient, since the most important decisions in the operating room are still made by human beings.

Thorsten Koch | +49 30 84185-213 | koch@zib.de

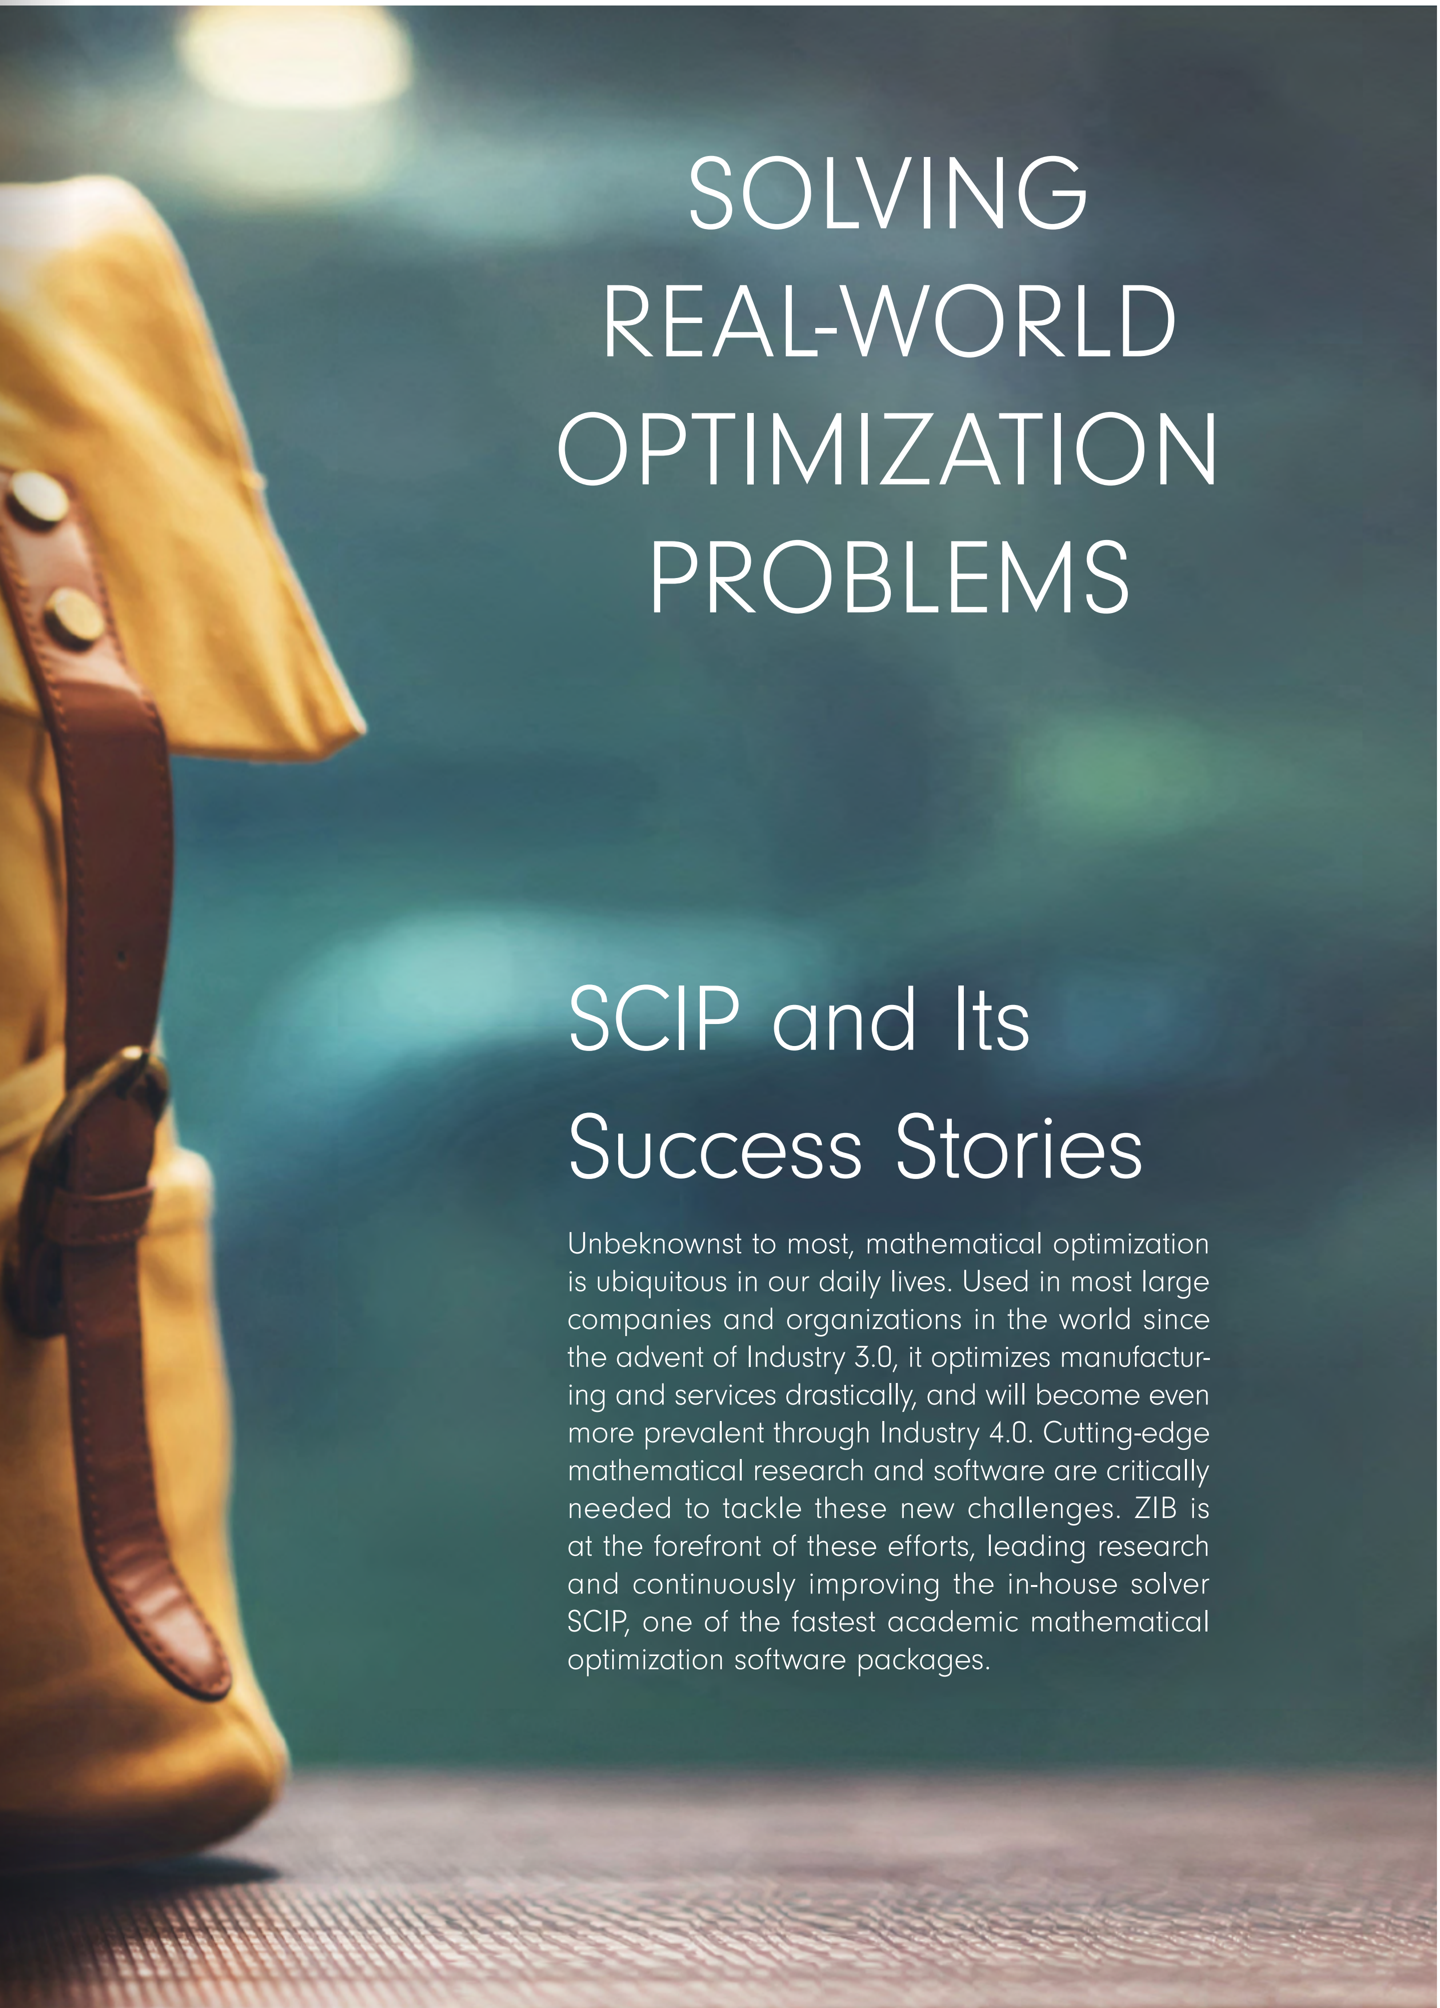

# SOLVING REAL-WORLD OPTIMIZATION PROBLEMS

## SCIP and Its Success Stories

Unbeknownst to most, mathematical optimization is ubiquitous in our daily lives. Used in most large companies and organizations in the world since the advent of Industry 3.0, it optimizes manufacturing and services drastically, and will become even more prevalent through Industry 4.0. Cutting-edge mathematical research and software are critically needed to tackle these new challenges. ZIB is at the forefront of these efforts, leading research and continuously improving the in-house solver SCIP, one of the fastest academic mathematical optimization software packages.

# Mathematical Optimization

Determining the best course of action in a given situation is a fundamental task performed by humans on a daily basis. In the modern world, it often involves complex, large-scale decisions: What is the best use of limited production resources? How to construct a schedule that achieves maximal efficiency while avoiding conflicts? Which connections should be added when expanding a power network so that the costs and environmental impact are minimized? These decisions may involve thousands and even millions of variables, making computational tools necessary in order to obtain high-quality solutions. Providing such tools is the aim of mathematical optimization, which is the science of finding the best choice under given restrictions.

A mathematical optimization problem is described in terms of variables, which represent decisions and their consequences, and parameters, which reflect the knowledge about the system to be optimized. The goal is to find the maximum or minimum of an objective function (in some cases, several objective functions) while satisfying constraints. Some common types of constraints are equations, inequalities, logical relations, variable domains, and requirements that certain variables have only integer values.

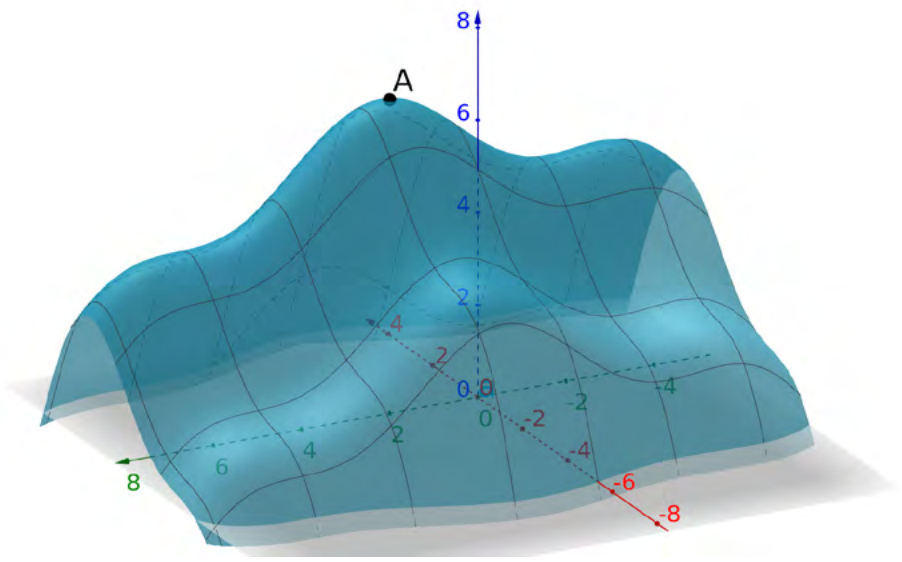


Graph of a nonlinear function. The global maximum is achieved at point A (image created in GeoGebra).

# The Knapsack Problem

One well-known optimization problem is the knapsack problem, which has numerous applications such as portfolio optimization, warehouse space management, and raw material cutting, to name a few. The name is derived from an everyday situation where a traveler must decide which items to pack into his or her knapsack by considering their usefulness and weight.

More formally, a set of indivisible items is given, where each item has a given weight and value. The combined weight of the chosen items is constrained to be below a given weight limit. Each item has a variable associated to it, encoding the decision of whether to choose this item (variable has value 1 in the solution) or not (value 0). No fractional solutions are allowed, since the items are indivisible. The objective is to maximize the combined value of the chosen items. A slightly different version of this problem considers types of items, and the variables represent the numbers of items chosen from each type.

This is a combinatorial optimization problem because it deals with a finite set of objects. The main difficulty here lies in the fact that the number of possible combinations that comprise a feasible solution grows exponentially in the size of the problem.

Combinatorial optimization is one of several interesting and challenging areas of optimization. Other common problem classes include mixed-integer linear programs (MIPs) that contain both discrete and continuous variables, nonlinear programs (NLPs) where some constraints are described with the use of nonlinear functions, and the combination of these two classes, mixed-integer nonlinear programs (MINLPs).




An illustration of the knapsack problem. The objective is to choose the best selection of items given weight restrictions and item values.

# Solving Constraint Integer Programs

Given the practical relevance of optimization problems, researchers at ZIB have developed the general-purpose solver SCIP to solve large-scale problems to global optimality. SCIP is a framework for constraint integer programming oriented towards the needs of mathematical programming experts who want to have total control of the solution process and access detailed information down to the guts of the solver. SCIP can not only be used to solve difficult optimization problems, but also serves the optimization community to implement, test, and evaluate their algorithmic ideas in a state-of-the-art solver. Indeed, with over 14,000 downloads per year from over 100 countries, SCIP is one of the most used research frameworks for branch-cut-and-price algorithms worldwide – also because SCIP is currently one of the fastest MIP and MINLP solvers that is openly accessible in source code. Additionally, starting from version 3.2, SCIP's main new developments and features are documented in release reports [1, 2, 3, 4].

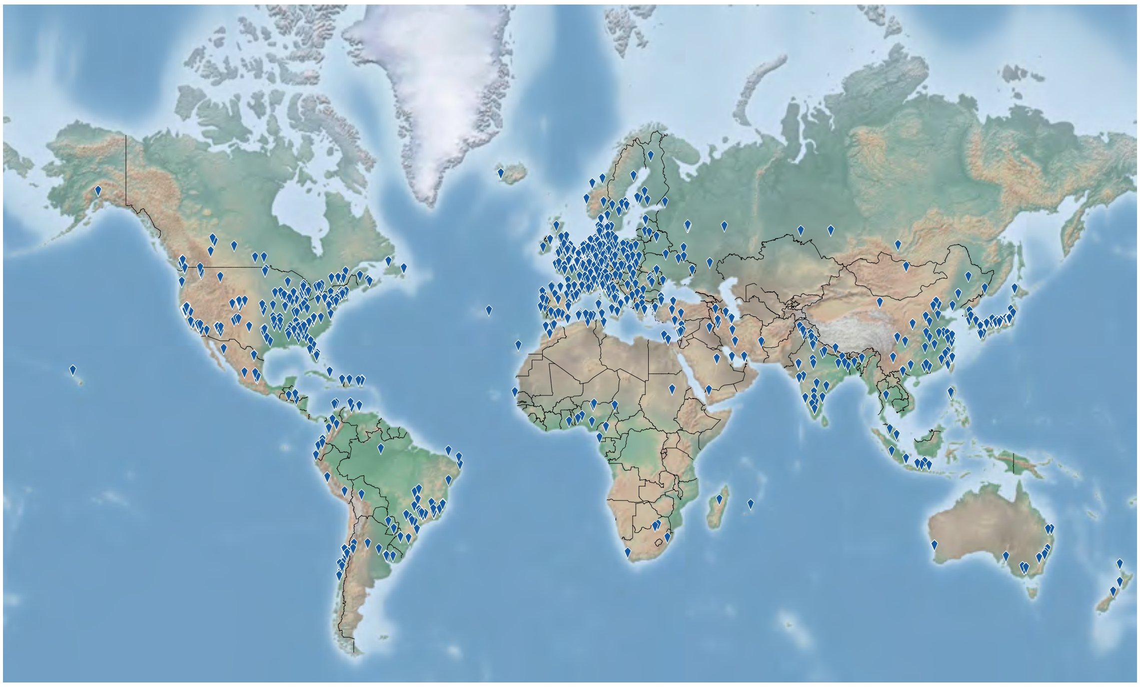

A world map showing all locations of registered SCIP downloads.

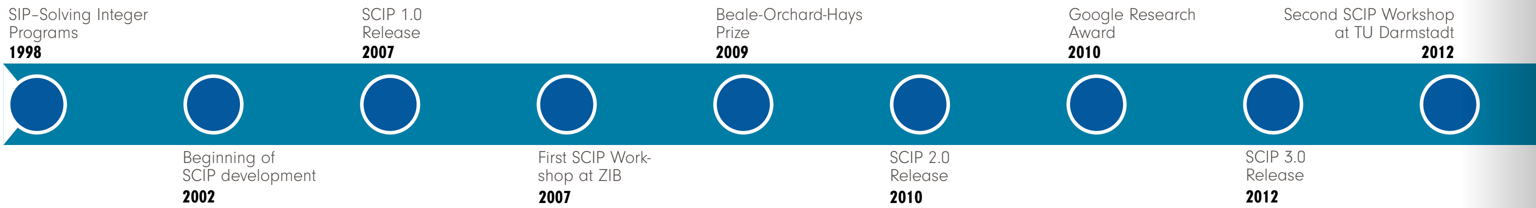


The history of SCIP's development over the last two decades.

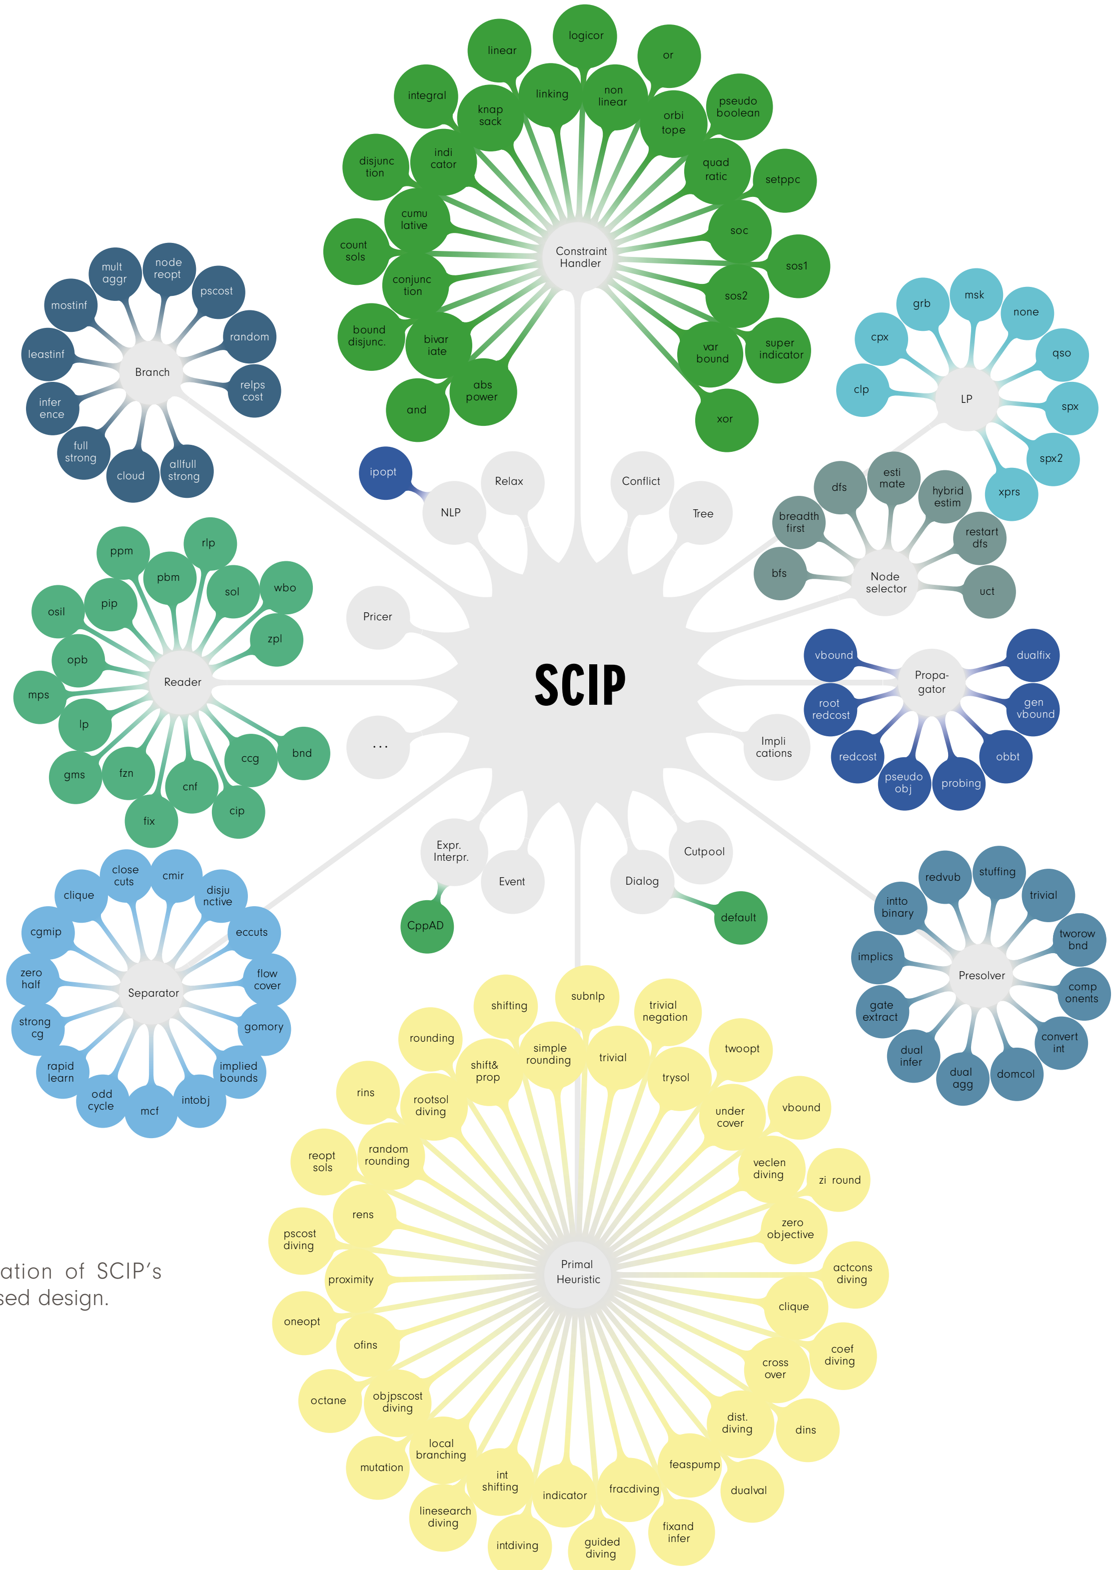


A visualization of SCIP's plug-in-based design.

