



ZIB


2016 Annual Report

BIGGER DATA,
BETTER HEALTH
PAGE 30

SHAPE-BASED
DATA ANALYSIS –
METHODS AND
RESULTS
PAGE 44

RESEARCH
CAMPUS MODAL:
REPORT FROM
THE GASLAB
PAGE 52

SUPER-
COMPUTING
AT THE LIMIT
PAGE 84



ZIB

WWW.ZIB.DE

PREFACE

The year 2016 was the 75th anniversary of the invention of the computer by Konrad Zuse, after whom ZIB is named. On May 12, 1941, Konrad Zuse presented the first fully functional digital computer, his Z3, in Berlin and the digital age began. ZIB celebrated this historic breakthrough by organizing a series of public events in which the computer as an instrument, computing, and digital science in general were presented in many of their fascinating aspects. The main event, the international conference “The Digital Future – 75 Years Zuse Z3 and the Digital Revolution” assembled world-renowned experts and more than 1,600 participants in Berlin (see page 16 for details).

The anniversary was an ideal occasion for ZIB to reflect upon its strategic position and its research activities in a world where digital technology is exponentially growing, getting more complex, promising unprecedented insights into nature, economy, or society at large, but also into our personal lives. As Berlin's competence center for Digital Science, with its rich portfolio of collaborative projects in Applied Mathematics and Applied Computer Science, its regional, national, and international partners, and its expertise and capacities in data and information science, ZIB seems to be ideally placed, both scientifically and geographically, as a part of Berlin's thriving research and technology landscape. The members of ZIB are very proud of their institute and its outstanding performance, and strongly believe that success is only sustainable through continuous adaptation and development. In 2016, the initiation of the next round of the German Excellence Initiative presented a perfect platform for discussing ZIB's future research strategy together with its closest partner institutions. Debates began in almost all research fields regarding the respective grand challenges and related research opportunities. Members of ZIB participated in many of these discussions, and many new initiatives and partnerships origi-

nated from them. Through this process, ZIB has identified several new research priorities and, in turn, several of ZIB's research groups are part of Berlin-based Clusters of Excellence initiatives. One of the consequences of these developments, for example, will be that ZIB strengthens its research activities regarding digital humanities and computational social sciences.

Apart from strategic planning, ZIB continues to be a place booming with excellent research and first-rate scientific services and infrastructure. The year 2016 again broke several all-time records. For example, in total,

8 MILLION EUROS' WORTH OF THIRD-PARTY FUNDING

was acquired, which marked an increase for the fifth year in a row and a new record in ZIB's history! With more than one hundred ongoing research projects, ZIB grew to have over 200 employees, complemented by more than 70 research fellows and long-term guests. In several cases, the rewards of long-term efforts could be reaped. For example, a nontoxic

pain-relief drug, which has been developed purely based on computational methods in a cooperation between ZIB and Berlin's university hospital Charité within the last five years, has successfully passed preclinical trials, has led to a publication in *Science*, the formation of a new spin-off company, and the beginning of clinical tests (see page 22 for details).

This annual report provides insight into a variety of other success stories and gives a general overview of ZIB's organization and key factors for its successful development. In particular, six feature articles highlight aspects of our work: "Bigger Data, Better Health" gives an overview of ZIB-based data-driven research and related Web-based deployment of data-analysis and data-management tools for applications in biomedicine and health care. "Shape-Based Data Analysis – Methods and Results" features our research activities around shapes and geometry reconstruction with its diversity of applications ranging from medical therapy planning to digital humanities. The article "Computing the Truth" takes a completely different perspective by addressing the question of whether absolute evidence can be computed, at least in the realm of mathematical proofs. Turning back from philosophical questions to real-life applications, the article "Research Campus MODAL: Report from

the GasLab" reviews progress and challenges in the largest of ZIB's projects with industry that may even be the largest industry-related research project in all of German mathematics. Switching fields again, "Supercomputing at the Limit" explains why software modernization is one of the essential factors for the future of supercomputing and how this topic is addressed at ZIB. Finally, the article "Performance in Data Science" reports on ZIB's strategy to increase the performance of data-analysis frameworks by monitoring and improving data-access pattern, layout, and placement.

Obviously, these feature articles can only be appetizers. Readers interested in the specifics of ZIB's activities can find further articles and the full technical details of our research and services, of our projects, and the statistics about publications, lectures, and so on, on ZIB's Web pages: the starting point is www.zib.de.

Last but not least, ZIB is grateful for having many supporters and friends. Thanks to all those who provide us with encouragement, assistance, and advice.

Berlin, May 2017
Christof Schütte
President of ZIB

30

BIGGER DATA, BETTER HEALTH

Generating Medical Knowledge from
Vast Amounts of Data

4

ZUSE INSTITUTE BERLIN

Preface | Executive Summary |
Organization | ZIB Structure |
ZIB in Numbers | The Digital Future –
75 Years Zuse Z3 and the Digital
Revolution | The Hall of Fame |
Core Facility@ZIB | Computational
Design of a Nontoxic Painkiller |
Economic Situation | Spin-Offs |
Number of Employees

44

SHAPE-BASED DATA ANALYSIS – METHODS AND RESULTS

Towards Understanding Empirical
Shape Collections

CONTENTS

52

RESEARCH CAMPUS MODAL: REPORT FROM THE GASLAB

Meeting the Challenges of
Transporting Natural Gas

72

PERFORMANCE IN DATA SCIENCE

Digging Deep to Uncover Treasures

64

COMPUTING THE TRUTH

A Man-Machine Search for the
Absolute Evidence

84

SUPERCOMPUTING AT THE LIMIT

How Technology Innovations
Impact Software Design

104

ZIB PUBLICATIONS REFERENCES IMPRINT

EXECUTIVE SUMMARY

MATHEMATICS FOR LIFE AND FOR MATERIAL SCIENCES

The focus shift towards data-driven research gained momentum during the last year, with many activities in model-based data-analysis, inverse problems, data-driven and network-based modeling of multiscale processes, uncertainty quantification, and geometric image analysis for extracting information and deriving scientific knowledge from data. In this context, considerable effort went into the development of Web-based deployment of data-analysis and management tools; the collaboration with the FU Core Facility BioSupraMol on microscopy and the MODAL smart health platform are two prominent examples.