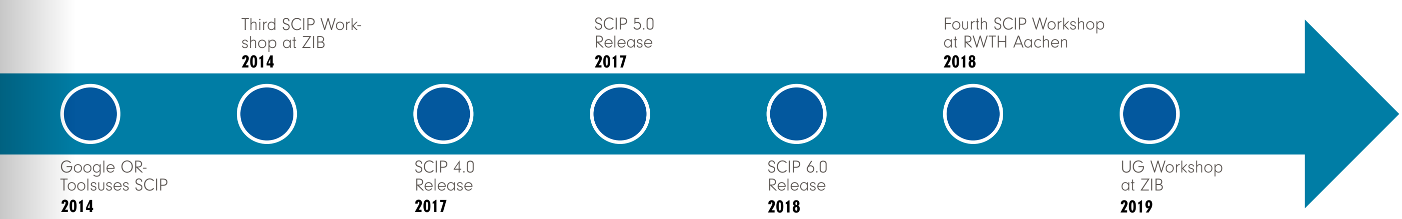

## Interfaces

A natural way of formulating an optimization problem is to use a modeling language. Besides the modeling language ZIMPL, which is developed at ZIB, there are several other modeling tools with a direct interface to SCIP. These include Comet, a modeling language for constraint programming, AMPL and GAMS, which are well-suited for modeling mixed-integer linear and nonlinear optimization problems, and CMPL for mixed-integer linear problems. Additionally, SCIP provides interfaces to several programming languages that are developed and maintained on GitHub in order to provide extensions and patches faster and to collaborate on them more easily. Besides the popular interfaces for Python and Java, there is also an interface for Julia available.

## Development and Quality Control

With more than 500,000 lines of code, the source code of SCIP is not trivially maintainable. To facilitate and organize the development, Git and GitLab are used as a source code management system. The code is organized in branches on which bugs are fixed or features are developed that are then merged to the main master and bug-fix branches. Git decentrally and hierarchically preserves the history of code changes. GitLab provides a platform for discussion, including description of bugs and issues. Ensuring code quality is of the highest importance; for this we use a Jenkins server that regularly runs our test suite on different platforms and environments. Additionally, before applying a code change to one of the main branches, a short series of checks is automatically executed and results are reported back to the developer. Finally, the continuous improvement of the performance of SCIP is ensured by regular performance runs. These consist of the execution of SCIP on the compute cluster on a set of problem instances. The log files are scanned for different measures, for example speed and solution quality. Since the evaluation of this data can be a difficult task, a custom software tool, Rubberband/IPET, is used to store the log files, collect the data, and facilitate and standardize the analysis. The components interact with the developers as well as with each other on demand and combine to form an elaborate system for quality control.

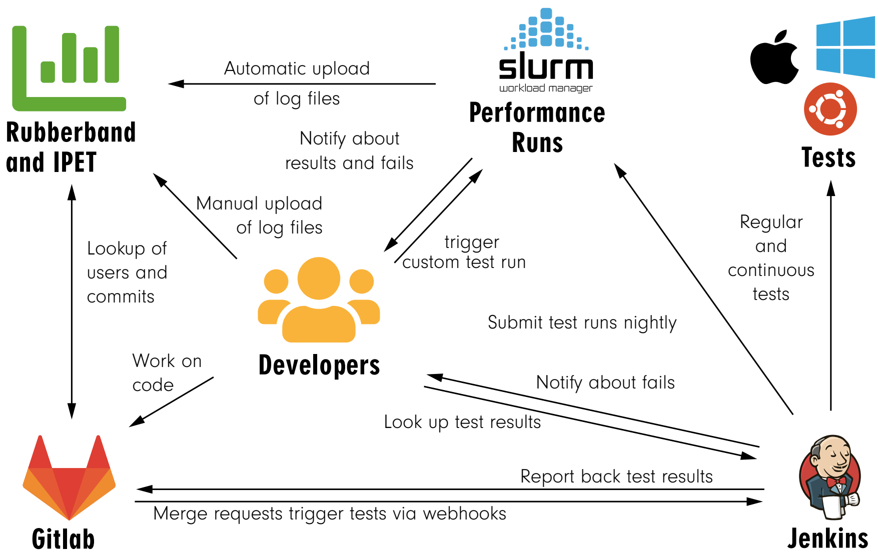


Interplay of components for developing and testing of SCIP.

## CAREERS

- Ten awards for master's and Ph.D. theses at MOS, EURO, GOR, and DMV
- Nine former developers are now building commercial optimization software at CPLEX, FICO Xpress, Gurobi, MOSEK, and GAMS

## RECENT AND PLANNED SCIP WORKSHOPS

**UG workshop**

Workshop on parallel algorithms in tree search and mathematical optimization:

- 15 speakers
- 46 registered attendees
- 16 talks
- Two programming sessions

**Google × SCIP workshop**

Purpose: For Google engineers to better understand how SCIP works, for SCIP developers to better understand how Google uses MIP, and for both parties to recognize each other's priorities and identify opportunities for future collaboration.

- 11 lectures with discussions

**Combinatorial optimization at work 2020**

- Two-week summer school to be held online on September 14–26, 2020
- 27 confirmed speakers at the moment

**Online SCIP workshop 2020**
To be hosted by the University of Exeter on June 3–4, 2020

## The SCIP Optimization Suite

# The Power of Collaboration

SCIP is the main component of a toolbox for generating and solving mixed-integer nonlinear programs. The SCIP Optimization Suite consists of SCIP, SoPlex, ZIMPL, UG, and GCG. The SCIP Optimization Suite is the result of a continuous collaboration between researchers located at ZIB within the context of SynLab, one of the labs of the Research Campus MODAL, at research centers across Germany, such as TU Darmstadt, FAU Erlangen-Nürnberg, RWTH Aachen, as well as other countries in Europe such as the UK and the Netherlands. Overall, the development team consists of 11 postdocs and professors, 13 Ph.D. students, and five student assistants.

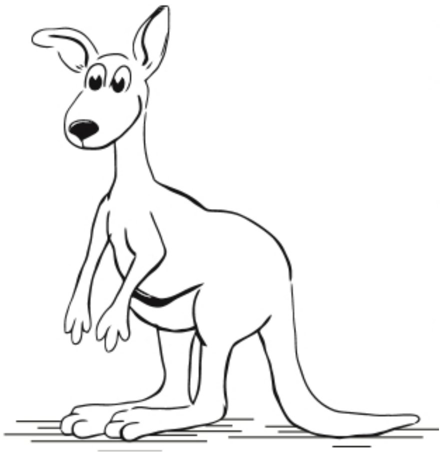

The mascot of the SCIP Optimization Suite.

## SUCCESS STORIES AND APPLICATIONS

Large-scale mathematical optimization problems arise in various projects developed at ZIB. Here, we present a selection of projects for which SCIP was critical for their success.

### ZIMPL

Zuse Institute Mathematical Programming Language (ZIMPL) is an algebraic modeling language. ZIMPL supports the creation of linear as well as nonlinear mixed-integer mathematical programs.

https://zimpl.zib.de/

### SOPLEX

Sequential object-oriented simPlex (SoPlex) is an efficient simplex-based linear programming solver. SoPlex allows linear programs to be solved exactly.

https://soplex.zib.de/

### GCG

Generic Column Generation (GCG) is a generic branch-cut-and-price solver for mixed-integer linear programs. Based on SCIP, GCG solves MILPs using Dantzig-Wolfe decomposition. The decomposition is based on a structure provided by the user or detected by GCG itself.

https://gcg.or.rwth-aachen.de/

### UG

Ubiquity Generator (UG) is a framework for parallelization of branch-and-bound solvers in a distributed or shared memory computing environment. UG has been used to create massive parallel solvers by parallelizing SCIP, CPLEX, and Xpress. ParaSCIP solved 21 open instances from different MIPLIBs. ParaXpress solved three open instances from MIPLIB2010 and one from MIPLIB2017.

https://ug.zib.de/

# The SAP Collaboration

Supply chain management (SCM) deals with the optimization of a company's value chain. This begins with the choice of suppliers and extends to the delivery of produced goods to the customer. In this process, decisions are made about purchasing, transport, storage, and further processing of the goods. The value chain is depicted symbolically in the figure below.

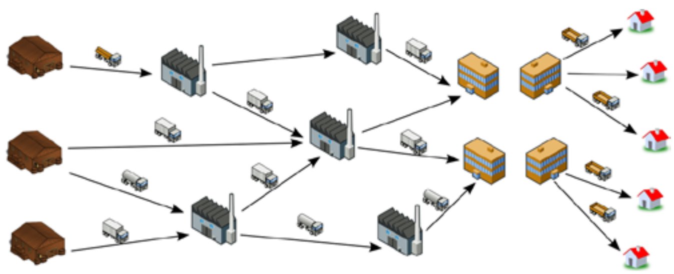

A visualization of a supply chain management instance.

Production planning plays a central role here: based on customer orders or the forecasted demand, it decides when and where which products are produced. In doing so, it ensures that the materials required for production are available in sufficient quantities at the respective location, takes into account capacity restrictions, such as the number of available workers, and attempts to reduce storage costs by producing as late as possible.

The integrated planning of the entire value creation process contains a high potential for optimization and is therefore of great importance for the profitability of companies. Our industry partner SAP is developing software that enables companies to model their value chain in detail. This includes precise descriptions of production models, capacity restrictions, as well as costs of production, storage, transport, and delivery. In addition, fixed incoming stock, customer orders, and demand forecasts are specified for different time periods (for example, days or weeks). A mixed-integer linear optimization problem (MIP) is then generated from this data to optimize the supply chain.

The SAP project deals with the fast solution of especially numerically demanding and large MIP instances that arise from such supply chain management problems. Thereby the primary goal is not the optimal solution, but rather the best possible solution within fixed time limits. The SCIP Optimization Suite, which was extended with new components added during the course of the project, is used to solve the MIP instances.

The complexity and the large size of the problems motivated several SCIP methodological improvements [5]. For example, several problems exhibit underlying structures that could allow for more efficient, decomposition-based solution algorithms. A special decomposition API is thus being developed within SCIP in order to read user decompositions in a more flexible way and allow for better structure exploitation in the developed algorithms.

# Energy Grid Berlin-Adlershof

As part of the energy strategy Berlin-Adlershof 2020, the high-tech location Adlershof has itself set the target of reducing its primary energy demand by 30 percent by 2020. As part of various projects, several energy-related concepts and measures have been implemented at the Adlershof campus. The project "Energienetz Berlin Adlershof," implemented as a joint project of the three partners TU Berlin (TUB), Siemens AG (in collaboration with ZIB), and Hochschule für Technik und Wirtschaft Berlin (HTW Berlin), made a significant contribution to achieving this goal. The project aimed at improving energy efficiency at the property and district level as well as creating planning bases for the efficient energy supply of urban neighborhoods. The project demonstrated a networked energy system that included heat, cold, and electricity as forms of energy. Based on this supply system, a cross-media energy management system has been implemented.

In this context, optimization problems need to be solved, which can be formulated mathematically as mixed-integer programs with quadratic constraints, MIQCPs for short. Quadratic constraints become necessary when accurately modeling nonlinear physical processes, in this case, the heat transfer between ice storage, heat storage, and the connecting thermal networks. To be able to treat the projected detail levels in realistic problem dimensions, the performance of current MIQCP algorithms had to be improved. To this end, the MINIP solver SCIP, developed at ZIB, has been further developed to generically exploit emerging structures such as bilinear mixing conditions and time-indexed decision variables. Performance on the given MIQCPs has been greatly improved, and SCIP was successfully integrated into the energy management system, helping to plan and improve the given energy system.

Schematic topology of a smart grid with various energy mediums.

Ice storage at the Adlershof site. (copyright Anja Hanßke).

# plan4res

The general objective of the EU Horizon 2020 project plan4res is to fill the gaps between the increasing complexity of the future energy system planning and operational problems, and the currently available system analysis tools. In a cooperation of seven partners from research and industry, enhanced end-to-end planning and operational tools dealing with technological and market uncertainty, emerging technologies, and increased sector coupling of multienergy vectors, such as heat, cold, and transport, will be assembled. This would result into a synergistic approach to support European system planners, operators, decision-makers, and regulators.

plan4res [6] requires a holistic approach to accomplish its ambitious objective. This approach results in very challenging integrated models that have a high level of detail at European scale with precise spatial and temporal resolutions and that take uncertainty into account. Solving these models requires using decomposition and aggregation methods together with improved solvers. The SCIP Optimization Suite is going to be used by plan4res' planning and operational tools to solve the LPs and MIPs arising from optimization models of energy systems, either separately or as part of a large stochastic optimization methodology.

plan4res

The logo of the plan4res project.

# The GasLab: Gas-Flow Prognosis

In the framework of the Research Campus MODAL, the GasLab addresses the main challenges of natural-gas dispatching in light of the German energy transition (Energiewende) and the crucial role that Germany plays as a transit country in the European cross-border natural-gas supply. The research is conducted in a close partnership with Open Grid Europe (OGE), one of Germany's largest transmission system operators, who is responsible for operating about 12,000 kilometers of natural-gas pipelines in Germany: roughly the same length as the overall German Autobahn network. Through large-scale software, real-time decision support is provided incorporating mathematical models that depict the technical aspects as well as the nontechnical operational measures of the gas transmission network. Additionally, machine-learning techniques are thoroughly investigated, alongside efficient solution algorithms for the respective optimization problems.

An important component in the GasLab is to improve OGE's hourly forecasts for gas demand and supply. The goal is twofold: to predict as precisely as possible the average hourly gas flows for the upcoming gas day and to overcome the diversity in the data characteristics in order to deliver predictions that are appropriate for all different types of entries and exits (e.g. connections to other networks/countries, industrial users, municipal consumers, etc.). A multiregression prognosis methodology is developed for this purpose by training a prediction model using a comparably short historical time. In order to accurately describe the given historical time series, a function with predefined feature terms is formulated by optimizing over its coefficients (weights) to reach the smallest error over the training horizon. Moreover, in order to avoid the potential risk of overfitting, that is computing a function that perfectly represents the past but is very poor in predicting the future, the prognosis method is further extended to include coefficients' selection decisions. This development ensures that at most k coefficients of the prediction function are nonzero, by selecting the most significant ones. Indeed, this higher level of precision comes at the expense of introducing more functional terms and discrete decision variables to the optimization model; thus, solving a large-scale mixed-integer linear program.

SCIP is currently used at the base solver for this vital optimization problem. The models are solved to optimality within an acceptable time frame for the industrial requirements. Furthermore, in case of symmetric solutions, it has been noted that SCIP delivered the solutions that perform best in terms of prediction accuracy for this application. This optimization-based multiple regression method can now be used for predicting the future values of all kinds of time series.

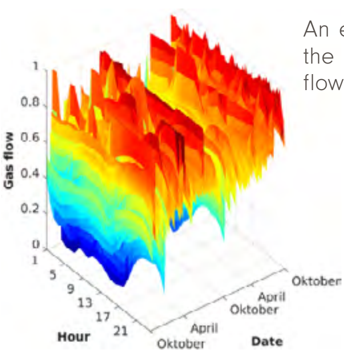


An example of the hourly gas-flow variations.

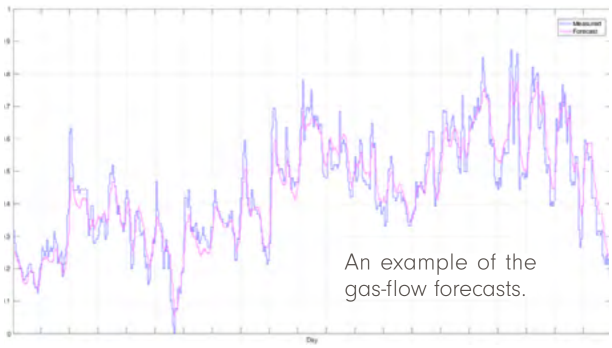

An example of the gas-flow forecasts.

The logo of the transmission system operator.

An example of a natural-gas compression station.
(© Open Grid Europe GmbH)

An example of a gas station.
(© Open Grid Europe GmbH)

# SCIP-JACK – Specializations of SCIP for Hard Combinatorial Optimization Problems

Many optimization problems arising in practice are based on an underlying (physical) network. Consider, for example, the construction of telecommunications networks. Operators are faced with the task of connecting customers and deciding where to route new cables in the most cost-efficient way. Similar scenarios are also found in many other applications. They lead to one of the most studied problems in combinatorial optimization and computer science, the Steiner tree problem in graphs. Mathematically, such problems can be modeled by graphs, abstract representations of "networks." A graph consists of a set of so-called vertices (corresponding to locations) that are connected or can be connected by so-called edges (cables). One can now associate a non-negative weight (distance or cost) with each edge. A subset of the vertices is designated as terminals (customers). A tree – a connected subgraph without any cycles – is called a Steiner tree if it connects all terminals. One is usually interested in computing a Steiner tree of minimum length or cost. This seemingly simple problem called the Steiner tree problem in graphs, or SPG for short, has proven to be notoriously hard to solve, both in theory and in practice. Indeed, simply checking each possible Steiner tree and choosing the best among them becomes quickly intractable; for problems with a few hundred edges, this approach would already take millions of years on the fastest currently available computers.

Moreover, many variations and generalizations of the SPG arise from practical applications. The diverse fields where one encounters such problems encompass computational biology, computer chip design, computer vision, the deployment of drones, and recently also machine learning. To efficiently solve many of these problems, the Steiner tree framework SCIP-Jack has been developed at ZIB as part of the SCIP Optimization Suite. SCIP-Jack can currently handle the classical SPG and 13 related problems. Despite this versatility, SCIP-Jack is the fastest solver for most of the 14 problems it can solve. For example, at the PACE Challenge 2018, dedicated to so-called fixed-parameter tractable SPGs (a subclass of SPGs), SCIP-Jack was the winning solver in one category, and took second and third place in the remaining two. Even though SCIP-Jack (unlike the other participants) does not provide any specialized routines for these subproblems. Steiner tree problems can also be modeled and solved as general MIPs. However, the highly specialized algorithms used in SCIP-Jack perform much better: problems with a few thousand edges can usually be solved within seconds by SCIP-Jack, but take weeks to be solved even by the fastest commercial MIP solvers – or cannot be solved at all. The largest problem solved by SCIP-Jack so far, derived from a cancer application, has more than 64 million edges.

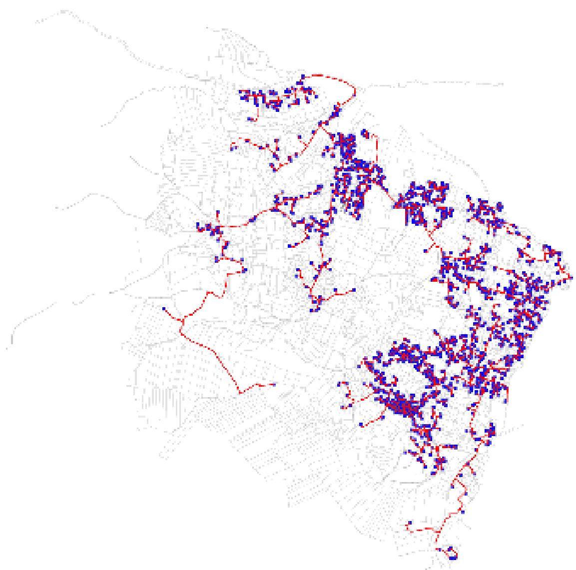

A real-world telecommunications network from an Austrian city with an optimal Steiner tree as a solution marked in red.

# Exact IP/LP

The vast majority of linear and mixed-integer solvers rely on floating-point arithmetic, due to its superior computational capabilities. The occurring small imprecision is controlled by the careful handling of error tolerances. In most real-world applications, it is sufficient to find a result that is no more than one millionth from the exact optimum. However, there exist problems with the requirement of an exact optimal solution, without any numerical imprecision. The exact versions of SCIP and SoPlex achieve this by employing a combination of fast floating-point arithmetic and slower but exact symbolic computations.

Use cases that require exact solutions stretch from research to industry applications that do not allow for any violations of the problem constraints. Exact linear programming was used in molecular biology where multiscale reactions led to problems that could not be handled by floating-point solvers [7]. This sparked the development of the exact solving mode in ZIB's linear solver SoPlex. Exact SCIP was developed and used at ZIB to investigate two famous open conjectures formulated by Péter Frankl [8] and Václav Chvátal [9], respectively. On the industry side, in the design verification of integrated circuit designs, exact solvers are necessary to avoid the accumulation of errors in the numerous different components [10].

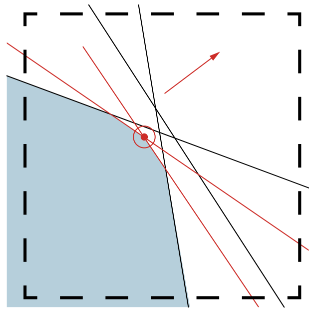

Visualization of the iterative refinement principle used to solve linear programs exactly.

T. Maritima: use case of exact LP in molecular biology (copyright Huber, Hannig, 2004).

Portrait of V. Chvátal, Kyoto, 2007.

# MACHINE LEARNING AND BIG DATA

C. Schütte | +49-30-841-8104 | schuette@zib.de

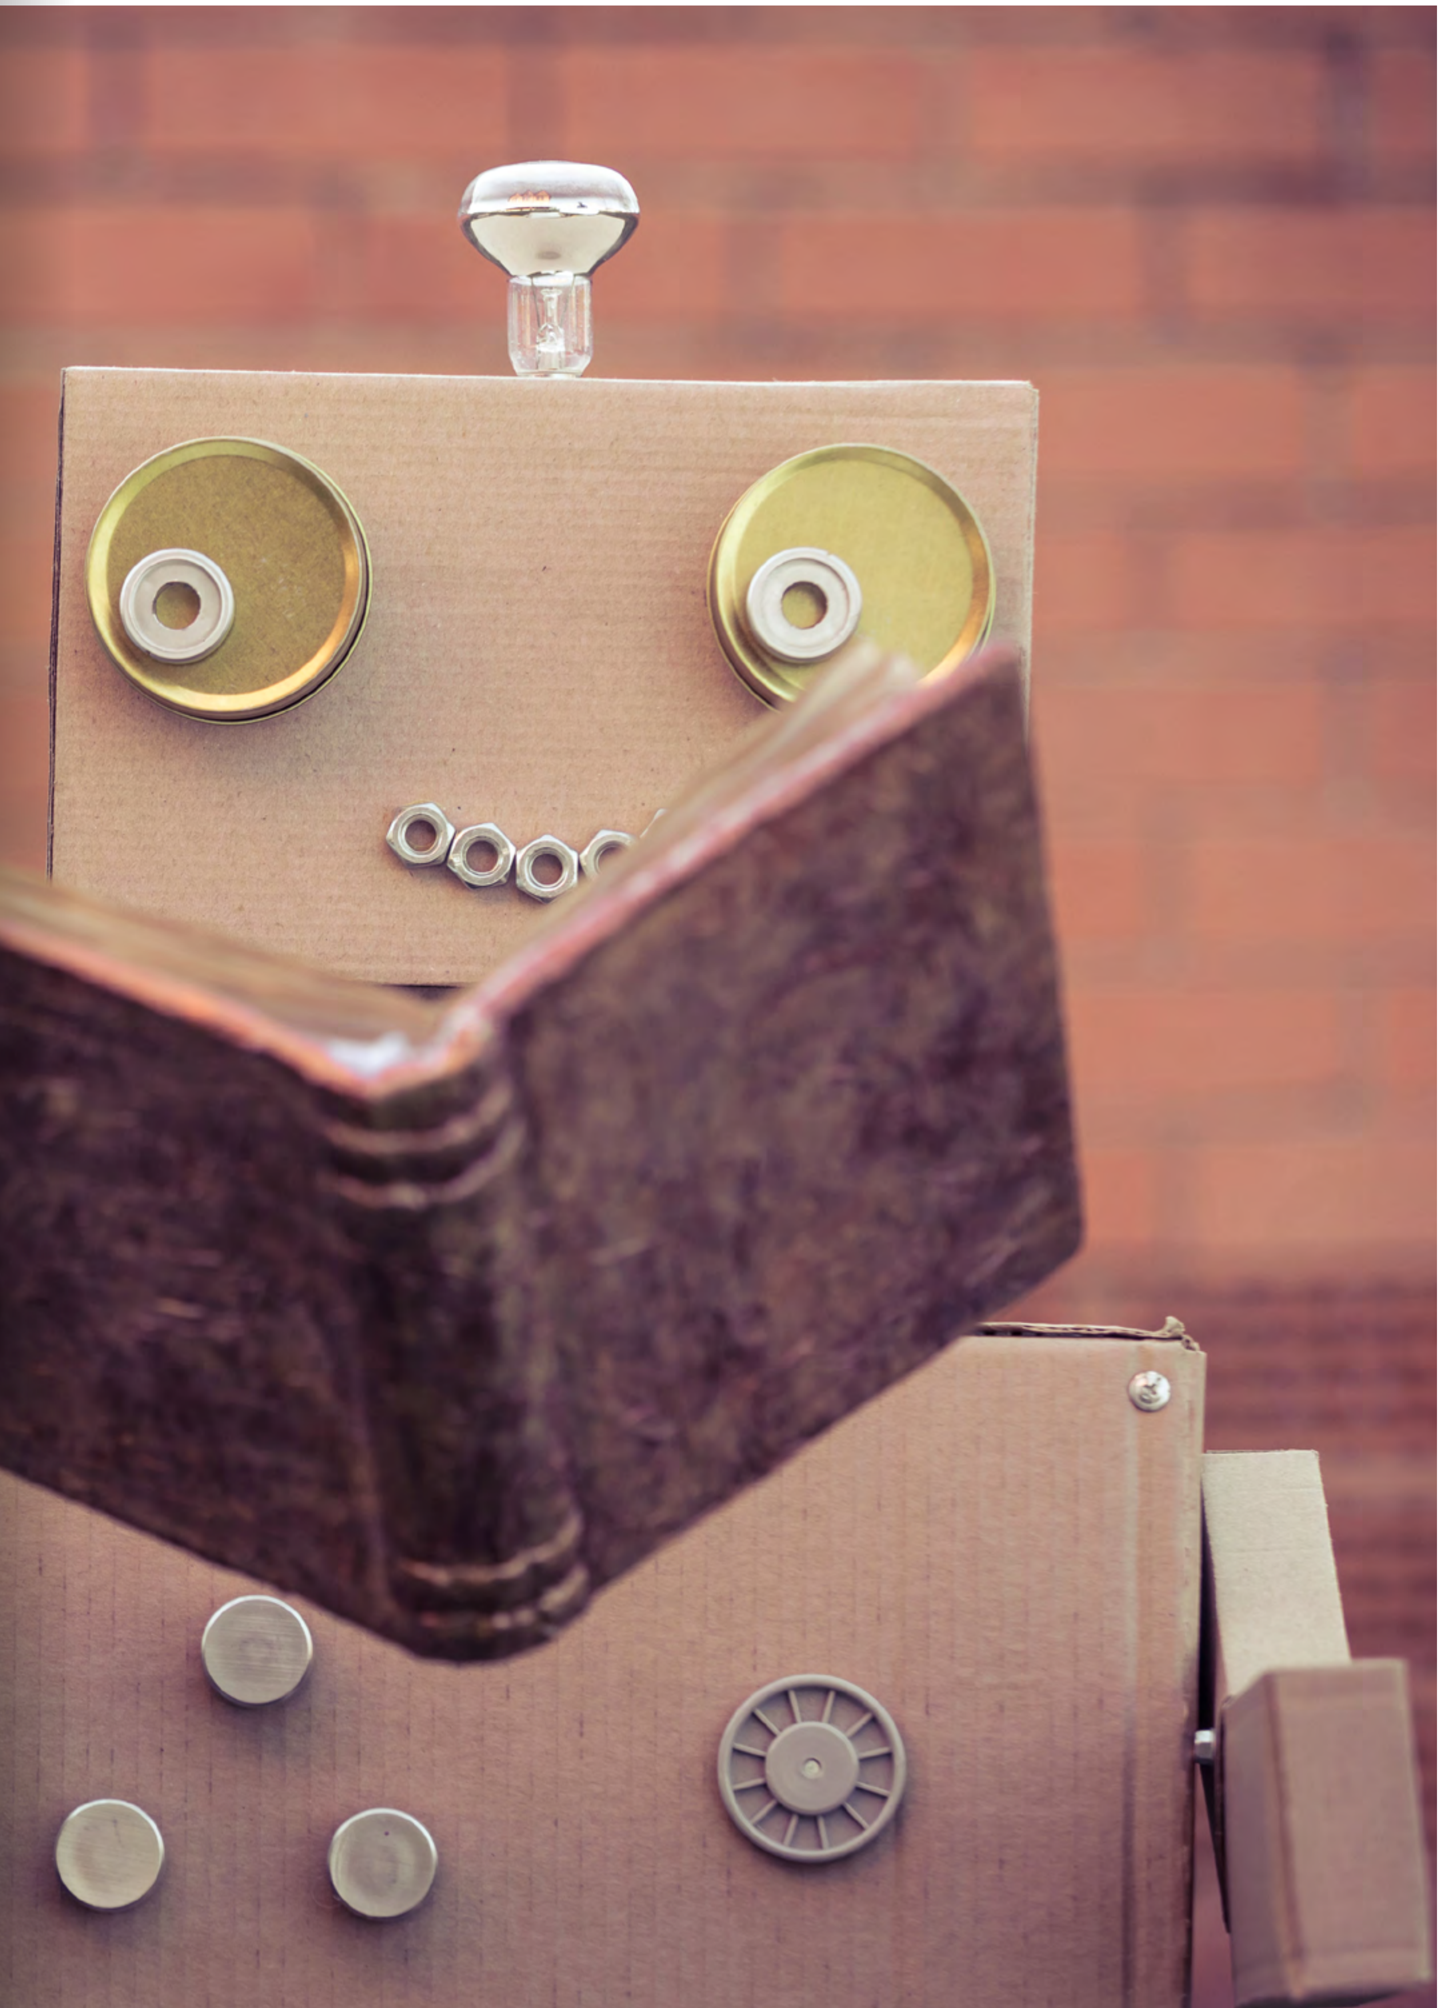

Big data (BD) and machine learning (ML) (together with decision-making) are the key pillars of artificial intelligence (AI). The analysis of very large and heterogeneous amounts of data has the potential to revolutionize many areas of our lives, from the sciences, production, transport, and energy, to political and social processes. However, dealing with BD and ML requires highly specialized knowledge and infrastructure and poses many challenges: it requires computer skills, mathematical skills, and engineering skills. It further demands a societal discussion regarding its ethical implications. BD and ML will disrupt all levels of society and will inspire completely new applications fueling innovation and strengthening the economy. The impact is already evident today: we entered the "fourth paradigm" in the sciences and the economy is talking about the next industrial revolution.



# Research at ZIB

Research in ML and BD has been an integral part of research activities at ZIB for ten years now. With recent success stories and the strong growth of international research activities, ML and BD are becoming a cornerstone of our research strategy. Therefore, in fall 2019, we decided to create a new research department, "AI in Society, Science, and Technology," at ZIB that is strongly integrated with our other departments. In addition, a new Berlin-based Competence Center, the "Berlin Institute for the Foundations of Learning and Data" (BIFOLD), was founded that integrates the former Berlin Center for Machine Learning (BZML) and the Berlin Big Data Center (BBDC) into a new structure that is funded by the German Federal Ministry of Education and Research (BMBF). ZIB, having been part of both, BZML and BBDC, is part of the consortium supporting BIFOLD and will implement new research groups in this context. In the following, we report on four sample research projects on BD and ML that show the breadth of our research activities in these fields and which we are going to strengthen and intensify.

BIFOLD

# Decision Support for Gas Network Operation

Would you feel comfortable driving a car using your rearview mirror and your driving experience only? Open Grid Europe (OGE) is a transmission system operator responsible for the delivery of more than 25% of the German primary energy consumption by operating a natural gas network of comparable length to the German Autobahn. To meet all demands, the 24/7 dispatching center operates more than 100 compressor units, almost 300 regulators, and more than 3,000 valves to control the network. Still, the network is operated mainly based on measured data, which is only available in the "rearview mirror," and the expert knowledge of the dispatcher. To improve this situation and to anticipate and prevent critical situations in the network, researchers at ZIB, together with a group of experts at OGE, have developed a smart, forward-looking, analytics-based decision support system. For this to work, it was necessary to utilize three types of analytics:

- Descriptive: modeling and simulating the gas flow in the network
- Predictive: predicting future gas supply and demand at the entries and exits of the network
- Prescriptive: recommending network control measures to ensure safe and efficient operation of the network

## The GasLab in MODAL

In the early 2000s, the European Union changed the regulative framework for the gas market toward fulfilling the main European energy objectives: competitiveness, security of supply, and sustainability. In the unbundled gas market, demands, supply, and storage facilities are beyond the control of the transmission system operator. As part of the decarbonization of the energy system, short-term transport demands for the rapid supply of gas-fired power plants as a supplement to fluctuating electricity generation from renewable energies are increasing sharply. These fluctuations require the dispatching center to have quick reactions, which can only succeed if the network is prepared to deal with such short-notice transports. To tackle this challenge, five years ago, together with OGE, ZIB started the GasLab as part of the Research Campus MODAL, funded by the German Federal Ministry of Education and Research. The goal of the GasLab is to develop a decision support system that predicts the future gas flows, analyses control options and risks, and intelligently computes recommended actions to secure safe operation while minimizing energy consumption and equipment wear.

# Machine Learning Joins Optimization for Decision Support

Every 15 minutes, the optimization core computes recommendations based on the current state of the network, its past, and its technical capabilities. A single station in the network can have more than 1,250 discrete operational modes. Each mode may include target values for continuous quantities such as the target pressure of a regulator. First, an hourly forecast for the more than 1,000 entry and exit points of the network for the next 24 hours is computed employing a mixture of machine learning and optimization on a preselected set of features most suitable for each entry or exit. Based on these forecasts, the recommended operational measures are computed. However, an exact transient model, including all discrete and continuous variables for the full network is computationally intractable. Therefore, we employ a two-phase approach, decomposing the different aspects of the problem by exploiting the topology of the gas network: a coarse model computes amounts and directions of flows, computing when to transport how much gas on which path through the network. Then, detailed models for the individual compressor stations are solved in parallel. These validate and complement the amounts and directions calculated by the coarse model and compute precise action recommendations. The objective is to fulfill demands while minimizing costly mode changes considering many additional constraints needed to ensure the practical feasibility of the recommended action. Finally, the results are summarized and displayed to the dispatchers on dedicated iPads, providing them with a set of directions.

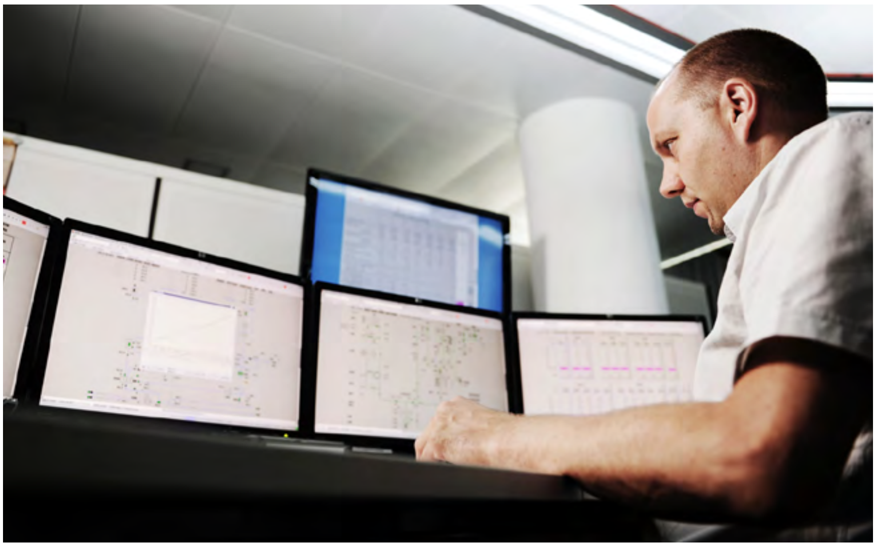

# Consistent State Management for Data Stream Analysis

## Big Data Stream Processing

Data stream processing frameworks such as Flink [1], Spark [2], and Hadoop require the data analysis pipeline to be expressed as a directed dataflow graph. Unfortunately, this is sometimes difficult or even impossible. In some cases, the flow graph is not static but changes its pathways depending on the input data. Despite much research in science and industry, data access and the data handling in stream processing systems is still based on a few, overly simplistic interfaces such as parallel file systems, the Hadoop Distributed File System HDFS, or the Google File System GFS [3, 4, 5].

None of these systems allows to communicate nonfunctional properties of the pending data accesses. Relevant nonfunctional properties include data access patterns, access concurrency, repeated data accesses, required data redundancy, location restrictions, or similar aspects. Clearly, data stream processing could be made much more flexible, portable, adaptable, and robust against single component failures if the nonfunctional properties were expressed before starting an analysis job and taken into consideration by the stream processing environment at runtime.



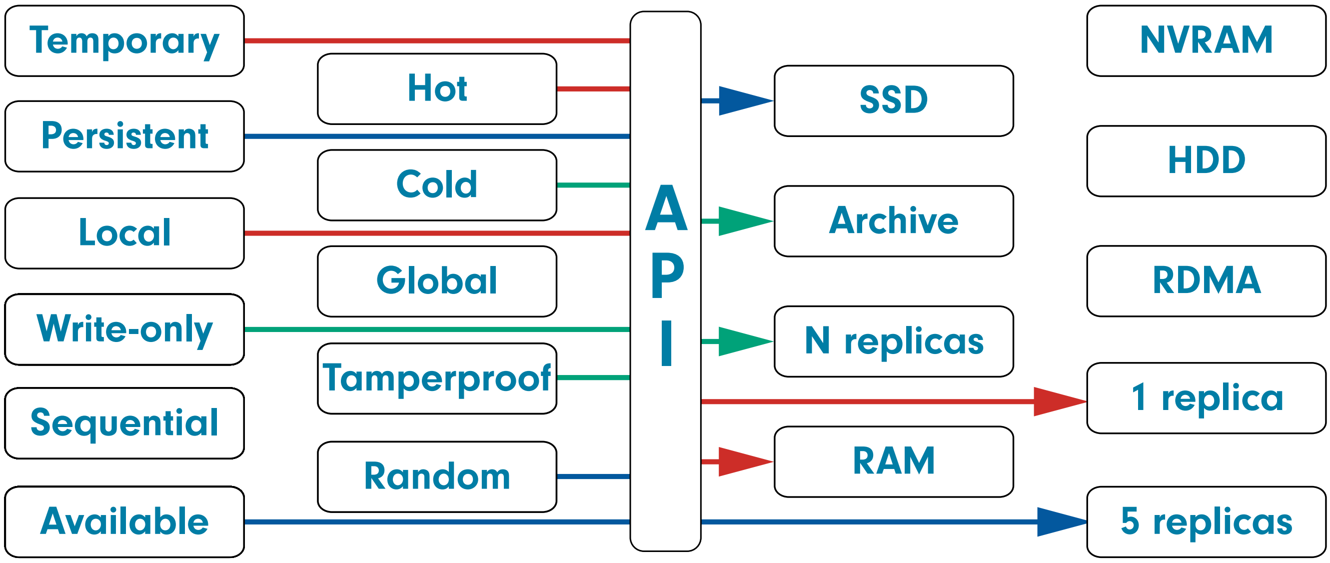


Fig. 5: Declarative API mapping requirements to executable implementations, with three mappings in **red**, **blue** and **green**.

# Consistent State Management – What Is It Good for?

For the current and future big data analysis pipelines with further increasing complexity, it becomes more and more clear that the concept of independently treated data streams leads to workflows with unnecessary synchronization and inefficient resource utilization. With a well-designed and properly implemented consistent state management, stream processing can be much enhanced, leading to better utilized resources and faster job turn-around times. All subtasks of an analysis workflow must be allowed to access the shared state, or a part of it, in order to exchange information among themselves, such as threshold values for search algorithms, previously learned classifiers, or similar.

The consistent and efficient management of the concurrently accessed shared state information, which is distributed over all accessed resources, is a difficult problem – both in theory and practice. Especially in the context of multiple, potentially widely distributed and heterogeneous resources and the calculation on-site, new challenges arise.