The successful extension of the Einstein Center for Mathematics enabled a focus

on the upcoming excellence initiative guiding many new developments, especially in new topics such as modeling and simulation of opto-molecular interactions and tissue-scale processes, or agent-based modeling in humanities and social science. Furthermore, during the last year, the division's research activity was shaped by an institute-wide trend towards the development and utilization of advanced machine learning techniques.

The continuation of established MSO activities in optical nanostructures, metabolic networks, and musculoskeletal diseases led to many achievements, with the extension of the DFG-CRC 787 as just one example.



MATHEMATICAL OPTIMIZATION AND SCIENTIFIC INFORMATION

The last year was marked by considerable progress in the optimization department's ability to deal with problems that require the integration of methods from different mathematical fields. Such problems are of paramount importance from both a theoretical and a practical point of view. Examples of such developments are successes in interactive combinatorial theorem proving using exact integer programming (see the respective feature article), first steps towards discrete-continuous aircraft trajectory routing algorithms, and striking results for the optimization of gas network operations using demand forecasts based on machine learning; the latter project was awarded the prestigious EURO Excellence in Practice Award 2016. ZIB also hosted the International Congress on Mathematical Software (ICMS 2016), a high-level conference that aims at an integration of mathematics and computer science.

Promoting open-access to scholarly research remained a major task for the scientific information department. The DeepGreen project was started with the aim to automatically transfer scientific publications into open-access repositories. The release of version 4.5 of the open-source repository software OPUS 4 marked a cornerstone of the department's ongoing commitment to openness. KOBV and digiS launched the digital preservation system EWIG which combines strategies and actions that address the still unresolved issue of ensuring long-term access to digital cultural heritage.

PARALLEL AND DISTRIBUTED COMPUTING

75 years after Konrad Zuse's invention of his Z3, the first digital, programmable computer, we prepared the procurement of the next HLRN supercomputer which will be more than 100,000,000,000,000 times faster than the Z3. Computing technology is still advancing at impressive speed, recently through technological advancements such as many-core CPUs, deeper memory hierarchies and wider SIMD vector units. Consequently, the difficult process of procuring a supercomputer has to be combined with further improvements of our HPC consultancy, which is a focal point in the efficient use of modern supercomputers. This includes code modernization in tight cooperation with application developers to achieve better performance on new system architectures.

Handling large data volumes plays an increasingly important role in high performance computing and large-scale data analytics. Within our two research projects GeoMultiSens and Berlin Big Data Center, for example, we develop methods to improve the performance of data-analysis pipelines. In some cases, better data placement, improved scheduling, and enhanced code optimization workflow resulted in tremendous runtime reductions.



ADMINISTRATIVE BODIES

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

President of ZIB is
PROF. DR. CHRISTOF SCHÜTTE

Vice President is
N.N.

The Board of Directors was composed in 2016 as follows

PROF. DR. PETER FRENSCH
Vice President, Humboldt-Universität zu Berlin (Chairman)

PROF. DR. CHRISTIAN THOMSEN
President, Technische Universität Berlin (Vice Chairman)

PROF. BRIGITTA SCHÜTT
Vice President, Freie Universität Berlin

DR. JUTTA KOCH-UNTERSEHER
until September 18, 2016
Senatsverwaltung für Wirtschaft, Technologie und Forschung
since September 19, 2016
Der Regierende Bürgermeister von Berlin Senatskanzlei Wissenschaft und Forschung

N.N.
since September 19, 2016
Senatsverwaltung für Wirtschaft, Energie und Betriebe

STS. STEFFEN KRACH
until September 18, 2016
Senatsverwaltung für Bildung, Jugend und Wissenschaft

PROF. MANFRED HENNECKE
Bundesanstalt für Materialforschung und -prüfung (BAM)

THOMAS FREDERKING
Helmholtz-Zentrum Berlin für Materialien und Energie (HZB)

PROF. DR. HEIKE WOLKE
Max-Delbrück-Centrum für Molekulare Medizin (MDC)

The Board of Directors met on May 27, 2016, and December 9, 2016.

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. ALFRED K. LOUIS
Universität des Saarlandes, Saarbrücken

PROF. DR. RAINER E. BURKARD
Technische Universität Graz, Austria

PROF. DR. MICHAEL DELLNITZ
Universität Paderborn

LUDGER D. SAX
Grid Optimization Europe GmbH

PROF. DR. ANNA SCHREIECK
BASF SE, Ludwigshafen

PROF. DR. REINHARD UPPENKAMP
Berlin Chemie AG, Berlin

PROF. DR. KERSTIN WAAS
Deutsche Bahn AG, Frankfurt am Main

The Scientific Advisory Board met on July 4 and 5, 2016, at ZIB.

THE STATUTES

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

ORGANIZATION

SCIENTIFIC ADVISORY BOARD

CHAIRMAN PROF. DR. JÖRG-RÜDIGER SACK

| *Ottawa*

PROF. DR. RAINER E. BURKARD | *Graz*

PROF. DR. MICHAEL DELLNITZ | *Paderborn*

PROF. DR. ALFRED K. LOUIS | *Saarbrücken*

LUDGER D. SAX | *Essen*

PROF. DR. ANNA SCHREIECK

| *Ludwigshafen*

PROF. DR. REINHARD UPPENKAMP | *Berlin*

PROF. DR. KERSTIN WAAS

| *Frankfurt am Main*

BOARD OF DIRECTORS

CHAIRMAN: PROF. DR. PETER FRENSCH

Humboldt-Universität zu Berlin (HUB)

PRESIDENT

PROF. DR. CHRISTOF SCHÜTTE

VICE PRESIDENT

N.N.