In high-security applications like information-based medicine, for example, the data must be processed and analyzed on local resources at the respective institution (clinic or laboratory) and may only be further processed as aggregated, anonymized data. Even so, it should be possible to derive meaningful analysis results by combining aggregated data across the various resources of the participating institutions with better representativeness of the total number of cases and to avoid or consider a possible bias of the sources of an institution.



# Research Questions

Consistent management of distributed state information is at the core of distributed algorithms. Several requirements and research questions arise from this problem:

1. Where should data be placed in order to enable the efficient processing of iterative algorithms in geographically distributed environments?
2. Which description of data locality, degree of aggregation, and anonymization is required?
3. How can failure tolerance and distribution restrictions be specified and realized?
4. How can heterogeneity, dynamic data-space partitioning, self-management, scalability, bandwidth utilization, and incremental data modification be efficiently integrated and implemented?

To answer these questions, we draw on our research and development experience with the two distributed data management systems developed at ZIB: **XtreemFS**, a failure-tolerant distributed file system, and **Scalaris**, a distributed failure-tolerant NoSQL database. A suitable declarative abstraction layer will be developed for various requirements of state management, e.g. regarding failure tolerance, data placement, I/O characteristics, and scalability, which will then be realized automatically with the available resources and transparently optimized for the application.

Our improved Paxos algorithm with in-place consensus sequences [6] will serve as a basis for new algorithms. Additionally, modern hardware features such as one-sided communication with remote direct memory access (RDMA) and persistent memory (nonvolatile memory, NVM) will play an important role in deriving new solutions. While all these new hardware features may be beneficial for specific computer science systems, they pose additional challenges for distributed systems, such as atomicity and visibility of operations with concurrent accesses in the face of compute nodes operating under the crash recovery failure model.

# Extracting Dynamical Laws from Data

In all areas of science, technology, and society, vast amounts of data are becoming available that can be utilized to analyze, control, and understand the underlying complex processes. The past decade has seen tremendous progress regarding algorithmic methods for data-driven prediction of processes. In classical machine-learning settings, it is often acceptable to have a black-box prediction method, provided that it performs efficiently and with small error. In science and technology, the essential objective is often to obtain an interpretable model that leads to a fundamental understanding of the observed process. One aims to derive a theory that allows for successful predictions in regimes that are very different from those in which the underlying data was gathered. Estimation methods that can generate understanding and deliver physically interpretable and mathematically tractable models from data are still in their infancy.

The real challenge lies in inferring effective equations in mesoscopic regimes of complex systems, where the microscopic dynamical laws are not known or do not lead to a practical description and human intuition is often overwhelmed by the sheer complexity of the system. Examples include developmental biology, where the development and cell differentiation of an organism can be tracked by time-dependent microscopy with single-cell resolution, or social or societal processes, where, in most cases, a fundamental theory that would allow a model for the dynamics to be derived, are simply not available. At ZIB, our aim is to combine machine learning and mathematical process simulation in order to derive effective dynamical laws and to obtain models that generate understanding and insight.

# Data-Driven Transfer Operator Techniques

Data-driven approaches for the analysis of complex dynamical systems – be it methods to approximate transfer operators for computing metastable or coherent sets, methods to learn physical laws, or methods for optimization and control – have been steadily gaining popularity over the last years. Algorithms such as EDMD [7, 8], SINDy [9], and their various kernel-, tensor-, or neural-network-based extensions and generalizations have been successfully applied to a plethora of different problems, including molecular and fluid dynamics, meteorology, as well as mechanical and electrical engineering. Most of the aforementioned techniques turn out to be strongly related, with the unifying concept being transfer operator theory, whether Koopman or Perron-Frobenius or their stochastic variants.

Inferring dynamical laws from data mostly means estimating the right-hand side of a (mostly stochastic, perhaps partial) differential equation given rich enough time series data. SINDy constitutes a milestone for such a purely data-driven discovery of dynamical systems. Its main idea is to approximate the right-hand side of the differential equation via a sparse linear combination of terms coming from a rich library of ansatz functions. While SINDy performs very reliably for deterministic dynamics, there are some pitfalls for stochastic settings or large amounts of noise. In order to overcome these obstacles, we developed a method for estimating the generator of the Koopman operator from data [10], based on an ansatz space spanned by a rich library of nonlinear functions as in SINDy.

# First Results

To this end, in one of our BIFOLD projects, we developed a general framework for computing an approximation of the Koopman generator, both for deterministic and stochastic systems, and started exploring a range of applications [10]:

We illustrated that the governing equations of deterministic as well as stochastic dynamical systems can be obtained from empirical estimates of the generator.

In addition, we explored two powerful applications of the approximated Koopman generator. We showed that the resulting method, termed gEDMD, can be used to identify coarse-grained models based on data of the full system, which is a highly relevant topic across different research fields, like molecular dynamics simulations for instance. Moreover, we apply the Koopman generator to control dynamical systems, providing flexible and efficient model predictive control strategies.

In the next step, we will investigate how gEDMD can be applied to social and societal processes, e.g. in opinion formation or mobility decision processes.

Photo by Ryan Quintal on Unsplash

# Stream Processing with Blockchain Technologies

## Auditability and Reproducibility in Stream Processing

Reproducibility and traceability play a central role in justifying and supporting credibility not only in science but also in business applications. To realise both for data-intensive research applications, for example in the field of machine learning, that handle dynamic unbounded data streams, several methodological and technical challenges have to be solved.

Traceability for small data sets can be reached with the promising technology of blockchains, which store and distribute intermediate states of a sequence of data manipulations so that they cannot be tampered with by individuals, as in a ledger. To process large amounts of data without the digital ledger growing to an unmanageable size, data can be stored off chain. The challenge here is to protect the off-chain data from manipulation. It becomes even more difficult if not only the processing of static, closed data objects is to be made credible and auditable, but also the results of a continuous data analysis on data streams, which is to be made credible and auditable up to a given point in time. For this problem, no suitable solutions are known yet.



## Ongoing Research

We are developing a distributed data management system that supports and enables traceability and reproducibility of the results of distributed continuous data analysis on data streams. To achieve this goal, we are coupling blockchain technology with cluster data management. More specifically, we are researching, developing, and evaluating methods for manipulation protection of off-chain data in distributed environments, including algorithms for failure-tolerant, coordinated snapshot generation for distributed data stream processing. Detection and reliable deletion of data that is no longer required is another difficult task that relies on distributed snapshot processing. Our work gives answers to the following research questions:

1. How can continuously running distributed data stream applications be made auditable and reproducible?
2. How can blockchain technology and distributed cluster data management be combined to enable tamper-free, consistent snapshots on a group of data objects (e.g. files), to support distributed data stream processing applications? The application should be influenced as little as possible, i.e. the coordination of a corresponding snapshot creation should be asynchronous.
3. How and when can data objects that are no longer referenced be cleaned up automatically in order to keep storage requirements at a moderate size? Time specifications for deletion periods could also be taken into account.
4. Which data structures and layouts are suitable to enable efficient access to the latest state as well as to snapshots with low storage overhead in respect to storage capacity?

# WHY IS THERE A NEW WAVE OF ATTENTION FOR FPGAs?

Thomas Steinke | +49 30 84185-144 | steinke@zib.de

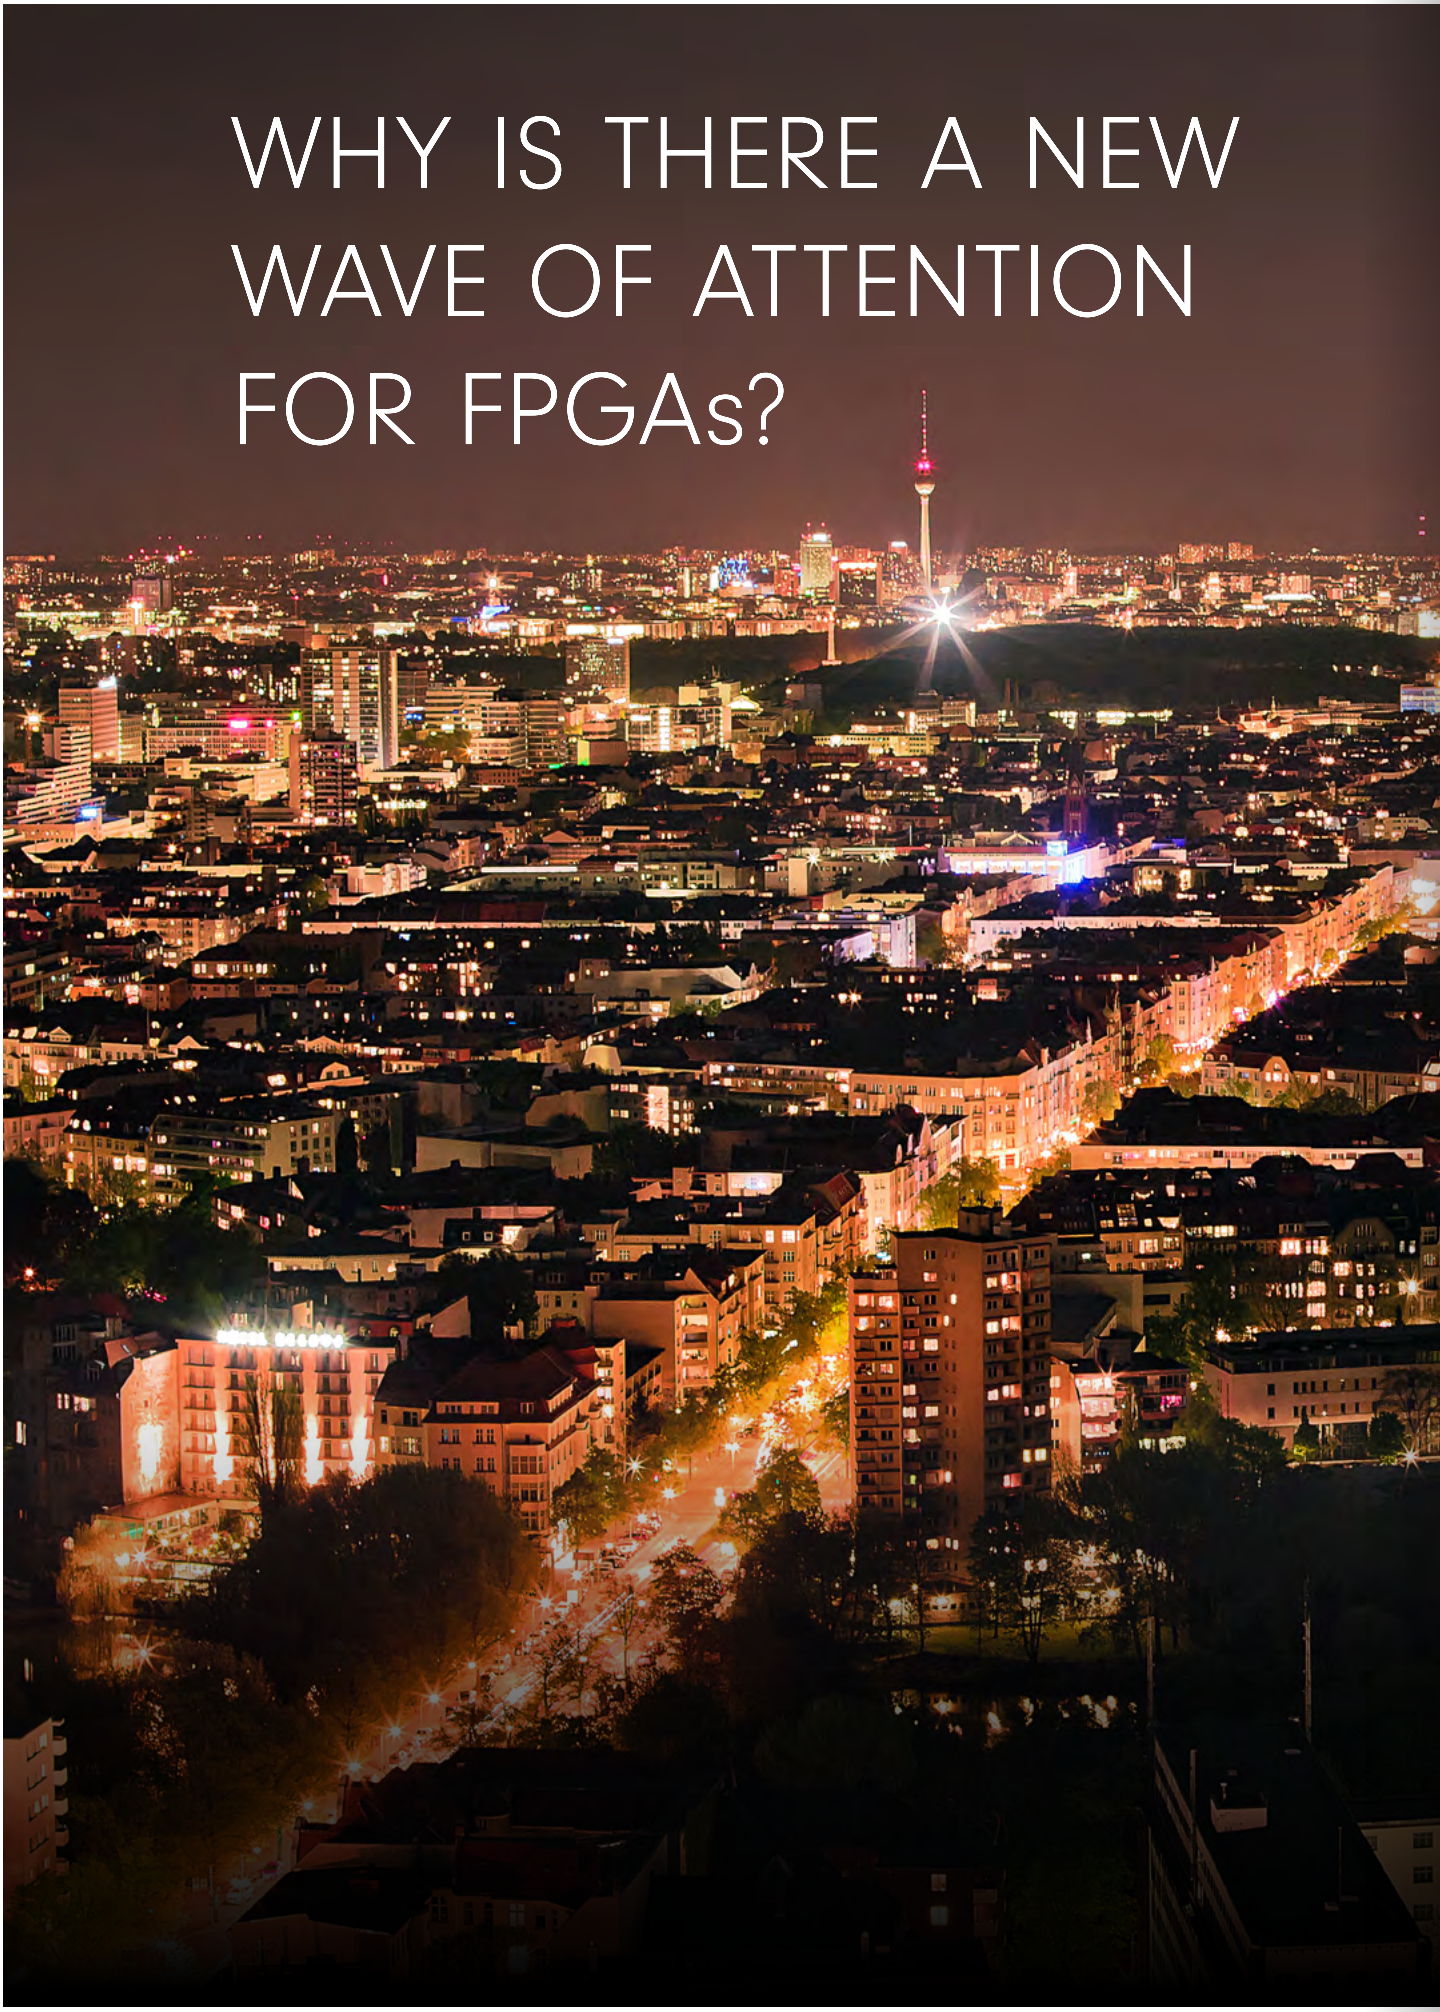

Currently, reconfigurable computing with field-programmable gate arrays (FPGA) is experiencing a new wave of attention, particularly in the HPC community. On FPGAs, highly data-parallel architectures can be implemented to target more power-efficient solutions. Here are some of the challenges of "programming" an FPGA.

# Reconfigurable Computing with Field-Programmable Gate Arrays

**Reconfigurable computing** is a class of custom computing where algorithms are not implemented in software for a fixed logic device like a CPU but on a device, which allows high-level logic functions to be defined by the customer depending on the algorithmic needs. Field-programmable gate arrays (FPGA) are such devices enabling a custom circuit design that can be loaded within microseconds – the configuration step. This ability to reconfigure an FPGA enables the power-efficient implementation, since only the functions are configured that the algorithm really needs.

Modern FPGAs provide basic logical elements like lookup tables (LUT), flip-flops (FF), and block RAM (BRAM) but increasingly also hard-code functional units like DSPs or CPU cores. Furthermore, interfaces to the chip-surrounding infrastructure as memory (DRAM, SRAM), communication (Ethernet, PCIe), and XXX require corresponding IP core that are integrated by the design tools on demand onto the FPGA. This way, in particular large FPGAs converge to system-on-a-chip (SoC) devices. Fig. 1 gives an overall schema and Table 1 summarizes key features of a modern FPGA suitable for high-performance computing and data analytics.

Since the birth of FPGAs in the mid-1980s, their capabilities have grown exponentially and more features have been added over time. FPGAs made their way into communication and encryption systems. Not surprisingly, within the last decade, FPGAs have become attractive for system providers and software developers seeking for energy-efficient operation and solutions, respectively. Life science (in particular genomics), search engines, high-frequency trading, or further user-defined functions near communication paths become new application domains beyond the traditional FPGA success stories. Whereever a low latency and limited power are critical demands, FPGA might be the right choice as a target platform. IoT applications (edge computing) and self-driving cars are recent examples. FPGAs are shining with highly data-parallel algorithms, since the lower clock rate of some hundreds of megahertz needs to be compensated.

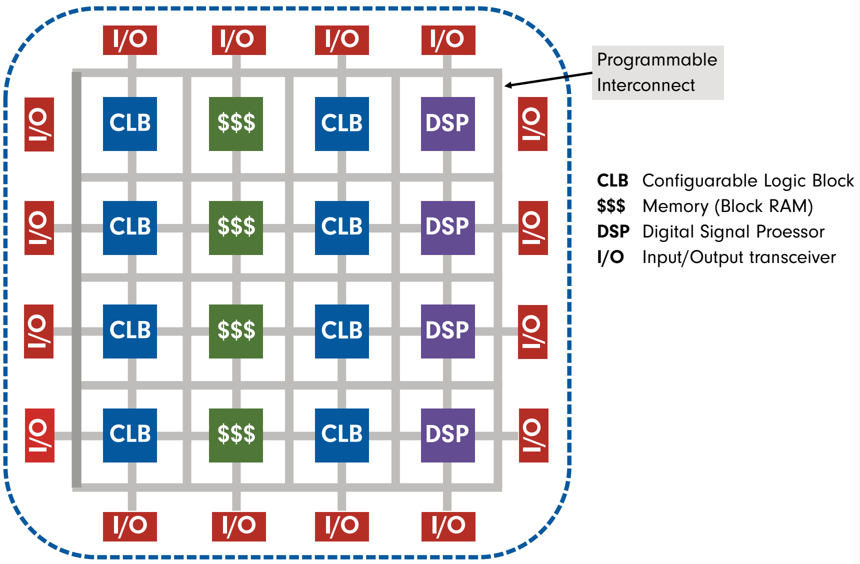


Fig. 1: Simplified schema of an FPGA. Configurable logic blocks (CLB) consisting of lookup tables (LUT) and flip-flops (FF), block RAM (BRAM), and hard-coded IP like digital signal processors (DSP) can be connected through a programmable interconnection. Transceiver blocks (I/O) enable high-speed communication to the outside world (e.g. PCIe, Ethernet).

Typical FPGA platform infrastructure for HPC and DA workloads include multiple memory banks and a host interface, e.g. PCIe for system integration. A unique feature of FPGA devices is multiple serial links for low-latency, high-bandwidth communication, which is attractive for custom configurable high-speed interconnections linking multiple FPGA platforms or providing communication links to external data sources (experiments, IoT).

For selected workloads and steps in computational workflows reconfigurable computing can help to implement energy-efficient solutions. This way, HPC configurations with FPGA partitions can help in exascale systems to offer the required computational power while reducing the energy demands. To extend the usage profile of FPGAs to a broader range, i.e. making reconfigurable computing accessible to the majority of HPC software developers, one has to tackle the challenge of "programming" this device class. Some of the approaches are described in the next sections.

| Key Features of a Modern FPGA |
|---|
| More than **900,000** configurable logic blocks |
| **5,760** DSPs for fixed point and IEEE 754 single precision floating-point operations |
| **11,721** SRAM blocks, configurable in data width and number of access ports |
| Up to **96** serial transceivers with up to **28.3 Gbit/s** |
| Integrated DDR4 memory controllers |
| Typical design frequencies **300–600 MHz** |
| Typical power consumption **50–225 W** |

Table 1: Typical key features of a modern FPGA (here an Intel Stratix 10 GX2800 FPGA).

# Programming FPGAs – Really?



With so many interesting features, it is a surprise why FPGAs are not in widespread use. In contrast to CPUs, where, over the last six decades, a hardware abstraction via high-level programming languages and powerful compilers evolved, migrating an algorithm onto an FPGA means circuit design (a special engineering domain). Typical hardware description languages (VHDL, Verilog) have to be implemented and are still mainly used by engineers who want to squeeze the last bit of performance out of the hardware. But, in the last decade, approaches have become available that provide a higher abstraction level to the underlying hardware like high-level-synthesis (HLS) languages or recently OpenCL.

HLS languages are typically C/C++ derivatives supporting a subset of the C/C++ language-plus-vendor-specific compiler directives to help the C-to-HDL compiler in the mapping and optimization process. More details about the inner workings of an HLS is presented in the next section.

OpenCL [1] is a framework that provides a standardized API and runtime system for heterogeneous platforms. Together with the vendor-specific and platform-dependent infrastructure, it provides one of the highest abstracted programming models and a convenient approach to program FPGAs. With OpenCL, the software approach to "programming" an FPGA is today becoming the dominant method for a wide group of developers. Fig. 2 schematically classifies performance versus development time for the different programming approaches.

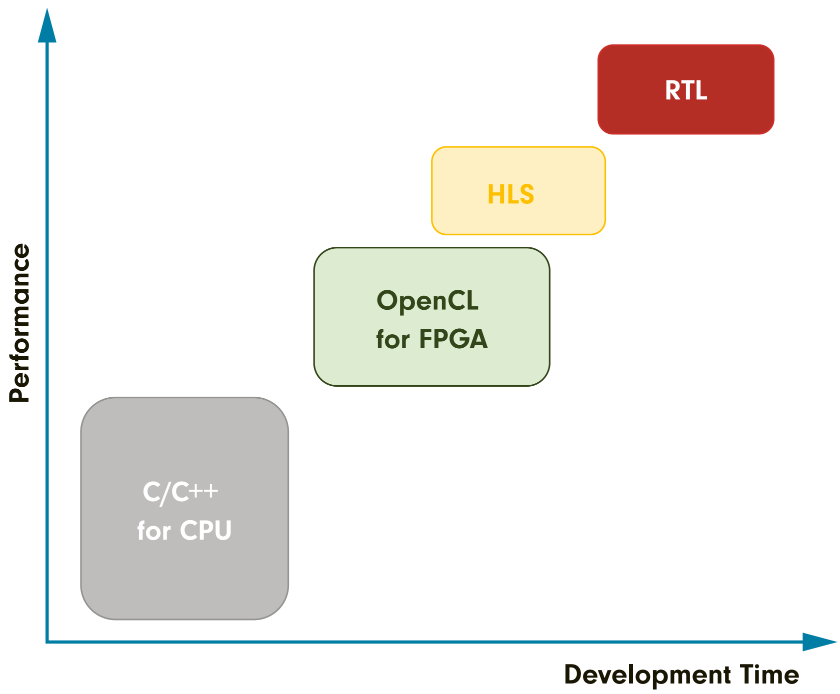


Idea taken from: Hoozemans, J., de Jong, R., van der Vlugt, S. et al. Frame-based Programming, Stream-Based Processing for Medical Image Processing Applications. J Sign Process Syst 91, 47–59 (2019). https://doi.org/10.1007/s11265-018-1422-3

Fig. 2: High-level synthesis (HLS) aims to reduce development time compared to traditional RTL (register-transfer level) circuit design approach. Using the OpenCL approach, it costs less time to synthesize the first working design to FPGA, but more time-consuming test, debug, and optimization cycles are needed before the performance is comparable to an HLS design (idea taken from [2]).

One could imagine that bringing an algorithm onto an FPGA is now, with these high-level approaches mentioned above, comparable to the common software development cycle for CPUs and GPUs. Unfortunately, there is still a very time-consuming step that is part of all toolchains present in the chosen programming language: the so-called place-and-route (P&R) step. The P&R step can take hours, even when using multiple processing threads on modern hardware. Fig. 3 shows an overall view of the traditional and the OpenCL design and tool flow.

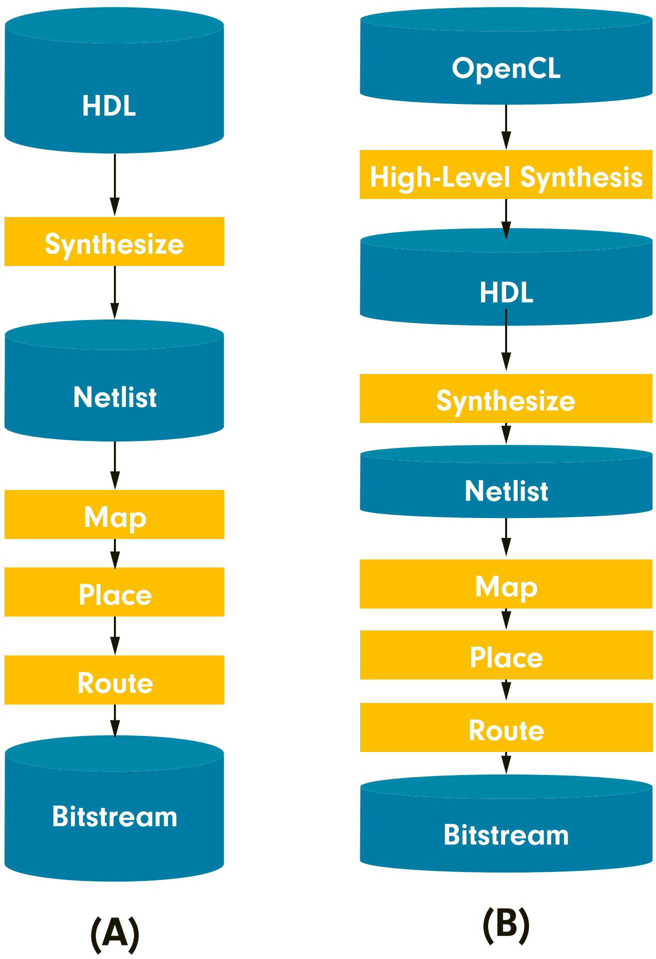


Fig. 3: (A) The traditional design flow using a hardware description language (HDL). The synthesis tool generates a netlist of basic logic elements, which is then mapped to the available basic logic on the FPGA, placed on the chip, and the connecting signals are routed through the interconnection network. The results of the tool flow is the FPGA configuration file (bitstream).
(B) In the OpenCL tool flow, the first step is the transformation of the C++14-based kernel code to be offloaded onto the FPGA into a synthesizable HDL code. Afterwards, the steps of (A) are processed. The host code is identified by the OpenCL compiler and generates the executable for the host CPU. This executable man ages the FPGA device and issues the FPGA kernel calls (schemas from [3]).

Schemas taken from: T. Kenter, OpenCL design flows for Intel and Xilinx FPGAs – common optimization strategies, design patterns and vendor-specific differences, Tutorial, 2019 conference, March 25, 2019, https://pc2.uni-paderborn.de/fileadmin/pc2/presentations/Kenter-OpenCL-FPGA-Tutorial-2019-03-25-Website-compressed-fixed.pdf

The place-and-route step is about placing the higher functional logic blocks on the FPGA and connecting these blocks in the most achievable, optimal way so that the design requirements such as highest possible frequency (shortest/longest routing path) or smallest possible area usage can be met.

To shorten the development process for program developers (and circuit designers), software emulators running on standard hardware for incremental code development and functional debugging are used. Only at certain stages of the last time-consuming step, the generation of the FPGA configuration file (bitstream) is triggered.

# High-Level Synthesis – a Look Under the Hood



In order to use FPGAs as accelerators for HPC, it is important that they can be programmed by software engineers without extensive hardware design knowledge. This is where high-level synthesis (HLS) comes into play, the translation of a high-level language (HLL) like C/C++ into a hardware description language (HDL), which can be further synthesized to configure the FPGA accelerator.

Even though the compiler is capable of many automatic optimizations, as for every other target (e.g. CPU or GPU), it is important to design and express the algorithm in a manner that allows the advantages of that platform to be utilized. To be able to accomplish this and to gain trust in the compiler, one needs to roughly understand the inner workings of the HLS.

Most modern HLS tools (e.g. Intel and Xilinx) are LLVM-based, they will use a pre-existing compiler front end (e.g. Clang for C/C++) to translate the HLL into a target-independent intermediate representation (IR). The LLVM IR is similar to a typed RISC assembly and it follows the static single assignment (SSA) form. This form will assign each variable only once and create a copy if it is further modified. It is meant to simplify dependencies analysis but also suits the FPGA as a target very well, as there is no strictly limited set of registers (like with traditional CPUs) and operations are often just chained without saving them to registers in-between.

On the IR-level, a set of pre-existing and custom optimization paths will be executed that fit the FPGA as a target. A prominent example is loop unrolling.

The compiler-frameworked back end will normally lower the IR into the target's instruction set and allocate the variables in SSA form to the registers of the CPU. For the FPGA, this compares the translation to an HDL and the allocation of hardware resources (like DSPs). In general, a state machine going through the instructions is created. In the simplest (but also the least efficient) method, this will go through the instructions one by one. To exploit the parallel architecture of the FPGA, it is necessary to schedule the instructions in parallel.

This is typically achieved by viewing the program as a control-data flow graph, while the nodes represent instructions and the edges define the dependencies. The dependencies are then modeled as a set of integer difference constraints ($x - y \leq b$). Further resource and timing constraints are added and finally the system is ready for a linear scheduling objective [4]. The result will assign the appropriate state to each instruction.

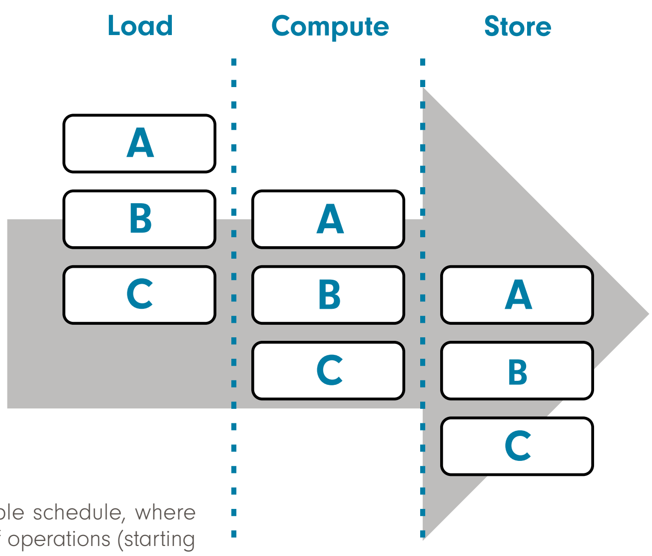


Fig. 4: A simplified example schedule, where two independent chains of operations (starting with a load) get scheduled in parallel and later merge during the third operation before the result gets stored in the memory.

# HLS Optimization – a Glimpse into the Best Practices

One of the greatest advantages of HLS over an HDL is that the development of the algorithm and iterative functional testing can be achieved by emulation on the CPU, which is much faster than simulation or even synthesis. The next step is to optimize based on the reports of the HLS, and only in the last step the performance will be tweaked based on real synthesis results. These reports include loop analysis and area estimates that will guide the following optimization practices.

Besides designing the algorithm in an FPGA-friendly way and structuring the code accordingly, there is a set of annotations to accommodate the source code for the FPGA. These range from annotations to provide additional information or guarantees to the compiler (to help the automatic optimization) up to annotations to instruct the compiler to handle constructs in a certain way. One of the most important candidates is the annotation to unroll loops, as this directly transforms sequential logic (feedback loop) to combinatorial logic. This mainly increases throughput per clock cycle with the trade-off of a higher area demand.

In case it is not feasible to unroll a loop, the compiler will automatically try to pipeline it, for example by overlapping load, computation, and store. Due to loop-carried dependencies, it might not be possible to launch a new loop iteration every clock cycle. This delay is called initiation interval (II). The goal is to minimize the II in order to fully utilize the hardware. This is mainly achieved by introducing local buffers to avoid or at least relax these dependencies.

Even though memory accesses will be automatically prefetched, burst-coalesced, or cached (depending on the pattern), it is always advisable to introduce manual buffering. This is preferable in registers (single-cycle access) or otherwise in block RAM (SRAM, multiple cycles but fixed latency). This has the advantage of pipelines not being stalled or it even allowing loops to be unrolled.

To imply registers over block RAM, it is important to access the buffers statically, so it can basically be hardwired instead of an address-based dynamic access. A static access, and hence buffering in local registers, is typically archived by implying a shift-register-like structure, where the new element always gets pushed in at one end of the buffer (fixed location).

To squeeze out the last bit of performance and create a safe area, it is possible to use custom types, like arbitrary precision integers with a custom number of bits or relax floating point operations, so the order can be changed for them to be executed in parallel.

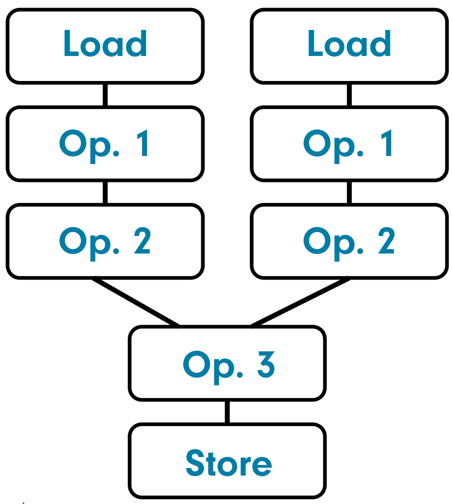


Fig. 5: A simplified pipeline with three stages (columns) shown over time (rows). For example, while the first element A is in the second cycle already in the compute stage, the next element B gets loaded in parallel.

# OpenMP Offloading on FPGAs

Even though HLS greatly simplifies the effort of hardware design, it typically only provides an IP block meant to be integrated into the FPGA infrastructure. This still leaves the developer to use the hardware design tools to build a framework for it, e.g. with host (PCI) and memory connections.

The OpenCL toolchains greatly reduce the knowledge needed by providing a fixed infrastructure based on partial reconfiguration, and handling the integration of the user's component. But, as there is more OpenMP than OpenCL legacy code used within the HPC community and more users are familiar with it, we set out for the quest to bring OpenMP to the FPGA world as part of the ORKA-HPC project [5].

So far, only a limited amount of work was done in that direction. A survey [6] showed that either OpenMP was only handling the configuration of the host side, an old OpenMP version without target pragma was used, or the work only had the character of a prototype and is no longer supported.

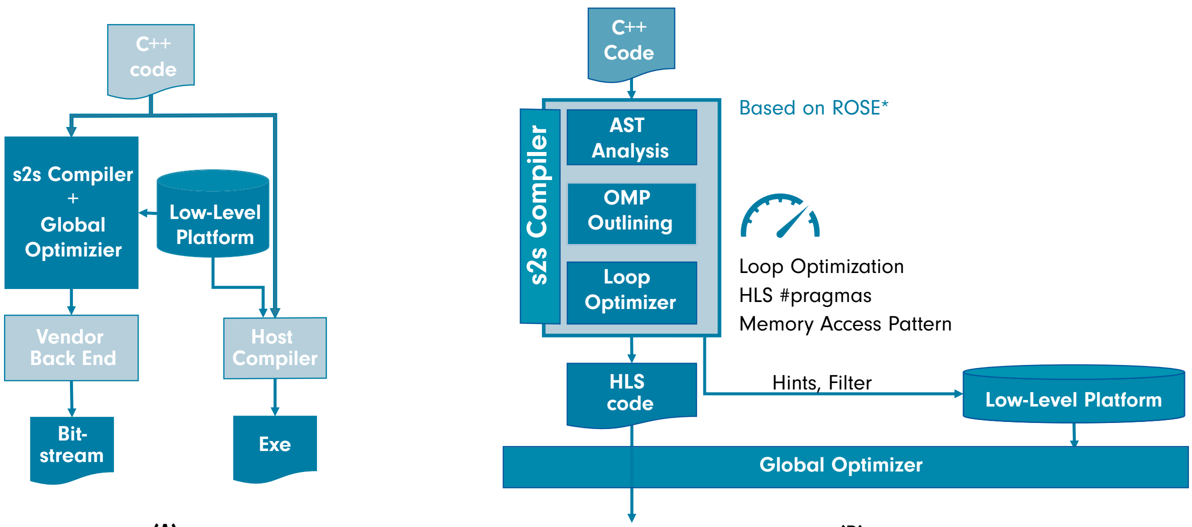


**(A)** **(B)**

Fig. 6: Schematic overview of the ORKA-HPC tool flow to support the OpenMP programming model for FPGAs. **(A)** The design goal of ORKA-HPC is a fully automated workflow without any required user interaction. The vendor toolchains are used as backends. **(B)** The source-to-source compiler is based on the ROSE compiler framework and manages the OpenMP outlining, the local optimizations at loop level (pipelining, unrolling), and emits platform-dependent HLS code. If possible, basic memory access patterns are identified to support the selection of an appropriate low-level platform.

A first approach was to utilize the pre-existing OpenCL toolchains [7] of the vendors by outlining the annotated OpenMP regions with Clang, feeding the resulting IR into the OpenCL toolchain, and replacing the host code with OpenCL API calls. This proved to work out but showed limitations regarding the flexibility of the proprietary toolchains.

Therefore, a custom solution was created. It uses the ROSE compiler framework for the outlining task and the standard HLS tools for the generation of the IP blocks representing the user-written C++ functions. These IP blocks are then embedded into a low-level platform created by an ORKA-HPC tool using information provided by the OpenMP annotations. The final implementation of the ORKA-HPC toolchain (see Fig. 6) by integrating the developed individual tools into the overall tool flow is in progress.



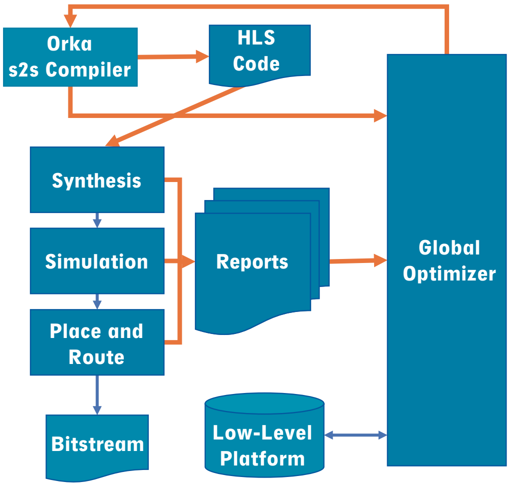


(C)

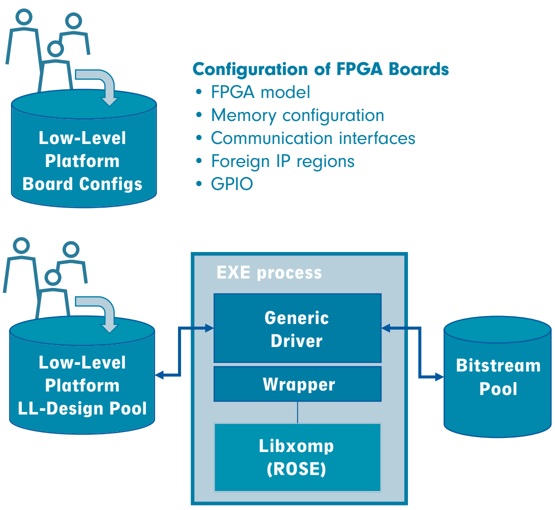


(D)

**(C)** The global optimization step realizes a heuristic exploration of design space (clock rate, resource utilization) and gives feedback to the src-to-src compiler (orange lines). The reports of the synthesis, simulation, place, and route tools guide the global optimization for which evolutionary algorithms are used.

**(D)** The low-level platform and runtime system manages predesigned low-level platforms. It includes a generic driver and runtime API for the host – FPGA board communication and, for example, the monitoring of platform parameters (temperature, voltages).

# REFERENCES


## GREEN ENERGY

[1] C. Becker et al. Nanophotonic Enhanced Two-Photon Excited Photoluminescence of Perovskite Quantum Dots. ACS Photonics 5, 4668 (2018).

[2] M. Hammerschmidt et al. Field Heterogeneities and Their Impact on Photocatalysis: Combining Optical and Kinetic Monte Carlo Simulations on the Nanoscale. J. Phys. Chem. C 124, 3177, (2020).

[3] P. Schneider et al. Benchmarking Five Global Optimization Approaches for Nano-optical Shape Optimization and Parameter Reconstruction. ACS Photonics 6, 2726, (2019).

[4] P. Manley et al. Absorption Enhancement for Ultrathin Solar Fuel Devices with Plasmonic Gratings. ACS Appl. Energy Mater. 1, 5810, (2018).

[5] J. Wen et al. Membrane Orientation of the FMO Antenna Protein from Chlorobaculum Tepidum as Determined by Mass Spectrometry-Based Footprinting. Proc. Nat. Acad. Sci. 106, 6134, (2009).

[6] J. M. Olson. The FMO Protein. Discoveries in Photosynthesis. Springer, Dordrecht (2005).

[7] T. Brixner et al. Two-Dimensional Spectroscopy of Electronic Couplings in Photosynthesis. Nature 434, 625, (2005).

[8] T. Kramer et al. Efficient Calculation of Open Quantum System Dynamics and Time-Resolved Spectroscopy with Distributed Memory HEOM (DM-HEOM). J. Comp. Chem. 39, 1779, (2018).