MATHEMATICS FOR LIFE AND
MATERIAL SCIENCES

Prof. Dr. Christof Schütte

MATHEMATICAL OPTIMIZATION AND
SCIENTIFIC INFORMATION

Prof. Dr. Ralf Borndörfer
Prof. Dr. Thorsten Koch

PARALLEL AND
DISTRIBUTED COMPUTING

Prof. Dr. Alexander Reinefeld

ADMINISTRATION
AND LIBRARY

Annerose Steinke

ADMINISTRATION AND LIBRARY
A. Steinke

ZIB STRUCTURE

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

Each of the scientific divisions is composed of two departments that are further subdivided into research groups (darker bluish color) and research service groups (lighter bluish color). Click on the respective box to get the unit of interest:

LEGEND

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

BRAIN BERLIN
RESEARCH AREA
INFORMATION
NETWORK
C. Schäuble

ZIB IN NUMBERS


PROMOTION OF YOUNG SCIENTISTS:



PEER-REVIEWED PUBLICATIONS IN INTERNATIONAL SCIENTIFIC JOURNALS

SEMINARS

GIVEN BY ZIB SCIENTISTS AT UNIVERSITIES



LECTURES

GIVEN BY ZIB SCIENTISTS AT UNIVERSITIES


3 PROFESSORSHIPS OFFERED TO ZIB RESEARCHERS



61 DISTINGUISHED


150 INVITED

12,903



OF MIP SOLVER SCIP

50 PB DATA ARCHIVE AT ZIB
TOTAL CAPACITY ON
16,600 TAPES



67 OUTREACH EVENTS FOR
SCHOOL CLASSES AND
THE GENERAL PUBLIC

€7,981,000

€2,493,770
INDUSTRIAL THIRD-PARTY PROJECTS



€5,487,230
PROJECT-RELATED, PUBLIC
THIRD-PARTY FUNDS



AWARDS


7,350

INTERNATIONAL GUESTS
AT ZIB IN 2016



105

CONFERENCES AND
WORKSHOPS AT ZIB



**IN JUNE 2016:
ZIB SUPERCOMPUTER IS
NO. 95 IN TOP-500 LIST**

THE DIGITAL FUTURE – 75 YEARS ZUSE Z3 AND THE DIGITAL REVOLUTION

Celebrating the 75th Anniversary of the Invention of the Computer; a Perfect Occasion for Science Communication and Public Outreach

On May 11, 1941, Konrad Zuse presented the first fully functional digital computer, his Z3, in Berlin and the digital age began.

75 years later, the world has changed completely: Konrad Zuse's invention started the digital revolution that brought computers into almost a billion households, created the Internet, mobile communication and smart devices, high performance computing, big-data analytics, as well as, with it, the information society and an avalanche of other innovations and transformations. Whether

this digital revolution is seen as a curse or blessing very much depends, however, on the perceptions and actions of humans: Do we possess the tools to utilize it for our goals, or will we be overwhelmed and even controlled by these advances? The challenge of mastering this transformation affects individuals as well as companies or organizations but also scientific fields and disciplines.

The Zuse Institute Berlin, feeling committed to the legacy of its name patron, aims at supporting debates between the public and the progressing digital

sciences, and therefore decided to utilize the organized 75th anniversary of the invention of the Z3 for organizing a series of related public events.

The main event, the scientific conference "The Digital Future – 75 years Zuse Z3 and the Digital Revolution" took place in Berlin on May 11, 2016, with more than 1,600 participants and a series of more than 40 presentations by internationally renowned experts in the following five sessions:



Horst Zuse and Governing Mayor of Berlin Michael Müller



“The Digital Future 2016”

1. Simulation, Optimization, Visualization
2. Data Analysis, Big Data, and Security/Privacy
3. The Future of Computing
4. Networks and Mobility
5. Communication, Digital Society, and Gaming

The scientific program was complemented by addresses of the Governing Mayor of Berlin, Michael Müller, the Head of the Government of the State of Berlin, and by the Chief of Staff of

the German Chancellery and Federal Minister for Special Affairs, Peter Altmaier.

The conference was closely followed by the media. For example, 25 German and international newspapers published articles about the conference, the anniversary, and ZIB’s work, and several TV channels broadcasted related reports, documentaries, and interviews.

In terms of outreach into the public, another media activity even outclassed

the effect of the conference: on each of the 75 days before the conference, ZIB’s media partner *Der Tagesspiegel*, one of Germany’s largest newspapers, published a portrait of a pioneer of the digital age and of their scientific achievements. These portraits were written by science journalists based on information provided by ZIB, were published on the second most read page of the newspaper, and created a significant awareness for the scientists presented as well as for ZIB.



Left to right: Christof Schütte, Horst Zuse, Leslie Greengard, Michele Parrinello, Governing Mayor of Berlin Michael Müller, Klaus-Robert Müller, Martin Grötschel



The digital pioneers presented in this article series, starting with Konrad Zuse and Alan Turing but also including technology drivers like Bill Gates or Steve Jobs, were all selected and inducted into the “Hall of Fame of the Digital Age” by an academic jury. In its electronic form, this Hall of Fame can be found at www.zib.de/de/hall-of-fame.

At the beginning of the conference, Michael Müller, Governing Mayor of Berlin, honored four digital pioneers for their remarkable achievements: Michele Parrinello, Leslie Greengard, Klaus-Robert Müller, and Konrad Zuse – who’s

son Horst Zuse accepted the award on his father’s behalf, and gave a presentation about the life of the inventor of the computer.