[9] M. Noack et al. DM-HEOM: A Portable and Scalable Solver-Framework for the Hierarchical Equations of Motion. IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), 201, (2018).

[10] J. Dostál et al. In Situ Mapping of the Energy Flow through the Entire Photosynthetic Apparatus. Nat. Chem. 8, 705, (2016).

[11] T. Kramer et al. Two-Dimensional Electronic Spectra of the Photosynthetic Apparatus of Green Sulfur Bacteria. Sci. Rep. 7, 45245, (2017).

[12] T. Kramer, et al. Effect of Disorder and Polarization Sequences on Two-Dimensional Spectra of Light-Harvesting Complexes. Photosynth. Res., just accepted (2019).

[13] D. Chen et al. Nanophotonic Light Management for Perovskite-Silicon Tandem Solar Cells. J. Photonics Energy 8, 022601, (2018).

[14] P. Tillmann et al. Minimising the Levelised Cost of Electricity for Bifacial Solar Panel Arrays Using Bayesian Optimisation. Sustain. Energy Devices 4, 254, (2020).

## RIEMANNIAN ANALYSIS OF TIME-VARYING SHAPE DATA – UNDERSTANDING GEOMETRIC EVOLUTIONS

[1] F. Ambellan, S. Zachow, C. v. Tycowicz. A Surface-Theoretic Approach for Statistical Shape Modeling. Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI), Part IV, pp. 21–29, Vol.11767, Lecture Notes in Computer Science (2019).

[2] F. Ambellan, S. Zachow, C. v. Tycowicz. An as-invariant-as-possible $GL^+(3)$-based Statistical Shape Model. Proc. 7th MICCAI workshop on Mathematical Foundations of Computational Anatomy (MFCA), pp. 219–228, Vol. 11846, Lecture Notes in Computer Science (2019).

[3] F. Ambellan, S. Zachow, C. v. Tycowicz. Statistical Shape Models – Understanding and Mastering Variation in Anatomy. Biomedical Visualisation, P. M. Rea (Ed.), Springer Nature Switzerland AG, 1, pp. 67–84, (2019).

[4] C. Brandt, C. v. Tycowicz, K. Hildebrandt. Geometric Flows of Curves in Shape Space for Processing Motion of Deformable Objects. Computer Graphics Forum, 35(2), (2016).

[5] DG. Kendall. A Survey of the Statistical Theory of Shape. Statistical Science 4:2, pp. 87–99, (1989).

[6] P. Muralidharan, P. T. Fletcher. Sasaki metrics for analysis of longitudinal data on manifolds. IEEE Conference on Computer Vision and Pattern Recognition, pp. 1027–1034, (2012).

[7] E. Nava-Yazdani, H-C. Hege, C. v. Tycowicz, T. Sullivan. Geodesic Analysis in Kendall's Shape Space with Epidemiological Applications. Journal of Mathematical Imaging and Vision, accepted for publication (2020).

[8] E. Nava-Yazdani, H-C. Hege, C. v. Tycowicz. A Geodesic Mixed Effects Model in Kendall's Shape Space. Proc. 7th MICCAI workshop on Mathematical Foundations of Computational Anatomy (MFCA), Lecture Notes in Computer Science (2019).

[9] C. v. Tycowicz, F. Ambellan, A. Mukhopadhyay, S. Zachow. An efficient Riemannian statistical shape model using differential coordinates: With application to the classification of data from the Osteoarthritis Initiative. Medical Image Analysis, 43, 1–9, (2018).

[10] C. v. Tycowicz. Towards Shape-based Knee Osteoarthritis Classification using Graph Convolutional Networks. IEEE 17th International Symposium on Biomedical Imaging (ISBI 2020) pp. 750–753, (2020).

## SOLVING REAL-WORLD OPTIMIZATION PROBLEMS


[1] A. Gleixner, M. Bastubbe, L. Eifler, T. Gally, G. Gamrath, R. L. Gottwald, G. Hendel, C. Hojny, T. Koch, M. E. Lübbecke, S. J. Maher, M. Miltenberger, B. Müller, M. E. Pfetsch, C. Puchert, D. Rehfeldt, F. Schlösser, C. Schubert, F. Serrano, Y. Shinano, J. M. Viernickel, M. Walter, F. Wegscheider, J. T. Witt, J. Witzig. The SCIP Optimization Suite 6.0, ZIB-Report 18-26, (2018).

[2] A. Gleixner, L. Eifler, T. Gally, G. Gamrath, P. Gemander, R. L. Gottwald, G. Hendel, C. Hojny, T. Koch, M. Miltenberger, B. Müller, Marc E. Pfetsch, C. Puchert, D. Rehfeldt, F. Schlösser, F. Serrano, Y. Shinano, J. M. Viernickel, S. Vigerske, D. Weninger, J. T. Witt, J. Witzig. The SCIP Optimization Suite 5.0, ZIB-Report 17-61, (2017).

[3] S. J. Maher, T. Fischer, T. Gally, G. Gamrath, A. Gleixner, R. L. Gottwald, G. Hendel, T. Koch, M. E. Lübbecke, M. Miltenberger, B. Müller, M. E. Pfetsch, C. Puchert, D. Rehfeldt, S. Schenker, R. Schwarz, F. Serrano, Y. Shinano, D. Weninger, J. T. Witt, J. Witzig. The SCIP Optimization Suite 4.0, ZIB-Report 17-12, (2017).

[4] G. Gamrath, T. Fischer, T. Gally, A. M. Gleixner, G. Hendel, T. Koch, S. J. Maher, M. Miltenberger, B.Müller, M. E. Pfetsch, C. Puchert, D. Rehfeldt, S. Schenker, R. Schwarz, F. Serrano, Y. Shinano, S. Vigerske, D. Weninger, M. Winkler, J. T. Witt, J. Witzig. The SCIP Optimization Suite 3.2, ZIB-Report 15-60, (2016).

[5] G. Gamrath, A. Gleixner, T. Koch, M. Miltenberger, D. Kniasew, D. Schlögel, D. Weninger, Tackling Industrial-Scale Supply Chain Problems by Mixed-Integer Programming (2016).

[6] Description of Model Interconnections D. Beulertz et al. (2018). (plan4res Deliverable No. D3.1) Retrieved from the project website of plan4res: https://www.plan4res.eu/wp-content/uploads/2019/02/D3.1-ModelsInterconnections_Summary.pdf

[7] J. A. Lerman, D. R. Hyduke, H. Latif et al. In Silico Method for Modelling Metabolism and Gene Product Expression at Genome Scale, Nat Commun 3, 929 (2012).

[8] J. Pulaj, Cutting Planes for Union-Closed Families, TU Berlin, doctoral thesis, (2017).

[9] L. Eifler, A. Gleixner, J. Pulaj. Chvátal's Conjecture Holds for Ground Sets of Seven Elements, ZIB-Report 18-49, (2018).

[10] T. Achterberg, Constraint Integer Programming, TU Berlin, doctoral thesis (2009).


## MACHINE LEARNING AND BIG DATA


[1] P. Carbone, A. Katsifodimos, S. Ewen, V. Markl, S. Haridi, and K. Tzoumas: Apache Flink™: Stream and Batch Processing in a Single Engine. In IEEE Data Eng. Bull. 38(4): 28-38, (2015).

[2] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica: Spark: Cluster Computing with Working Sets. In Proceedings of the 2nd USENIX Conference on Hot Topics in Cloud Computing, 10, (2010).

[3] S. Ghemawat, H. Gobioff, and S.-T. Leung: The Google File System. In ACM Symposium on Operating System Principles (SOSP), (2003).

[4] F. Chang, J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. E. Gruber: Bigtable: A Distributed Storage System for Structured Data. In ACM Transactions on Computer Systems (TOCS), 26(2), (2008).

[5] R. Schmidtke, G. Laubender, and T. Steinke: Big Data Analytics on Cray XC Series DataWarp Using Hadoop, Spark and Flink. In Proceedings of CUG, (2016).

[6] J. Skrzypczak, F. Schintke, T. Schütt: Linearizable State Machine Replication of State-Based CRDTs without Logs, Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, 455-457, (2019).

[7] M. O. Williams, C. W. Rowley, and I. G. Kevrekidis. A Kernel-Based Method for Data-Driven Koopman Spectral Analysis. Journal of Computational Dynamics, 2(2):247-265. doi: 10.3934/jcd.2015005, (2015).

[8] S. Klus, P. Koltai, and C. Schuette. On the Numerical Approximation of the Perron-Frobenius and Koopman Operator. Journal of Computational Dynamics, 3(1):51-79, 2016. doi:10.3934/jcd.2016003, (2016).

[9] S. L. Brunton, J. L. Proctor, and J. N. Kutz Discovering Governing Equations from Data by Sparse Identification of Nonlinear Dynamical Systems. Proceedings of the National Academy of Sciences, 113(15):3932-3937. doi:10.1073/pnas.1517384113, (2016).

[10] S. Klus, F. Nüske, S. Peitz, J.H. Niemann, C. Clementi, Ch. Schütte. Data-Driven Approximation of the Koopman Generator: Model Reduction, System Identification, and Control. Accepted for Physica D. (2019).

[1] The Khronos Group: OpenCL. https://www.khronos.org/opencl/

[2] J. Hoozemans, R. de Jong, S. van der Vlugt et al. Frame-Based Programming, Stream-Based Processing for Medical Image Processing Applications. J Sign Process Syst 91, pp. 47–59. https://doi.org/10.1007/s11265-018-1422-3 (2019).

[3] T. Kenter, OpenCL Design Flows for Intel and Xilinx FPGAs – Common Optimization Strategies, Design Patterns and Vendor-Specific Differences, Tutorial, DATE 2019 conference. https://pc2.uni-paderborn.de/fileadmin/pc2/presentations/Kenter-OpenCL-FPGA-Tutorial-2019-03-25-Website-compressed-fixed.pdf (2019).

[4] J. Cong, Z. Zhang, An Efficient and Versatile Scheduling Algorithm Based on SDC Formulation, 2006 43rd ACM/IEEE Design Automation Conference, pp. 433–438 (2016).

[5] V. Achter, M. Hellmann, M. Knaust, U. Lang, F. Mayer, L. Nieroda, M. Philippsen, O. Rodikow, S. Sadasivam, J. Schwarzer, H. Schwichtenberg, Th. Steinke, ORKA-HPC. https://orka-hpc.uni-koeln.de

[6] F. Mayer, M. Knaust, M. Philippsen, OpenMP on FPGAs – a Survey, OpenMP: Conquering the Full Hardware Spectrum, pp. 94–108 (2019).

[7] M. Knaust, F. Mayer, T. Steinke, OpenMP to FPGA Offloading Prototype Using OpenCL SDK, 2019 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), pp. 387–390 (2019).

# PUBLICATIONS

## PEER-REVIEWED (154)

Aennes Abbas, Ilona Schneider, Anna Bollmann, Jan Funke, Jörg Oehlmann, Carsten Prasse, Ulrike Schulte-Oehlmann, Wolfram Seitz, Thomas Ternes, Marcus Weber, Henning Wesely, Martin Wagner (2019). What you extract is what you see: Optimising the preparation of water and wastewater samples for in vitro bioassays. Water Research, 152:47–60. https://doi.org/10.1016/j.watres.2018.12.049

Joshua L. Abbott, Arthur V. Straube, Dirk G. A. L. Aarts, Roel P. A. Dullens (2019). Transport of a colloidal particle driven across a temporally oscillating optical potential energy landscape. New J. Phys., 21:083027. https://doi.org/10.1088/1367-2630/ab3765

Tobias Achterberg, Robert E. Bixby, Zonghao Gu, Edward Rothberg, Dieter Weninger (2019). Presolve Reductions in Mixed Integer Programming. INFORMS Journal on Computing. (epub ahead of print 2019-11-08)

M. Alchikh, Tim Conrad, X. Ma, E. Broberg, P. Penttinen, J. Reiche, B. Biere, B. Schweiger, B. Rath, Ch. Hoppe (2019). Are we missing respiratory viral infections in infants and children? Comparison of a hospital-based quality management system with standard of care. Clinical Microbiology and Infection, 25(3):380.e9–380.e16.

Gustavo Alonso, Carsten Binnig, Ippokratis Pandis, Kenneth Salem, Jan Skrzypczak, Ryan Stutsman, Lasse Thostrup, Tianzheng Wang, Zeke Wang, Tobias Ziegler (2019). DPI: The Data Processing Interface for Modern Networks. In 9th Biennial Conference on Innovative Data Systems Research.

Felix Ambellan, Hans Lamecker, Christoph von Tycowicz, Stefan Zachow (2019). Statistical Shape Models – Understanding and Mastering Variation in Anatomy. In Paul M. Rea, editor, Biomedical Visualisation, volume 3 of Advances in Experimental Medicine and Biology, pp. 67–84, Springer Nature Switzerland AG. https://doi.org/10.1007/978-3-030-19385-0_5

Felix Ambellan, Stefan Zachow, Christoph von Tycowicz (2019). An as-invariant-as-possible GL+(3)-based Statistical Shape Model. In Proc. 7th MICCAI workshop on Mathematical Foundations of Computational Anatomy (MFCA), volume 11846 of Lecture Notes in Computer Science, pp. 219–228. https://doi.org/10.1007/978-3-030-33226-6_23

Felix Ambellan, Stefan Zachow, Christoph von Tycowicz (2019). A Surface-Theoretic Approach for Statistical Shape Modeling. In Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI), Part IV, volume 11767 of Lecture Notes in Computer Science, pp. 21–29. https://doi.org/10.1007/978-3-030-32251-9_3

Felix Ambellan, Alexander Tack, Moritz Ehlke, Stefan Zachow (2019). Automated Segmentation of Knee Bone and Cartilage combining Statistical Shape Knowledge and Convolutional Neural Networks: Data from the Osteoarthritis Initiative. Medical Image Analysis, 52(2):109–118. https://doi.org/10.1016/j.media.2018.11.009

Lovis Anderson, Benjamin Hiller (2019). A Sweep-Plane Algorithm for the Computation of the Volume of a Union of Polytopes. In Operations Research Proceedings 2018, Operations Research Proceedings. https://doi.org/10.1007/978-3-030-18500-8_12

Anna Andrle, Philipp Hönicke, Philipp-Immanuel Schneider, Yves Kayser, Martin Hammerschmidt, Sven Burger, Frank Scholze, Burkhard Beckhoff, Victor Soltwisch (2019). Grazing incidence x-ray fluorescence based characterization of nanostructures for element sensitive profile reconstruction. Proc. SPIE, 11057:110570M. https://doi.org/10.1117/12.2526082

Bahareh Banyassady (2019). Routing in polygonal domains. Computational Geometry, Theory and Applications. https://doi.org/10.1016/j.comgeo.2019.101593 (epub ahead of print 2019-11-13)

Daniel Baum, James C. Weaver, Igor Zlotnikov, David Knötel, Lara Tomholt, Mason N. Dean (2019). High-Throughput Segmentation of Tiled Biological Structures using Random-Walk Distance Transforms. Integrative And Comparative Biology. https://doi.org/10.1093/icb/icz117

Christiane Becker, Phillip Manley, Klaus Jäger, David Eisenhauer, Johannes Sutter, Steve Albrecht, Kaibo Zheng, Tonu Pullerits, Sven Burger (2019). Nanophotonic enhanced perovskite-silicon solar cell devices. In Said Zhoudi, Antonio Topa, editors, The 10th International Conference on Metamaterials, Photonic Crystals and Plasmonics (META 2019), 858.

Peter Benner, Sara Grundel, Christian Himpe, Christoph Huck, Tom Streubel, Caren Tischendorf (2019). Gas Network Benchmark Models. In Applications of Differential-Algebraic Equations: Examples and Benchmarks, Differential-Algebraic Equations Forum book series (DAEF), pp. 171–197, Springer International Publishing. https://doi.org/10.1007/11221_2018_5

Timo Berthold, Boris Grimm, Markus Reuther, Stanley Schade, Thomas Schlechte (2019). Strategic Planning of Rolling Stock Rotations for Public Tenders. In Proceedings of the 8th International Conference on Railway Operations Modelling and Analysis – RailNorrköping 2019, Linköping Electronic Conference Proceedings (069):148–159.

Timo Berthold, Peter J. Stuckey, Jakob Witzig (2019). Local Rapid Learning for Integer Programs. In Integration of AI and OR Techniques in Constraint Programming. CPAIOR 2019, volume 11494 of Lecture Notes in Computer Science, pp. 67–83. https://doi.org/10.1007/978-3-030-19212-9_5

Daniel Beulertz, Sandrine Charousset, Dieter Most, Spyros Giannelos, Inci Yueksel-Erguen (2019). Development of a Modular Framework for Future Energy System Analysis. In 54th International Universities Power Engineering Conference (UPEC). https://doi.org/10.1109/UPEC.2019.8893472

Felix Binkowski, Lin Zschiedrich, Philipp-Immanuel Schneider, Martin Hammerschmidt, Xavier Garcia Santiago, Fridtjof Betz, Sven Burger (2019). Spectral Expansion of the Scattering Response of Resonant Nanostructures. In Kimmo Saastamoinen, editor, EOS Topical Meeting on Diffractive Optics, 1.

Felix Binkowski, Phlipp-Immanuel Schneider, Martin Hammerschmidt, Lin Zschiedrich, Sven Burger (2019). Numerical optimization of resonant nanophotonic devices. In Milica Matijević, Marko Krstić, Petra Beličev, editors, The Seventh International School and Conference on Photonics (PHOTONICA Belgrade), 10.

Felix Binkowski, Lin Zschiedrich, Martin Hammerschmidt, Sven Burger (2019). Modal analysis for nanoplasmonics with nonlocal material properties. Phys. Rev. B, 100:155406. https://doi.org/10.1103/PhysRevB.100.155406

Felix Binkowski, Lin Zschiedrich, Sven Burger (2019). Computing Optical Resonances in Plasmonic Systems Using a Contour Integral Method. In Rita Asquini, editor, Photonics & Electromagnetics Research Symposium Abstracts (PIERS Rome), 1068.

Felix Binkowski, Lin Zschiedrich, Sven Burger (2019). An auxiliary field approach for computing optical resonances in dispersive media. J. Eur. Opt. Soc.-Rapid, 15:3. https://doi.org/10.1186/s41476-019-0098-z

Ralf Borndörfer, Niels Lindner, Sarah Roth (2019). A Concurrent Approach to the Periodic Event Scheduling Problem. Journal of Rail Transport Planning & Management, 100175. https://doi.org/10.1016/j.jrtpm.2019.100175 (epub ahead of print 2019-12-14)

Ralf Borndörfer, Boris Grimm, Thomas Schlechte (2019). Re-optimizing ICE Rotations after a Tunnel Breakdown near Rastatt. In Proceedings of the 8th International Conference on Railway Operations Modelling and Analysis – RailNorrköping 2019, Linköping Electronic Conference Proceedings(069):160–168. (epub ahead of print 2019-09-13).

Ralf Borndörfer, Boris Grimm, Markus Reuther, Thomas Schlechte (2019). Optimization of handouts for rolling stock rotations. Journal of Rail Transport Planning & Management, 1–8. https://doi.org/10.1016/j.jrtpm.2019.02.001

Ralf Borndörfer, Heide Hoppmann, Marika Karbstein, Niels Lindner (2019). Separation of cycle inequalities in periodic timetabling. Discrete Optimization, 100552. https://doi.org/10.1016/j.disopt.2019.100552 (epub ahead of print 2019-08-20)

Jens Buchmann, Bernhard Kaplan, Samuel Powell, Steffen Prohaska, Jan Laufer (2020). Quantitative PA tomography of high resolution 3-D images: experimental validation in tissue phantoms. Photoacoustics, 17:100157. https://doi.org/10.1016/j.pacs.2019.100157 (accepted for publication 2019-12-05)

Jens Buchmann, Bernhard A. Kaplan, Samuel Powell, Steffen Prohaska, Jan Laufer (2019). 3D quantitative photoacoustic tomography using an adjoint radiance Monte Carlo model and gradient descent. Journal of Biomedical Optics, 24(6):066001. https://doi.org/10.1117/1.JBO.24.6.066001

Sven Burger, Philipp-Immanuel Schneider, Theresa Höhne, Lin Zschiedrich (2019). Numerical methods for simulation and optimization of the extraction efficiency from quantum-dot based single-photon emitters. In Stephan Reitzenstein, Tobias Heindel, editors, 7th International Workshop on Engineering of Quantum Emitter Properties (EQEP), 20.

Sven Burger, Felix Binkowski, Lin Zschiedrich (2019). Finite Element Simulations of Optical Resonances in Dispersive Nanoresonators. In Rita Asquini, editor, Photonics & Electromagnetics Research Symposium Abstracts (PIERS Rome), 699.

Sven Burger, Felix Binkowski, Lin Zschiedrich (2019). Computing resonances in nano-photonic devices using Riesz-projection-based methods. In Said Zhoudi, Antonio Topa, editors, The 10th International Conference on Metamaterials, Photonic Crystals and Plasmonics (META 2019), 92.

Robert Clausecker, Alexander Reinefeld (2019). Zero-Aware Pattern Databases with 1-Bit Compression for Sliding Tile Puzzles. In Proceedings of the Twelfth International Symposium on Combinatorial Search (SoCS 2019), 35–43.

Jon Cockayne, Chris Oates, T. J. Sullivan, Mark Girolami (2019). Bayesian Probabilistic Numerical Methods. SIAM Review, 61(4):756–789. https://doi.org/10.1137/17M1139357

Manuel Dibak, Christoph Fröhner, Frank Noé, Felix Höfling (2019). Diffusion-influenced reaction rates in the presence of pair interactions. The Journal of Chemical Physics, 151:164105. https://doi.org/10.1063/1.5124728

Rainald M. Ehrig, Markus O. Heller (2019). On intrinsic equivalences of the finite helical axis, the instantaneous helical axis, and the SARA approach. A mathematical perspective. Journal of Biomechanics, 84:4–10. https://doi.org/10.1016/j.jbiomech.2018.12.034

Franziska Erlekam, Sinaida Igde, Susanna Röblitz, Laura Hartmann, Marcus Weber (2019). Modeling of Multivalent Ligand-Receptor Binding Measured by kinITC. Computation, 7(3):46. https://doi.org/10.3390/computation7030046

Natalia Ernst, Konstantin Fackeldey, Andrea Volkamer, Oliver Opatz, Marcus Weber (2019). Computation of temperature-dependent dissociation rates of metastable protein–ligand complexes. Molecular Simulation, 45(11):904–911. https://doi.org/10.1080/08927022.2019.1610949

Ricardo Euler, Ralf Borndörfer (2019). A Graph – and Monoid-based Framework for Price-Sensitive Routing in Local Public Transportation Networks. In 19th Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2019), volume 75 of OpenAccess Series in Informatics (OASIcs), pp. 12:1–12:15. https://doi.org/10.4230/OASIcs.ATMOS.2019.12

Konstantin Fackeldey, Peter Koltai, Peter Nevir, Henning Rust, Axel Schild, Marcus Weber (2019). From metastable to coherent sets – Time-discretization schemes. Chaos: An Interdisciplinary Journal of Nonlinear Science, 29:012101–012101. https://doi.org/10.1063/1.5058128

Nando Farchmin, Martin Hammerschmidt, Philipp-Immanuel Schneider, Matthias Wurm, Markus Bär, Sebastian Heidenreich (2019). Efficient global sensitivity analysis for silicon line gratings using polynomial chaos. Proc. SPIE, 11057:110570J. https://doi.org/10.1117/12.2525978

Bernhard Fröhler, Tim Elberfeld, Torsten Möller, Hans-Christian Hege, Johannes Weissenböck, Jan De Beenhouwer, Jan Sijbers, Johann Kastner, Christoph Heinzl (2019). A Visual Tool for the Analysis of Algorithms for Tomographic Fiber Reconstruction in Materials Science. Computer Graphics Forum, 38(3):273–283. https://doi.org/10.1111/cgf.13688

Bernhard Fröhler, Lucas da Cunha Melo, Johannes Weissenböck, Johann Kastner, Torsten Möller, Hans-Christian Hege, Eduard M. Gröller, Jonathan Sanctorum, Jan De Beenhouwer, Jan Sijbers, Christoph Heinzl (2019). Tools for the analysis of datasets from X-ray computed tomography based on Talbot-Lau grating interferometry. In Proceedings of iCT 2019, (9th Conference on Industrial Computed Tomography, Padova, Italy – iCT 2019, February 13-15, 2019).

Fabio Furini, Emiliano Traversi, Pietro Belotti, Antonio Frangioni, Ambros Gleixner, Nick Gould, Leo Liberti, Andrea Lodi, Ruth Misener, Hans Mittelmann, Nikolaos V. Sahinidis, Stefan Vigerske, Angelika Wiegele (2019). QPLIB: A Library of Quadratic Programming Instances. Mathematical Programming Computation, 11(2):237–265. https://doi.org/10.1007/s12532-018-0147-4

Gerald Gamrath, Timo Berthold, Stefan Heinz, Michael Winkler (2019). Structure-driven fix-and-propagate heuristics for mixed integer programming. Mathematical Programming Computation, 1–24. https://doi.org/10.1007/s12532-019-00159-1

Patrick Gelß, Stefan Klus, Jens Eisert, Christof Schütte (2019). Multidimensional Approximation of Nonlinear Dynamical Systems. Journal of Computational and Nonlinear Dynamics, 14(6). https://doi.org/10.1115/1.4043148

Masoud Gholami, Florian Schintke (2019). Multilevel Checkpoint/Restart for Large Computational Jobs on Distributed Computing Resources. In 2019 IEEE 38th Symposium on Reliable Distributed Systems (SRDS). (accepted for publication 2019-06-24)

Ambros Gleixner, Nils-Christian Kempke, Thorsten Koch, Daniel Rehfeldt, Svenja Uslu (2019). First Experiments with Structure-Aware Presolving for a Parallel Interior-Point Method. In Operations Research 2019 Proceedings. (accepted for publication 2019-12-13)

Ambros Gleixner, Daniel E. Steffy (2019). Linear Programming using Limited-Precision Oracles. Mathematical Programming. https://doi.org/10.1007/s10107-019-01444-6 (epub ahead of print 2019-11-19)

Ambros Gleixner, Daniel E. Steffy (2019). Linear Programming using Limited-Precision Oracles. A. Lodi, V. Nagarajan (eds), Integer Programming and Combinatorial Optimization: 20th International Conference, IPCO 2019, 399–412. https://doi.org/10.1007/978-3-030-17953-3_30

Uwe Gotzes (2019). Ein neuer Ansatz zur Optimierung des Bilanzausgleichs in einem Gasmarktgebiet. Zeitschrift für Energiewirtschaft. https://doi.org/10.1007/s12398-019-00257-6

Leonid Goubergrits, Florian Hellmeier, Jan Bruening, Andreas Spuler, Hans-Christian Hege, Samuel Voss, Gábor Janiga, Sylvia Saalfeld, Oliver Beuing, Philipp Berg (2019). Multiple Aneurysms AnaTomy CHallenge 2018 (MATCH): Uncertainty Quantification of Geometric Rupture Risk Parameters. BioMedical Engineering OnLine, 18(35). https://doi.org/10.1186/s12938-019-0657-y

Florian Graf, Joshua Feis, Xavier Garcia Santiago, Martin Wegener, Carsten Rockstuhl, Ivan Fernandez-Corbaton (2019). Achiral, Helicity Preserving, and Resonant Structures for Enhanced Sensing of Chiral Molecules. ACS Photonics, 6:482. https://doi.org/10.1021/acsphotonics.8b01454

Boris Grimm, Ralf Borndörfer, Markus Reuther, Thomas Schlechte (2019). A Cut Separation Approach for the Rolling Stock Rotation Problem with Vehicle Maintenance. In Valentina Cacchiani, Alberto Marchetti-Spaccamela, editors, 19th Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2019), volume 75 of OpenAccess Series in Informatics (OASIcs), pp. 1:1–1:12. https://doi.org/10.4230/OASIcs.ATMOS.2019.1 (epub ahead of print 2019-09-14)

Boris Grimm, Ralf Borndörfer, Mats Olthoff (2019). A Solution Approach to the Vehicle Routing Problem with Perishable Goods. In Operations Research 2019 Proceedings. (accepted for publication 2019-12-13)

Pooja Gupta, Sarah Peter, Markus Jung, Astrid Lewin, Georg Hemmrich-Stanisak, Andre Franke, Max von Kleist, Christof Schütte, Ralf Einspanier, Soroush Sharbati, Jennifer zur Bruegge (2019). Analysis of long non-coding RNA and mRNA expression in bovine macrophages brings up novel 2 aspects of Mycobacterium avium subspecies paratuberculosis infections. Scientific Reports in Nature, 9. https://doi.org/10.1038/s41598-018-38141-x

Sebastian Götschel, Martin Weiser (2019). Compression Challenges in Large Scale Partial Differential Equation Solvers. Algorithms, 12(9):197. https://doi.org/10.3390/a12090197

Sebastian Götschel, Michael L. Minion (2019). An Efficient Parallel-in-Time Method for Optimization with Parabolic PDEs. SIAM J. Sci. Comput., 41(6):C603–C626. https://doi.org/10.1137/19M1239313

Tobias Haase, Vikram Sunkara, Benjamin Kohl, Carola Meier, Patricia Bußmann, Jessica Becker, Michal Jagielski, Max von Kleist, Wolfgang Ertel (2019). Discerning the spatio-temporal disease patterns of surgically induced OA mouse models, 14(4). https://doi.org/10.1371/journal.pone.0213734

Publications


Hassan Al Hajj, Mathieu Lamard, Pierre-Henri Conze, Soumali Roychowdhury, Xiaowei Hu, Gabija Marsalkaite, Odysseas Zisimopoulos, Muneer Ahmad Dedmari, Fenqiang Zhao, Jonas Prellberg, Manish Sahu, Adrian Galdran, Teresa Araujo, Duc My Vo, Chandan Panda, Navdeep Dahiya, Satoshi Kondo, Zhengbing Bian, Jonas Bialopetravicius, Chenghui Qiu, Sabrina Dill, Anirban Mukhopadyay, Pedro Costa, Guilherme Aresta, Senthil Ramamurthy, Sang-Woong Lee, Aurelio Campilho, Stefan Zachow, Shunren Xia, Sailesh Conjeti, Jogundas Armaitis, Pheng-Ann Heng, Arash Vahdat, Beatrice Cochener, Gwenole Quellec (2019). CATARACTS: Challenge on Automatic Tool Annotation for cataRACT Surgery. Medical Image Analysis, 52(2):24–41. https://doi.org/10.1016/j.media.2018.11.008

Martin Hammerschmidt, Lin Zschiedrich, Philipp-Immanuel Schneider, Felix Binkowski, Sven Burger (2019). Numerical optimization of resonant photonic devices. In Proc. SPIE, 11057:1105702. https://doi.org/10.1117/12.2534348

Carsten Hartmann, Christof Schütte, Wei Zhang (2019). Jarzynski's equality, fluctuation theorems, and variance reduction: Mathematical analysis and numerical algorithms. Journal of Statistical Physics, 175(6):1214–1261. https://doi.org/10.1007/s10955-019-02286-4

Marc Hartung, Florian Schintke (2019). Learned Clause Minimization in Parallel SAT Solvers. Pragmatics of SAT 2019, 1–11.

Marc Hartung, Florian Schintke, Thorsten Schütt (2019). Pinpoint Data Races via Testing and Classification. In 2019 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW); 3rd International Workshop on Software Faults (IWSF 2019), 386–393. https://doi.org/10.1109/ISSREW.2019.00100

Thomas Hildebrandt, Jan Joris Bruening, Hans Lamecker, Stefan Zachow, Werner Heppt, Nora Schmidt, Leonid Goubergrits (2019). Digital Analysis of Nasal Airflow Facilitating Decision Support in Rhinosurgery. Facial Plastic Surgery, 35(1):1–8. https://doi.org/10.1055/s-0039-1677720

Thomas Hildebrandt, Jan Joris Bruening, Nora Laura Schmidt, Hans Lamecker, Werner Heppt, Stefan Zachow, Goubergrits Leonid (2019). The Healthy Nasal Cavity – Characteristics of Morphology and Related Airflow Based on a Statistical Shape Model Viewed from a Surgeon's Perspective. Facial Plastic Surgery, 35(1):9–13. https://doi.org/10.1055/s-0039-1677721

Theresa Höhne, Peter Schnauber, Sven Rodt, Stephan Reitzenstein, Sven Burger (2019). Numerical Investigation of Light Emission from Quantum Dots Embedded into On-Chip, Low Index Contrast Optical Waveguides. Phys. Status Solidi B, 256:1800437. https://doi.org/10.1002/pssb.201800437

Theresa Höhne, Peter Schnauber, Sven Rodt, Stephan Reitzenstein, Sven Burger (2019). Numerical Studies of Integrated Single-Photon Emitters in Low-Index-Contrast Waveguides. In Stephan Reitzenstein, Tobias Heindel, editors, 7th International Workshop on Engineering of Quantum Emitter Properties (EQEP), 49.

Theresa Höhne, Peter Schnauber, Sven Rodt, Stephan Reitzenstein, Sven Burger (2019). Numerical Investigation of Integrated Single-Photon Emitters in Low-Index-Contrast Waveguides. In Lukas Novotny, Romain Quidant, editors, Nanophotonics: Foundations & Applications, CSF Conference Book of Abstracts, B11.

Sahar Iravani, Tim Conrad (2019). Deep Learning for Proteomics Data for Feature Selection and Classification. In A. Holzinger, P. Kieseberg, A. Tjoa, E. Weippl, editors, Lecture Notes in Computer Science, volume 11713 of Machine Learning and Knowledge extraction. https://doi.org/10.1007/978-3-030-29726-8_19

Robert Joachimsky, Lihong Ma, Christian Icking, Stefan Zachow (2019). A Collision-Aware Articulated Statistical Shape Model of the Human Spine. Proc. of the 18th annual conference on Computer – and Robot-assisted Surgery (CURAC).

Klaus Jäger, Marko Jošt, Johannes Sutter, Philipp Tockhorn, Eike Köhnen, David Eisenhauer, Phillip Manley, Steve Albrecht, Christiane Becker (2019). Improving Monolithic Perovskite/Silicon Tandem Solar Cells From an Optical Viewpoint. In OSA Advanced Photonics Congress, OSA Technical Digest, p. PM4C.2. https://doi.org/10.1364/PVLED.2019.PM4C.2

Günter Kewes, Felix Binkowski, Sven Burger, Lin Zschiedrich, Felix Stete, Wouter Koopman, Matias Bargheer, Oliver Benson (2019). Line Broadening Mechanisms in Hybrid Plasmonic Systems for Strong Coupling. In Rita Asquini, editor, Photonics & Electromagnetics Research Symposium Abstracts (PIERS Rome), 704.

Stefan Klus, Brooke E. Husic, Mattes Mollenhauer, Frank Noe (2019). Kernel methods for detecting coherent structures in dynamical data. Chaos: An Interdisciplinary Journal of Nonlinear Science, 29(12). https://doi.org/10.1063/1.5100267

Marius Knaust, Florian Mayer, Thomas Steinke (2019). OpenMP to FPGA Offloading Prototype Using OpenCL SDK. 2019 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), 387–390. https://doi.org/10.1109/IPDPSW.2019.00072

Péter Koltai, Han Cheng Lie, Martin Plonka (2019). Fréchet differentiable drift dependence of Perron–Frobenius and Koopman operators for non-deterministic dynamics. Nonlinearity, 32(11):4232–4257. https://doi.org/10.1088/1361-6544/ab1f2a

Tobias Kramer, Matthias Läuter (2019). Outgassing induced acceleration of comet 67P/Churyumov-Gerasimenko. Astronomy & Astrophysics, 630:A4. https://doi.org/10.1051/0004-6361/201935229

Tobias Kramer, Matthias Läuter, Stubbe Hviid, Laurent Jorda, Horst Uwe Keller, Ekkehard Kührt (2019). Comet 67P/Churyumov-Gerasimenko rotation changes derived from sublimation induced torques. Astronomy & Astrophysics, 630:A3. https://doi.org/10.1051/0004-6361/201834349

Tobias Kramer, Mirta Rodríguez (2019). Effect of disorder and polarization sequences on two-dimensional spectra of light harvesting complexes. Photosynthesis Research. https://doi.org/10.1007/s11120-019-00699-6

Martin Krämer, Marta Maggioni, Nicholas Brisson, Stefan Zachow, Ulf Teichgräber, Georg Duda, Jürgen Reichenbach (2019). T1 and T2* mapping of the human quadriceps and patellar tendons using ultra-short echo-time (UTE) imaging and bivariate relaxation parameter-based volumetric visualization. Magnetic Resonance Imaging, 63(11):29–36. https://doi.org/10.1016/j.mri.2019.07.015

P. Lalanne, W. Yan, A. Gras, C. Sauvan, J.-P. Hugonin, M. Besbes, G. Demésy, M. D. Truong, B. Gralak, F. Zolla, A. Nicolet, F. Binkowski, L. Zschiedrich, S. Burger, J. Zimmerling, R. Remis, P. Urbach, H. T. Liu, T. Weiss (2019). Quasinormal mode solvers for resonators with dispersive materials. J. Opt. Soc. Am. A, 36:686. https://doi.org/10.1364/JOSAA.36.000686

Annemarie Lang, Lisa Fischer, Marie-Christin Weber, Timo Gaber, Rainald Ehrig, Susanna Röblitz, Frank Buttgereit (2019). Combining in vito simulation and in silico modelling towards a sophisticated human osteoarthritis model. Osteoarthritis and Cartilage, 27:183. https://doi.org/10.1016/j.joca.2019.02.277

Tony Lelièvre, Wei Zhang (2019). Pathwise estimates for effective dynamics: the case of nonlinear vectorial reaction coordinates. Multiscale Modeling and Simulation, 1019–1051. https://doi.org/10.1137/18M1186034

Han Cheng Lie, T. J. Sullivan, Andrew Stuart (2019). Strong convergence rates of probabilistic integrators for ordinary differential equations. Statistics and Computing, 29(6):1265–1283. https://doi.org/10.1007/s11222-019-09898-6

Niels Lindner, Christian Liebchen (2019). New Perspectives on PESP: T-Partitions and Separators. In Valentina Cacchiani, Alberto Marchetti-Spaccamela, editors, 19th Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2019), volume 75 of OpenAccess Series in Informatics (OASIcs), pp. 2:1–2:18. https://doi.org/10.4230/OASIcs.ATMOS.2019.2

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

Matthias Läuter, Tobias Kramer, Martin Rubin, Kathrin Altwegg (2019). Surface localization of gas sources on comet 67P/Churyumov-Gerasimenko based on DFMS/COPS data. Monthly Notices of the Royal Astronomical Society, 483:852–861. https://doi.org/10.1093/mnras/sty3103

Heinz-Eberhard Mahnke, Tobias Arlt, Daniel Baum, Hans-Christian Hege, Felix Herter, Norbert Lindow, Ingo Manke, Tzulia Siopi, Eve Menei, Marc Etienne, Verena Lepper (2020). Virtual unfolding of folded papyri. Journal of Cultural Heritage, 41:264–269. https://doi.org/10.1016/j.culher.2019.07.007 (epub ahead of print 2019-07-30)

Phillip Manley, Martin Hammerschmidt, Sven Burger, Christiane Becker (2019). Nanophotonic Enhancement of Light Out-Coupling for Deep-UV LEDs. In Said Zhoudi, Antonio Topa, editors, The 10th International Conference on Metamaterials, Photonic Crystals and Plasmonics (META 2019), 745.

Florian Mayer, Marius Knaust, Michael Philippsen (2019). OpenMP on FPGAs—A Survey. OpenMP: Conquering the Full Hardware Spectrum, 94–108. https://doi.org/10.1007/978-3-030-28596-8_7

Gunther Mohr, Simon J. Altenburg, Alexander Ulbricht, Philipp Heinrich, Daniel Baum, Christiane Maierhofer, Kai Hilgenberg (2019). In-situ defect detection in laser powder bed fusion by using thermography and optical tomography – comparison to computed tomography. Metals. (accepted for publication 2019-12-20)

Pawel Mrowinski, Peter Schnauber, Philipp Gutsche, Arsenty Kaganskiy, Johannes Schall, Sven Burger, Sven Rodt, Stephan Reitzenstein (2019). Directional emission of a deterministically fabricated quantum dot – Bragg reflection multi-mode waveguide system. ACS Photonics, 6:2231. https://doi.org/10.1021/acsphotonics.9b00369

Lluis Miquel Munguia, Geoffrey Oxberry, Deepak Rajan, Yuji Shinano (2019). Parallel PIPS-SBB: Multi-Level Parallelism For Stochastic Mixed-Integer Programs. Computational Optimization and Applications. https://doi.org/10.1007/s10589-019-00074-0 (epub ahead of print 2019-02-15)

Benjamin Müller, Felipe Serrano, Ambros Gleixner (2019). Using two-dimensional Projections for Stronger Separation and Propagation of Bilinear Terms. Accepted for SIAM Journal on Optimization.

Esfandiar Nava-Yazdani, Hans-Christian Hege, Christoph von Tycowicz (2019). A Geodesic Mixed Effects Model in Kendall's Shape Space. In Proc. 7th MICCAI workshop on Mathematical Foundations of Computational Anatomy (MFCA), volume 11846 of Lecture Notes in Computer Science, pp. 209–218. https://doi.org/10.1007/978-3-030-33226-6_22

Mario Neumann, Olaf Hellwich, Stefan Zachow (2019). Localization and Classification of Teeth in Cone Beam CT using Convolutional Neural Networks. Proc. of the 18th annual conference on Computer – and Robot-assisted Surgery (CURAC).

Matthias Noack, Erich Focht, Thomas Steinke (2019). Heterogeneous Active Messages for Offloading on the NEC SX-Aurora TSUBASA. In 2019 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), Heterogeneity in Computing Workshop (HCW 2019).