In addition to the conference, ZIB and Freie Universität Berlin together organized a lecture series of 12 public presentations by renowned experts in the course of the 2016 summer semester. This series covered essential topics of digital science: scientific visualization (C. Hege), machine learning (K.R. Müller), optimization (M. Grötschel), digital health (E. Böttinger), digitization of cultural heritage (T. Koch), bioinfor-

matics (K. Reinert), women pioneers in computing (S. Krämer), digital humanities (V. Lepper), supercomputing (A. Reinfeld), computational drug design (M. Weber), efficient data transfer and video compression (T. Wiegand), and autonomous driving (R. Rojas). The 90-minute-long talks were all appropriate for a broad public audience and attracted more than 2,000 participants in total. Most of the presentations were followed by long discussions between the speaker and the audience that made the lectures particularly successful as communication events between scientific experts and the wider public.





CORE FACILITY

In 2016, the IT department of ZIB was restructured to become the new Core Facility “IT and Data Services.” Carsten Schäuble was hired as the head of the Core Facility. The reorganization merged several working groups and competence sets and aims at making ZIB’s IT infrastructure fit for the increasing requirements of the digital revolution. The ultimate goal of the Core Facility is to build and offer an IT environment that enables Open Science for the members of ZIB and collaborating institutions.



One way to handle 50 PB of data

@ZIB



© Zuse Institute Berlin

The IT service portfolio will be improved and extended. This includes a new identity management, which will allow better self-service and automation of common tasks such as work-group communication and management services. Internally, new security standards will be implemented, mainly effecting backup, storage, and firewall technology. Especially the access layer for the internal ZIB networks will be newly designed and installed. A new software-defined firewall system will allow the creation of two security zones and virtual private networks (VPNs) in which mission-critical and experimental services can be placed and run with different access patterns from the World Wide Web.

On the infrastructure side, ZIB's data-center will be fully modernized. New power-efficient server systems and racks will be installed, including new technology for climate, power supply, and distribution. Based on this, all internal resources such as compute and virtualization servers, HPC, and storage systems, as well as the network

infrastructure will be modernized and included in an institute-wide management system. This also includes the BRAIN (Berlin Research Area Information Network) system. Here, installation of the new 100 Gbit router was finished at the end of 2016. Until May 2017, BRAIN will take the Berlin research network to the next performance level, raising the possible network throughput from 10 Gbit/s up to 100 Gbit/s.

The modernization of services and infrastructure will lead to an innovative service portfolio that the new core facility can deliver and offer to internal members and external groups and institutions. Ultimately, this will enable more and advanced projects running within ZIB's data center. As one of the first steps, the Europe-wide authentication infrastructure DFN AAI AAI – Identity Federation was implemented which allows participating institutions to use ZIB services and enables ZIB members to use external DFN services with their normal ZIB account.


COMPUTATIONAL DESIGN OF A NONTOXIC PAINKILLER

ZIB's research group "Computational Molecular Design" successfully designed a candidate for a pain-relief drug. The design strategy was created in cooperation with clinical researchers at Berlin's university hospital Charité but was performed purely computationally. The drug candidate very successfully passed all preclinical tests; clinical trials will follow.



© Zuse Institute Berlin

One of the inventors, Marcus Weber, in front of ZIB's supercomputer that was used for parts of the development



NFEPP inside the binding pocket of the μ -opioid receptor

Worldwide, more than 1,000,000,000 people suffer from significant pain. In 50% of these cases, the pain is not adequately treated. Although effective analgesic drugs are available, their use is limited due to severely adverse side effects. Opioids, the most effective painkillers, produce sedation, apnea, addiction, and constipation mediated in the brain or gut. In the USA, the rates of overdose on prescription opioids increased 4-fold in the last decade, with almost 19,000 overdose deaths associated with prescription opioids in 2014. The famous singer Prince died from such an overdose in 2016. Similar alarming trends can also be observed in Europe.

In a long-lasting cooperation between ZIB and Charité, the idea was created to change the molecular structure of opioids in such a way that the painkillers are only active in inflamed tissue and inactive in other parts of the body, especially the brain or gut. This kind of inactivity would result in an absence of the adverse side effects, that is, in a nontoxic painkiller.

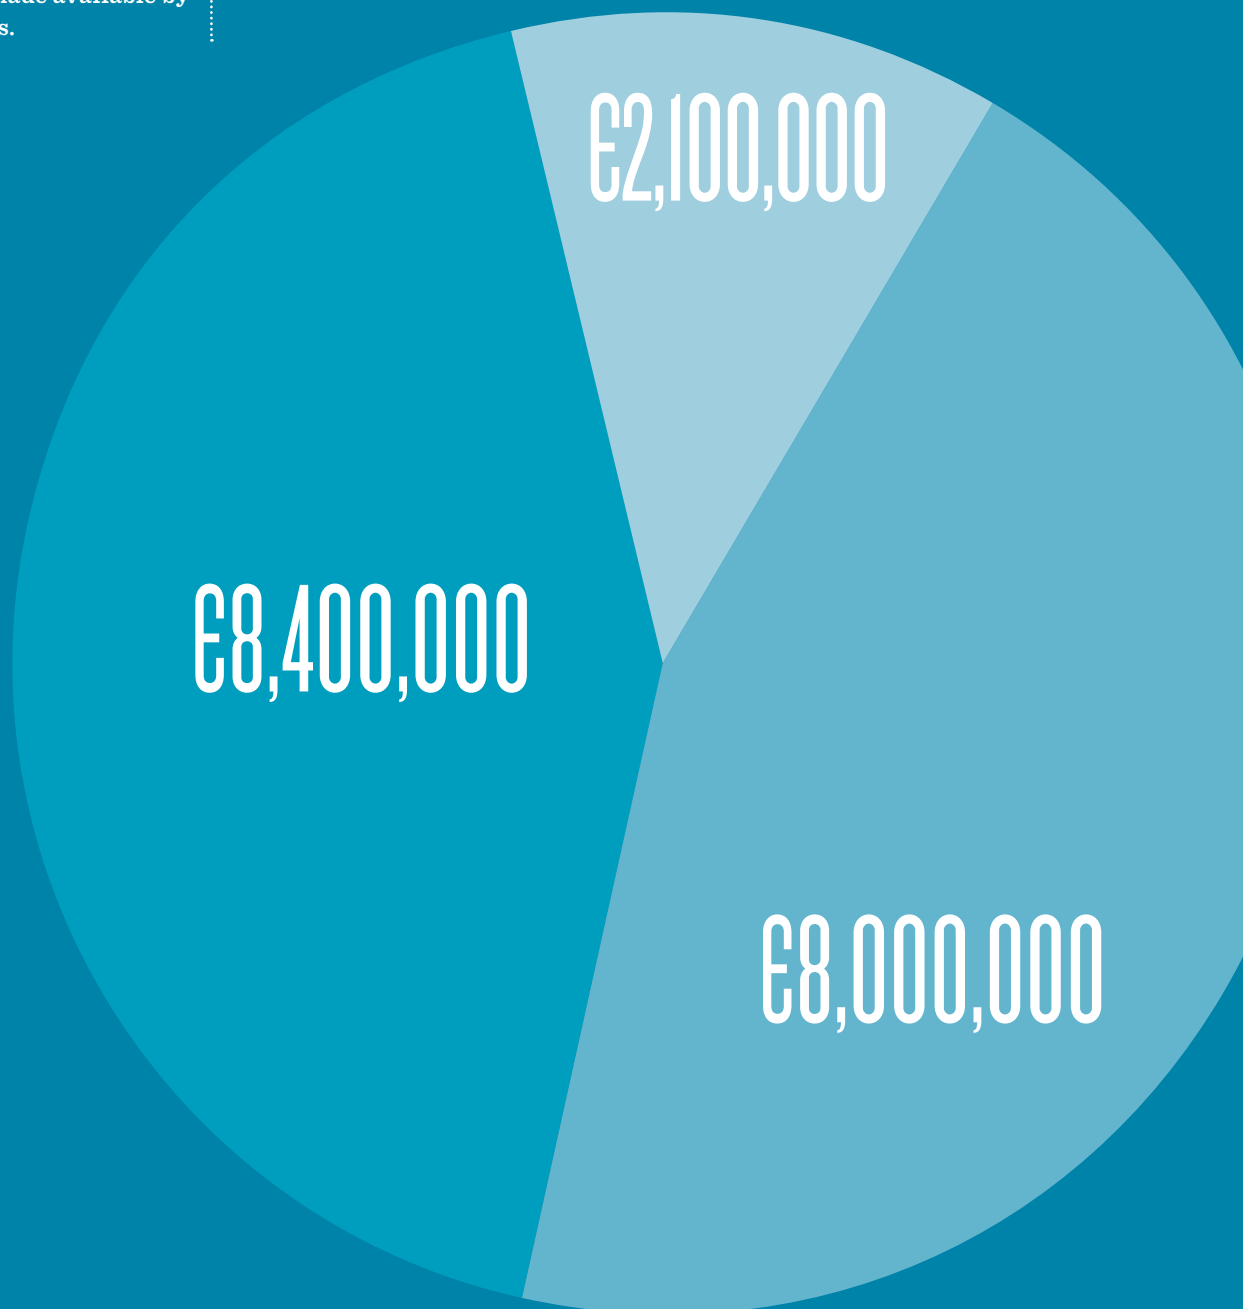
The search for an appropriate structural change was performed with purely computational means: under the working hypothesis that inflamed and healthy tissue exhibit different pH values, we screened for molecules that are good agonists of the μ -opioid receptor for normal pH but fail to do so for reduced pH values.

The computations resulted in the design of a fentanyl derivate (NFEPP). On the one hand, quantum chemical simulations showed that NFEPP would allow for the desired discrimination between healthy and inflamed tissue by pH values, while docking simulations using classical molecular dynamics in combination with conformation dynamics approaches developed at ZIB showed that NEEPP would be a good agonist of the μ -opioid receptor [CMD2]. After synthesizing NFEPP, it was used for animal tests at Charité that confirmed the properties predicted computationally [CMD1]. In 2016, a ZIB spin-off (DoloPharm UG) was founded with the aim to organize clinical trials of this drug candidate.

[CMD1] V. Spahn, G. Del Vecchio, D. Labuz, A. Rodriguez-Gaztelumendi, N. Massaly, J. Temp, V. Durmaz, P. Sabri, M. Reidelbach, H. Machelska, M. Weber, C. Stein: A nontoxic pain killer designed by modeling of pathological receptor conformations. *Science*, publication submitted August 2016, to appear in March 2017.

[CMD2] C. Stein, M. Weber, C. Zöllner, O. Scharkoj: Fentanyl derivatives as pH-dependent opioid receptor agonists. USA patent US 14/239,461 (2015).

In 2016, the total income of ZIB comprised 18.5 million euros. The main part of this was made available by the State of Berlin as the core budget of ZIB (8.4 million euros) including investments and Berlin's part of the budget of HLRN at ZIB. A similarly large part resulted from third-party funds (8.0 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 budgets of BRAIN (State of Berlin) and KOBV (mixed funding) as well as the part of the HLRN budget made available by other German states.



ECONOMIC SITUATION IN 2016

ZIB INCOME

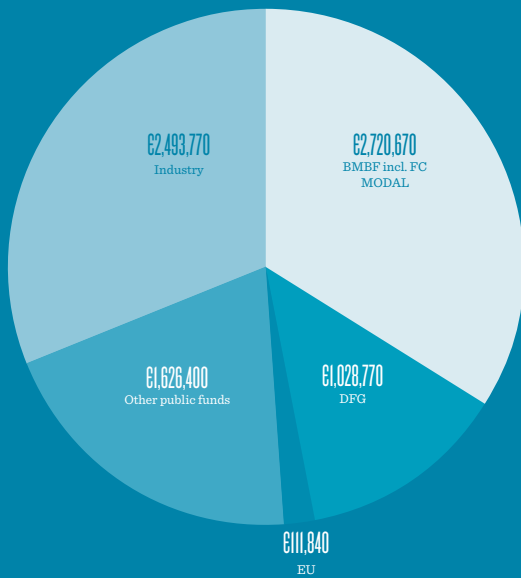
- 45% Core budget by State of Berlin
- 43% Third-party funds
- 12% Further grants

ECONOMIC SITUATION IN 2016

The Zuse Institute Berlin finances its scientific work via three main sources: the core budget of ZIB provided by the State of Berlin, third party funds from public sponsors, and those of industrial-cooperation contracts.