Matthias Noack (2019). Heterogeneous Active Messages (HAM) – Implementing Lightweight Remote Procedure Calls in C++. In Proceedings of the 5th International Workshop on OpenCL, The Distributed & Heterogeneous Programming in C/C++ (DHPCC++ 2019) Conference. https://doi.org/10.1145/3318170.3318195

John Nyakatura, Roxane Baumgarten, Daniel Baum, Heiko Stark, Dionisios Youlatos (2019). Muscle internal structure revealed by contrast-enhanced μCT and fibre recognition: The hindlimb extensors of an arboreal and a fossorial squirrel. Mammalian Biology, 99:71–80. https://doi.org/10.1016/j.mambio.2019.10.007 (epub ahead of print 2019-11-02)

Chris. J. Oates, T. J. Sullivan (2019). A modern retrospective on probabilistic numerics. Statistics and Computing, 29(6):1335–1351. https://doi.org/10.1007/s11222-019-09902-z

Steffen Oeltze-Jaffra, Monique Meuschke, Matthias Neugebauer, Sylvia Saalfeld, Kai Lawonn, Gabor Janiga, Hans-Christian Hege, Stefan Zachow, Bernhard Preim (2019). Generation and Visual Exploration of Medical Flow Data: Survey, Research Trends, and Future Challenges. Computer Graphics Forum, 38(1):87–125. https://doi.org/10.1111/cgf.13394

Elias Oltmanns, Tim Hasler, Wolfgang Peters-Kottig, Heinz-Günter Kuper (2019). Different Preservation Levels: The Case of Scholarly Digital Editions. Data Science Journal, 18(1(51)). https://doi.org/10.5334/dsj-2019-051

* Mehmet Neset Ozel, Abhishek Kulkarni, Amr Hasan, Josephine Brummer, Marian Moldenhauer, Ilsa-Maria Daumann, Heike Wolfenberg, Vincent Dercksen, Ferdi Ridvan Kiral, Martin Weiser, Steffen Prohaska, Max von Kleist, Peter Robin Hiesinger (2019). Serial synapse formation through filopodial competition for synaptic seeding factors. Developmental Cell, 50(4):447–461. https://doi.org/10.1016/j.devcel.2019.06.014 (Joint publication: Numerical Mathematics, Visual Data Analysis)

Olaf Paetsch (2019). Possibilities and Limitations of Automatic Feature Extraction shown by the Example of Crack Detection in 3D-CT Images of Concrete Specimen. In iCT 2019, iCT 2019 Conference Proceedings.

Sebastian Pokutta, M. Singh, A. Torrico (2019). On the Unreasonable Effectiveness of the Greedy Algorithm: Greedy Adapts to Sharpness. In OPTML Workshop Paper.

Jonad Pulaj (2019). Cutting planes for families implying Frankl's conjecture. Mathematics of Computation. https://doi.org/10.1090/mcom/3461 (epub ahead of print 2019-07-25)

Daniel Rehfeldt, Yuji Shinano, Thorsten Koch (2019). SCIP-Jack: An exact high performance solver for Steiner tree problems in graphs and related problems. In Modeling, Simulation and Optimization of Complex Processes HPSC 2018, Lecture Notes in Computer Science. (accepted for publication 2019-05-28)

Daniel Rehfeldt, Thorsten Koch (2019). Combining NP-Hard Reduction Techniques and Strong Heuristics in an Exact Algorithm for the Maximum-Weight Connected Subgraph Problem. SIAM Journal on Optimization, 29(1):369–398. https://doi.org/10.1137/17M1145963

Daniel Rehfeldt, Thorsten Koch, Stephen Maher (2019). Reduction Techniques for the Prize-Collecting Steiner Tree Problem and the Maximum-Weight Connected Subgraph Problem. Networks, 73:206–233. https://doi.org/10.1002/net.21857

Anika Rettig, Tobias Haase, Alexandr Pletnyov, Benjamin Kohl, Wolfgang Ertel, Max von Kleist, Vikram Sunkara (2019). SLCV-a supervised learning-computer vision combined strategy for automated muscle fibre detection in cross-sectional images. PeerJ. 2019;7:e7053. doi:10.7717/peerj.7053.

Bernhard Reuter, Konstantin Fackeldey, Marcus Weber (2019). Generalized Markov modeling of nonreversible molecular kinetics. The Journal of Chemical Physics, 17(150):174103. https://doi.org/10.1063/1.5064530

Mirta Rodríguez, Tobias Kramer (2019). Machine Learning of Two-Dimensional Spectroscopic Data. Chemical Physics, 520:52–60. https://doi.org/10.1016/j.chemphys.2019.01.002

Beate Rusch, Julia Boltze, Thomas Dierkes, Julia Alexandra Goltz-Fellgiebel, Hedda Staub (2019). DeepGreen – DeepGold: Open-Access-Transformation – Entwicklung und Perspektiven. GMS Medizin – Bibliothek – Information, 19(1–2). https://doi.org/10.3205/mbi000432

Guillaume Sagnol, Edouard Pauwels (2019). An unexpected connection between Bayes A-optimal designs and the group lasso. Statistical Papers, 60(2):215–234. https://doi.org/10.1007/s00362-018-01062-y

Daisuke Sakurai, Kenji Ono, Hamish Carr, Jorji Nonaka, Tomohiro Kawanabe (2019). Flexible Fiber Surfaces: A Reeb-Free Approach. In Hamish Carr, Issei Fujishiro, Filip Sadlo, Shigeo Takahashi, editors, Topological Methods in Data Analysis and Visualization V.

Farouk Salem, Thorsten Schütt, Florian Schintke, Alexander Reinefeld (2019). Scheduling Data Streams for Low Latency and High Throughput on a Cray XC40 Using Libfabric. In CUG Conference Proceedings.

Farouk Salem, Florian Schintke, Thorsten Schütt, Alexander Reinefeld (2019). Scheduling data streams for low latency and high throughput on a Cray XC40 using Libfabric. Concurrency and Computation Practice and Experience, 1–14. https://doi.org/10.1002/cpe.5563

Xavier Garcia Santiago, Martin Hammerschmidt, Sven Burger, Carsten Rockstuhl, Ivan Fernandez-Corbaton, Lin Zschiedrich (2019). Decomposition of scattered electromagnetic fields into vector spherical wave functions on surfaces with general shapes. Phys. Rev. B, 99:045406. https://doi.org/10.1103/PhysRevB.99.045406

Peter Schnauber, Johannes Schall, Samir Bounouar, Theresa Höhne, Suk-In Park, Geun-Hwan Ryu, Tobias Heindel, Sven Burger, Jin-Dong Song, Sven Rodt, Stephan Reitzenstein (2019). Deterministic integration of quantum dots into on-chip multi-mode interference couplers via in-situ electron beam lithography. In Conference on Lasers and Electro-Optics Europe and European Quantum Electronics Conference (CLEO/EQEC), OSA Technical Digest, p. EB2.2. https://doi.org/10.1109/CLEOE-EQEC.2019.8872583

Philipp-Immanuel Schneider, Xavier Garcia Santiago, Victor Soltwisch, Martin Hammerschmidt, Sven Burger, Carsten Rockstuhl (2019). Benchmarking five global optimization approaches for nano-optical shape optimization and parameter reconstruction. ACS Photonics, 6:2726. https://doi.org/10.1021/acsphotonics.9b00706

Philipp-Immanuel Schneider, Martin Hammerschmidt, Lin Zschiedrich, Sven Burger (2019). Using Gaussian process regression for efficient parameter reconstruction. Proc. SPIE, 10959:1095911. https://doi.org/10.1117/12.2513268

Robert Schulz, Kenji Yamamoto, André Klossek, Fiorenza Rancan, Annika Vogt, Christof Schütte, Eckart Rühl, Roland R. Netz (2019). Modeling of Drug Diffusion Based on Concentration Profiles in Healthy and Damaged Human Skin. Biophysical Journal, 117(5):998–1008. https://doi.org/10.1016/j.bpj.2019.07.027

Felipe Serrano (2019). Intersection cuts for factorable MINLP. In A. Lodi, V. Nagarajan (eds), Integer Programming and Combinatorial Optimization: 20th International Conference, IPCO 2019, volume 11480 of Lecture Notes in Computer Science, pp. 385–398. https://doi.org/10.1007/978-3-030-17953-3_29

* Borong Shao, Maria Bjaanaes, Aslaug Helland, Christof Schütte, Tim Conrad (2019). EMT network-based feature selection improves prognosis prediction in lung adenocarcinoma. PLOS ONE, 14(1). https://doi.org/10.1101/410472 (Joint publication: Numerical Mathematics, Visual Data Analysis)

Yuji Shinano, Daniel Rehfeldt, Tristan Gally (2019). An Easy Way to Build Parallel State-of-the-art Combinatorial Optimization Problem Solvers: A Computational Study on Solving Steiner Tree Problems and Mixed Integer Semidefinite Programs by using ug[SCIP-*,*]-libraries. In Proceedings of the 9th IEEE Workshop Parallel / Distributed Combinatorics and Optimization, 530–541. https://doi.org/10.1109/IPDPSW.2019.00095

Yuji Shinano, Daniel Rehfeldt, Thorsten Koch (2019). Building Optimal Steiner Trees on Supercomputers by Using up to 43,000 Cores. In Integration of Constraint Programming, Artificial Intelligence, and Operations Research. CPAIOR 2019, volume 11494 of Lecture Notes in Computer Science, pp. 529–539. https://doi.org/10.1007/978-3-030-19212-9_35

Luigi Delle Site, Christian Krekeler, John Whittaker, Animesh Agarwal, Rupert Klein, Felix Höfling (2019). Molecular Dynamics of Open Systems: Construction of a Mean-Field Particle Reservoir. Advanced Theory and Simulations, 2:1900014. https://doi.org/10.1002/adts.201900014

Jan Skrzypczak, Florian Schintke, Thorsten Schütt (2019). Linearizable State Machine Replication of State-Based CRDTs without Logs. In Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, 455–457. https://doi.org/10.1145/3293611.3331568

Arthur V. Straube, Josep M. Pagès, Pietro Tierno, Jordi Ignés-Mullol, Francesc Sagués (2019). Collective dynamics and conformal ordering in electrophoretically driven nematic colloids. Phys. Rev. Research, 1:022008. https://doi.org/10.1103/PhysRevResearch.1.022008

T. J. Sullivan (2019). Contributed discussion on the article "A Bayesian conjugate gradient method". Bayesian Analysis, 14(3):985–989. https://doi.org/10.1214/19-BA1145

Vikram Sunkara (2019). On the Properties of the Reaction Counts Chemical Master Equation. Entropy, 21(6):607. https://doi.org/10.3390/e21060607

Alexander Tack, Stefan Zachow (2019). Accurate Automated Volumetry of Cartilage of the Knee using Convolutional Neural Networks: Data from the Osteoarthritis Initiative. In IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019), 40–43.

Christine Tawfik, Sabine Limbourg (2019). A bilevel model for network design and pricing based on a level-of-service assessment. Transportation science. (accepted for publication 2019-01-17)

Julia Temp, Domonika Labuz, Rodger Negrete, Vikram Sunkara, Halina Machelska (2019). Pain and knee damage in male and female mice in the medial meniscal transection-induced osteoarthritis. Osteoarthritis and Cartilage. https://doi.org/10.1016/j.joca.2019.11.003 (epub ahead of print 2019-12-10)

* Denise Thiel, Natasa Djurdjevac Conrad, Evgenia Ntini, Ria Peschutter, Heike Siebert, Annalisa Marsico (2019). Identifying lncRNA-mediated regulatory modules via ChIA-PET network analysis. BMC Bioinformatics, 20(1471–2105). https://doi.org/10.1186/s12859-019-2900-8 (Joint publication: Numerical Mathematics, Visual Data Analysis)

Peter Tillmann, Klaus Jäger, Christiane Becker (2020). Minimising levelised cost of electricity of bifacial solar panel arrays using Bayesian optimisation. Sustain. Energy Fuels, 4:254. https://doi.org/10.1039/C9SE00750D (accepted for publication 2019-11-03)

Giovanna Del Vecchio, Dominika Labuz, Julia Temp, Viola Seitz, Michael Kloner, Roger Negrete, Antonio Rodriguez-Gaztelumendi, Marcus Weber, Halina Machelska, Christoph Stein (2019). pKa of opioid ligands as a discriminating factor for side effects. Nature Scientific Reports, 9:19344. https://doi.org/10.1038/s41598-019-55886-1

Narendra L. Venkatareddy, Patrick Wilke, Natalia Ernst, Justus Horch, Marcus Weber, Andre Dallmann, Hans G. Börner (2019). Mussel-glue inspired adhesives: A study on the relevance of L-Dopa and the function of the sequence at nanomaterial-peptide interfaces. Advanced Materials Interfaces, 6(13):1900501. https://doi.org/10.1002/admi.201900501

José Villatoro, Marcus Weber, Martin Zühlke, Andreas Lehmann, Karl Zechiowski, Daniel Riebe, Toralf Beitz, Hans Gerd Löhmannsröben, Oliver Kreuzer (2019). Structural characterization of synthetic peptides using electrospray ion mobility spectrometry and molecular dynamics simulations. International Journal of Mass Spectrometry, 436:108–117. https://doi.org/10.1016/j.ijms.2018.10.036

Tobias Weber, Sebastian Sager, Ambros Gleixner (2019). Solving Quadratic Programs to High Precision using Scaled Iterative Refinement. Mathematical Programming Computation, 11:421–455. https://doi.org/10.1007/s12532-019-00154-6

Marcus Weber (2019). Transformationsprodukte im Klärwerk: Mathematische Ansätze der Bewertung. KA Korrespondenz Abwasser, Abfall, 7:551–557.

C. Weinhold, A. Lackorzynski, J. Bierbaum, M. Küttler, M. Planeta, H. Weisbach, M. Hille, H. Härtig, A. Margolin, D. Sharf, E. Levy, P. Gak, A. Barak, M. Gholami, F. Schintke, T. Schütt, A. Reinefeld, M. Lieber, W. E. Nagel (2019). FFMK – A Fast and Fault-Tolerant Microkernel-Based System for Exascale Computing. In SPPEXA Symposium 2019, (accepted for publication 2019-10-07)

Florian Wende (2019). C++ Data Layout Abstractions through Proxy Types. In 2019 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), 14th International Workshop on Automatic Performance Tunings (iWAPT), 758–767. https://doi.org/10.1109/IPDPSW.2019.00126

Daniel Werdehausen, Sven Burger, Isabelle Staude, Thomas Pertsch, Manuel Decker (2019). Nanocomposites – A Route to better and smaller optical Elements?. In Optical Design and Fabrication 2019, OSA Technical Digest, p. OT2A.2. https://doi.org/10.1364/OFT.2019.OT2A.2

Daniel Werdehausen, Sven Burger, Isabelle Staude, Thomas Pertsch, Manuel Decker (2019). Dispersion-engineered nanocomposites enable achromatic diffractive optical elements. Optica, 6:1031. https://doi.org/10.1364/OPTICA.6.001031

Jon Wilson, Philipp Gutsche, Sven Herrmann, Sven Burger, Kevin McPeak (2019). Correlation of circular differential optical absorption with geometric chirality in plasmonic meta-atoms. Opt. Express, 27:5097. https://doi.org/10.1364/OE.27.005097

Jakob Witzig, Timo Berthold, Stefan Heinz (2019). A Status Report on Conflict Analysis in Mixed Integer Nonlinear Programming. In Integration of AI and OR Techniques in Constraint Programming. CPAIOR 2019, volume 11494 of Lecture Notes in Computer Science, pp. 84–94. https://doi.org/10.1007/978-3-030-19212-9_6

Wei Zhang (2019). Ergodic SDEs on submanifolds and related numerical sampling schemes, ESAIM:M2AN, 54(2), pp. 391–430.

*Wei Zhang, Stefan Klus, Tim Conrad, Christof Schütte (2019). Learning chemical reaction networks from trajectory data. SIAM Journal on Applied Dynamical Systems (SIADS), 18(4):2000–2046. (Joint publication: Numerical Mathematics, Visual Data Analysis)

Ralf F. Ziesche, Tobias Arlt, Donal P. Finegan, Thomas M.M. Heenan, Alessandro Tengattini, Daniel Baum, Nikolay Kardjilov, Henning Markötter, Ingo Manke, Winfried Kockelmann, Dan J.L. Brett, Paul R. Shearing (2019). 4D imaging of lithium-batteries using correlative neutron and X-ray tomography with a virtual unrolling technique. Nature Communications. https://doi.org/10.1038/s41467-019-13943-3 (accepted for publication 2019-12-04)

Lin Zschiedrich, Felix Binkowski, Sven Burger (2019). Light-matter Interaction in Optical Resonators: Spectral Expansion by Riesz Projection. In Rita Asquini, editor, Photonics & Electromagnetics Research Symposium Abstracts (PIERS Rome), p. 1067.

Lin Zschiedrich, Felix Binkowski, Sven Burger (2019). Spectral resonance expansion by Riesz projections. In Manfred Kaltenbacher, editor, 14 th International Conference on Mathematical and Numerical Aspects of Wave Propagation Book of Abstracts, 234. https://doi.org/10.34726/waves2019

Kinga Żołnacz, Anna Musiał, Nicole Srocka, Jan Große, Maximilian Schlösinger, Philipp-Immanuel Schneider, Oleh Kravets, Monika Mikulicz, Jacek Olszewski, Krzysztof Poturaj, Grzegorz Wójcik, Paweł Mergo, Kamil Dybka, Mariusz Dyrkacz, Michał Dłubek, Sven Rodt, Sven Burger, Lin Zschiedrich, Grzegorz Sęk, Stephan Reitzenstein, Wacław Urbańczyk (2019). Method for direct coupling of a semiconductor quantum dot to an optical fiber for single-photon source applications. Opt. Express, 27:26772. https://doi.org/10.1364/OE.27.026772

Cyrille W. Combettes, Sebastian Pokutta (2019). Blended Matching Pursuit. In Proceedings of NeurIPS.

Jelena Diakonikolas, Alej Carderera, Sebastian Pokutta (2019). Breaking the Curse of Dimensionality (Locally) to Accelerate Conditional Gradients. In OPTML Workshop Paper.

## NOT PEER-REVIEWED (10)


Daniel Baum (2019). An Evaluation of Color Maps for Visual Data Exploration. In Bettina Bock von Wülfingen, editor, Science in Color: Visualizing Achromatic Knowledge, Monograph, pp. 147–161, De Gruyter.

Christiane Becker, Sven Burger, Klaus Jäger (2019). Nanooptische Simulationen für optoelektronische Anwendungen. Photonik: Fachzeitschrift für die optischen Technologien, 2019(2):42–45.

Thomas Breuer, Michael Bussieck, Frederik Fiand, Karl-Kiên Cao, Hans Christian Gils, Manuel Wetzel, Ambros Gleixner, Thorsten Koch, Daniel Rehfeldt, Dmitry Khabi (2019). BEAM-ME: Ein interdisziplinärer Beitrag zur Erreichung der Klimaziele. OR-News : das Magazin der GOR, pp. 6–8.

Mark A. Girolami, Ilse C. F. Ipsen, Chris J. Oates, Art B. Owen, T. J. Sullivan (2019). Editorial: Special edition on probabilistic numerics. Statistics and Computing, 29(6):1181–1183. https://doi.org/doi:10.1007/s11222-019-09892-y

Marc Hartung (2019). ZIB_Glucose: Luby Blocked Restarts and Dynamic Vivification. In Proceedings of SAT Race 2019: Solver and Benchmark Descriptions, pp. 43–44.

Markus Kantner, Theresa Höhne, Thomas Koprucki, Sven Burger, Hans-Jürgen Wünsche, Frank Schmidt, Alexander Mielke, Uwe Bandelow (2019). Multi-dimensional modeling and simulation of semiconductor nanophotonic devices. In WIAS Preprint, volume 2653, pp. 1–35, https://doi.org/10.20347/WIAS.PREPRINT.2653

Ilja Klebanov, Ingmar Schuster, T. J. Sullivan (2019). A rigorous theory of conditional mean embeddings. arXiv. https://arxiv.org/abs/1912.00671

Farouk Salem, Florian Schintke, Thorsten Schütt, Alexander Reinefeld (2019). Improving the throughput of a scalable FLESnet using the Data-Flow Scheduler. CBM Progress Report 2018, pp. 149–150. https://doi.org/10.15120/GSI-2019-01018

Jan Skrzypczak, Florian Schintke, Thorsten Schütt (2019). Linearizable State Machine Replication of State-Based CRDTs without Logs. arXiv. https://arxiv.org/abs/1905.08733v1

Oliver Ernst, Fabo Nobile, Claudia Schillings, Tim Sullivan (2019). Uncertainty Quantification. Oberwolfach Rep. 16, 695–772. https://doi.org/doi: 10.4171/OWR/2019/12


## BOOKS (2)


Stefanie Winkelmann, Christof Schütte (2019). Stochastic Dynamics in Computational Biology. To appear in Springer's Frontiers in Applied Dynamical Systems.

Wolfgang Dalitz (2019). alsoMATH - A Database for Mathematical Algorithms and Software. Proceedings of the Conference on Intelligent Computer Mathematics CICM.


## DISSERTATIONS (5)


Isabel Beckenbach (2019). Matchings and Flows in Hypergraphs. Freie Universität Berlin.

Luca Donati (2019). Reweighting methods for Molecular Dynamics. Freie Universität Berlin.

Gauri Mangalgiri (2019). Development of Titanium Dioxide Metasurfaces and Nanosoupbowls for Optically Enhancing Silicon Photocathodes. Humboldt-Universität zu Berlin.

Mohamed Omari (2019). A Mathematical Model of Bovine Metabolism and Reproduction: Application to Feeding Strategies, Drug Administration and Experimental Design. Freie Universität Berlin.

Marco Reidelbach (2019). Optimal Network Generation for the Simulation of Proton Transfer Processes. Freie Universität Berlin.

# IMPRINT

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

**Phone** +49-30-84185-0
**Fax** +49-30-84185-125
**E-mail** info@zib.de
www.zib.de

**Layout**
Christiane Weimann Grafik + Produktion
www.schnittmarke.biz

**Printing**
Lochmann grafische Produktion GmbH

**ISSN**
0934-5892



Imprint