In 2016, ZIB raised third-party funding for a large number of projects. Project-related public third-party funds decreases slightly by 4 % to 5,487 k€ in 2016, and industrial third-party projects increased by more than 18% from 2,106 k€ in 2015 to 2,494 k€ in 2016.

In total, 7,981 k€ third-party funding marked a new record in ZIB's history; an increase for the fifth year in a row!




ZIB THIRD-PARTY FUNDS BY SOURCE

- 2% EU
- 13% DFG
- 20% Other public funds
- 31% Industry
- 34% BMBF incl. FC MODAL



ZIB THIRD-PARTY FUNDS IN EUROS

- INDUSTRY
- PUBLIC FUNDS



COMPUTING IN
TECHNOLOGY GMBH
(CIT)

1992 | www.cit-wulkow.de

Mathematical modeling and development of numerical software for technical chemistry

RISK-CONSULTING
PROF. DR. WEYER GMBH

1994 | www.risk-consulting.de

Database marketing for insurance companies

INTRANETZ GMBH

1996 | www.intranetz.de

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

AKTUARDATA GMBH

1998 | www.aktuardata.de

Development and distribution of risk-evaluation systems in health insurance

VISAGE IMAGING GMBH

(Originating from a spin-off of Visual Concepts GmbH)

1999 | www.visageimaging.com

Visualization and data-analysis (Amira etc.), especially medical visualization

ATESIO GMBH

2000 | www.atesio.de

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

BIT-SIDE GMBH

2000 | www.bit-side.com

Telecommunication applications and visualization

DRES. LÖBEL, BORNDÖRFER &
WEIDER GBR

2000 | www.lbw-berlin.de

Optimization and consulting in public transport

LENNE3-D GMBH

2005 | www.lenne3d.com

3-D landscape visualization, software development, and services

JCMWAVE GMBH

2006 | www.jcmwave.com

Simulation software for optical components

ONSCALE SOLUTIONS GMBH

2006 | www.onscale.de

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

LAUBWERK GMBH

2009 | www.laubwerk.com

Construction of digital plant models

1000SHAPES GMBH

2010 | www.1000shapes.com

Statistical shape analysis

TASK – Berthold Gleixner Heinz Koch
GbR

2010

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

QUOBYTE INC.

2013 | www.quobyte.com

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

KEYLIGHT GMBH

2015 | www.keylight.de

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

DOLOPHARM UG

2016

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

SPIN- OFFS

NUMBER OF EMPLOYEES

In the year 2016, 236 people were employed at ZIB; of these, 172 positions were financed by third-party funds.

1/1/2016			1/1/2017			
3	0	3	3	0	3	MANAGEMENT
18	87	105	15	99	114	SCIENTISTS
34	13	47	38	10	48	SERVICE PERSONNEL
7	8	15	8	8	16	KOBV HEADQUARTERS
0	54	54	0	55	55	STUDENTS
62	162	224	64	172	236	TOTAL
Permanent	Temporary	TOTAL	Permanent	Temporary	TOTAL	




Prof. Dr. Tim Conrad | conrad@zib.de | +49-30-84185-250

BIGGER DATA, BETTER HEALTH

© 2014 ZIB
All rights reserved.
Zentrum für Informationstechnik Berlin
www.zib.de



REPORT TO OFFICE AND STATE FOR MORE INFO
REPORT TO OFFICE AND STATE FOR MORE INFO
REPORT TO OFFICE AND STATE FOR MORE INFO
REPORT TO OFFICE AND STATE FOR MORE INFO
REPORT TO OFFICE AND STATE FOR MORE INFO



1

DO YOU REMEMBER THE LAST TIME A MEDICAL DOCTOR TOOK SOME DATA FROM YOU?

Chances are good it did not hurt at all and you might not even have noticed. However, if you take a moment and think about it, you probably remember an x-ray being taken or some blood test performed. Maybe you even remember that your body mass index was calculated and written down, or simply your weight being measured. It's possible you also underwent some gene testing procedure to check which drug you would react to best for a particular treatment.

One of the most noteworthy changes during the last years is that hospitals started to routinely collect and store all these data. What is yet missing though is the connection of all these data within and across hospitals, the aggregation, and

– finally – its deep analysis. Making use of these data will support medical doctors in the future to make more informed decisions. For example, comparing their own patients to millions of other cases of patients having the same conditions and (early) symptoms would allow to recognize a disease in a very early stage and select the best personalized treatment option.

At ZIB, we support the rise of this new type of information-based medicine in two ways. First, we develop IT infrastructure needed for secure storage and fast retrieval of these very large data collections. Second, we work on new algorithms and methods to analyze these data and turn it into medical knowledge.

1 A schematic representation of a DNA molecule and virus cells. The challenge is to find changes induced by these tiny intruders that can help identifying a particular disease.

DETECTING DISEASES FROM LARGE-SCALE OMICS DATA




Early detection of a disease usually leads to a significantly better outcome for the patient, compared to the same treatment given at a later time. This is true in particular for some types of cancer or cardiovascular diseases, where hours or days lost can make a substantial difference. Many traditional detection approaches based on molecular tests often fail in early detection due to a lack in efficiency. In many cases these methods simply fail because changes in the traditional, measurable biomarkers are simply too subtle.

BIOLOGY HAPPENS

Changes in the human body during progression of a disease happen on many biological levels, such as the genome (what can happen?), the transcriptome (what seems to happen next?), the proteome (what is currently happening?), or the metabolome (what happened?). Following the central dogma of molecular biology and its extensions, these levels are highly interconnected and depend on each other. Thus, tests for diagnosing a condition or classifying a disease's state should take all available data into account to build more reliable diagnostic predictors.

This can be especially crucial given the nature of biological data which are usually very noisy and contains errors (e.g. missing data). In this situation, multiple data types may compensate for missing or unreliable information in any one data type and thus yielding a more robust method, with regards to biological uncertainties. Further, different sources of data that point to the same gene, protein, or metabolite are less likely to be false positives and could indicate functionality.

THE CENTRAL DOGMA OF BIOLOGY





BIG DATA IN DISEASE DIAGNOSIS

Researchers at ZIB work on new methods based on the combination of multiple biomarkers, potentially from different biological omics levels and including new data types such as environmental factors or tracking devices. However, the more variables can be measured and hence enter the modeling equations, the more data is needed to determine their correct values. Or – in other words – to prevent overfitting. Until recently, acquisition of these large data sets was very time-consuming and certainly very expensive. Only large international scientific collaboration efforts could afford these large studies. Thanks to today’s high-throughput omics technologies, formerly hardly imaginable omics data volumes can be generated within days with a fraction of the budget needed some years ago. Using these technological advances, vast amounts of data have been created, providing a very detailed view on diseases and their diagnosis.

Today, we have all this information available in large databases and the computing power for detailed analysis. What is still missing are efficient methods and algorithms to extract knowledge. Knowledge, that is interpretable for practitioners and ultimately supports medical doctors to make better diagnoses.

REDUCING COMPLEXITY

The main problem when using large-scale data is that possible analysis results are also large scale. The goal is thus to generate answers that are of low complexity. Researchers at ZIB have found a new way to only extract the relevant (low-complexity) information from very large biological data sets. Using these new methods will allow us to generate solutions that can be easily understood by humans, and not only by large supercomputers.

For more details about efficient large-scale data-management methods developed by ZIB, see the feature article “Performance in Data Science.”

$$\max_{\omega \in \mathbb{R}^d} \sum_{i=1}^n y_i(x_{i_1} \omega) \text{ subject to } \|\omega\|_1 \leq \sqrt{\lambda} \text{ and } \|\omega\|_2 \leq 1$$

DATA ANALYSIS FOR UNDERSTANDING NEURONAL DEVELOPMENT AND MAINTENANCE

Stochastic mechanisms play an important role in the brain. A better understanding of the principles creates insight in disease mechanisms, which will contribute to developing effective therapies of neurological disorders. ZIB has been working together with NeuroCure on understanding principles of neuronal development and maintenance. Two specific examples are the analysis of protein trafficking during synaptic maintenance and filopodia dynamics in neuronal growth cones.

Understanding the mechanisms that cause neurological disorders such as multiple sclerosis and epilepsy could help developing therapies. Insight in the general principles are key to understanding specific disorders. We have been developing methods for image analysis and mathematical modeling that enable understanding stochastic mechanisms of maintenance in the adult brain and development in the growing brain.

Synapses are the neuronal connections in the brain. They remain functional over long periods, possibly the entire life span of an organism. The understanding of maintenance processes that enable the longevity of synapses is key to understanding the longevity of the brain as a whole. Seemingly random trafficking of proteins and cell organelles along axons can be observed using light microscopy.

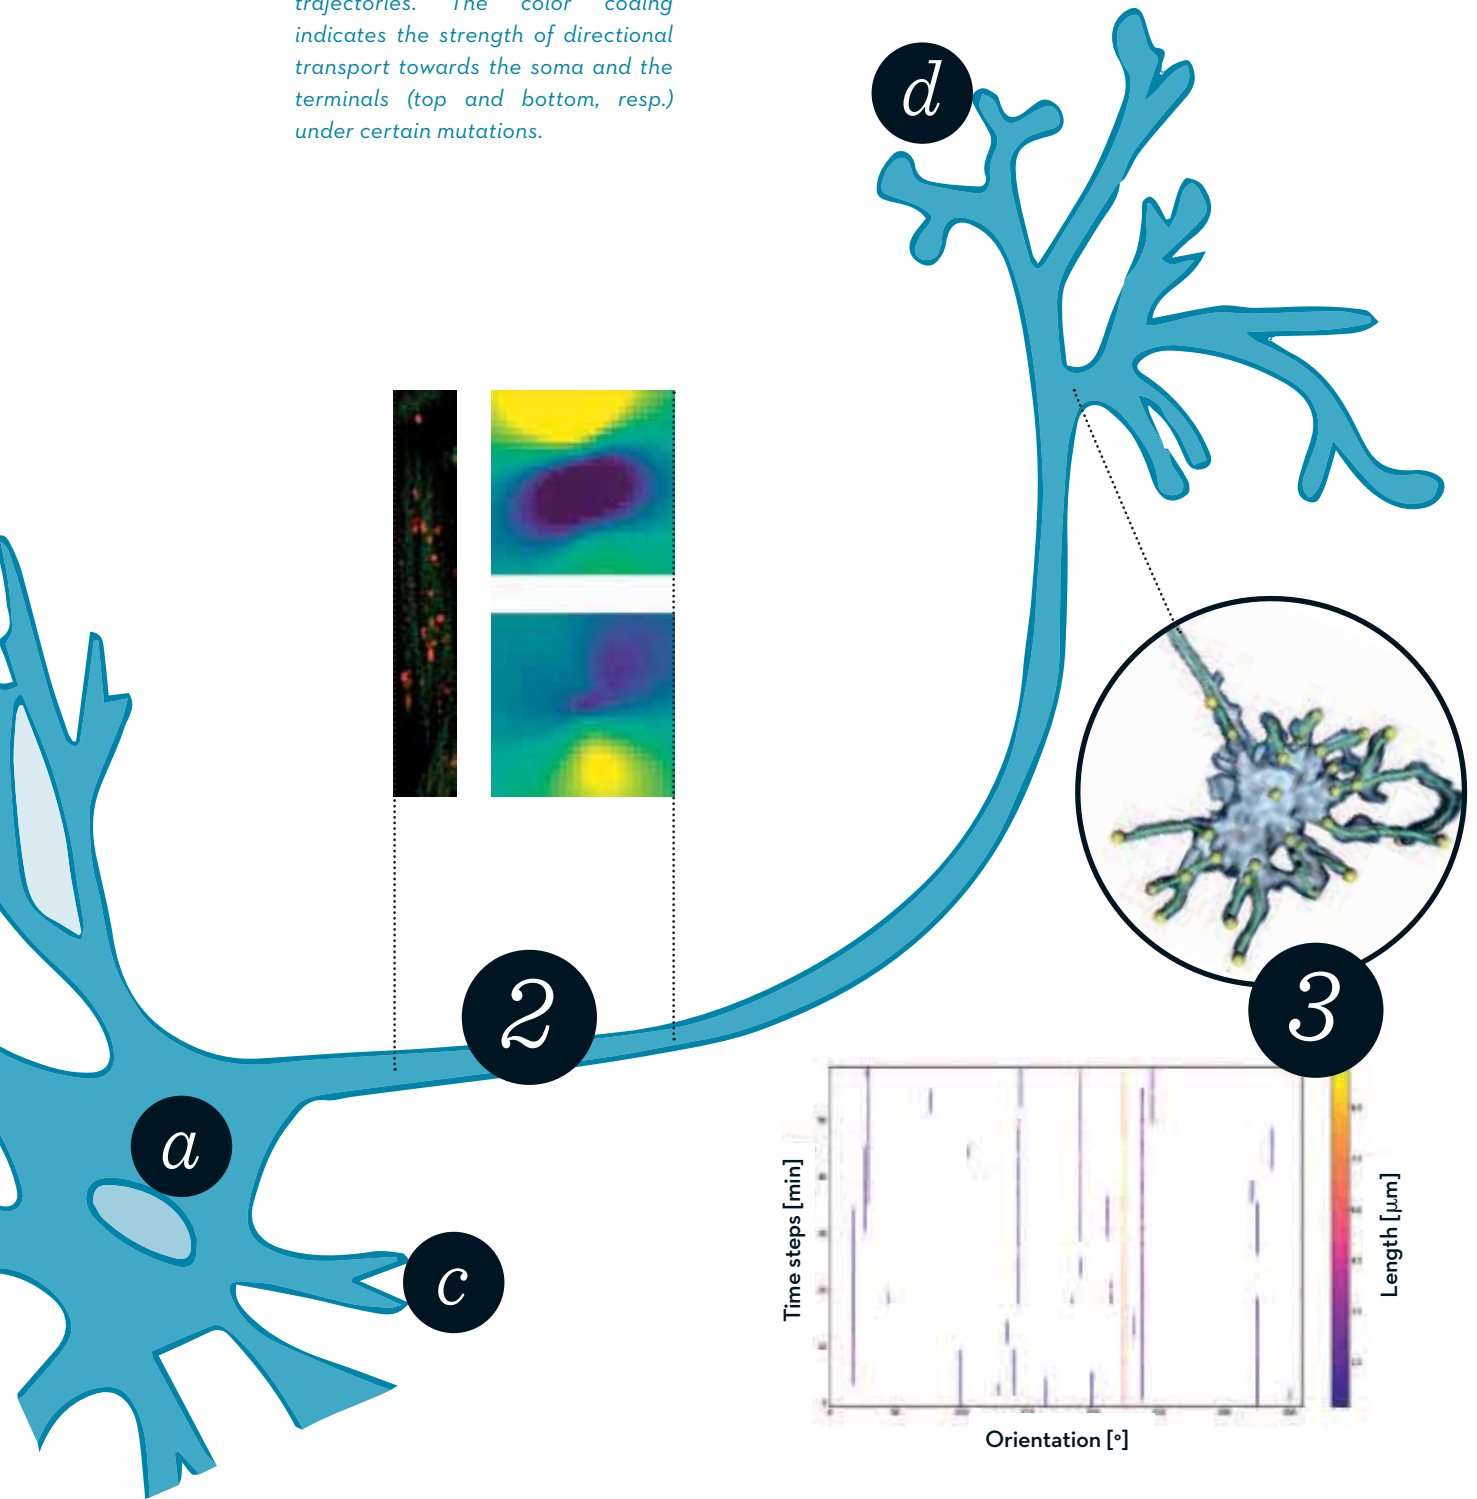
By applying image analysis, we are able to reconstruct and analyze particle dynamics, which can then reveal differences between cell mutations. (see figure 2). By applying stochastic modeling, we hope to explain the observed transport processes and understand how they lead to robust, long-term maintenance of synapses, whose underlying mechanism are yet unknown.

The brain develops complicated neuronal circuits. Long-term, high-resolution live imaging can be applied to visualize neuron growth during brain development. Growth cones exhibit seemingly random movements of finger-like structures, called filopodia, which eventually form connections with neighboring neurons. It is a remarkably robust process that reproduces complex neuronal circuits in individual brains. We have been developing image analysis methods that allow reconstructing the geometry of growth cones from microscopy images. Automatic image analysis is applied to create a geometric representation of the filopodia. The geometry is then manually corrected and fed back to the automatic analysis to incrementally refine the representation. The results are visually and statistically analyzed into gain insight in the neuron growth dynamics over time (see figure 2). The information is also key to validating stochastic mathematical models of the growth processes.



*The soma **a** collects the input signal from the dendrites **b** and generates a specific output signal. The axon **c** is the only output of a neuron. Its characteristic electrochemical signal is forwarded toward the axon terminals **d** which are connected via synapses with the dendrites of a succeeding neuron.*

2 Axonal trafficking. The microscopy image (left) shows transport vesicles (red), which are transported along the axon. The diagram (right) depicts a statistical analysis of particle trajectories. The color coding indicates the strength of directional transport towards the soma and the terminals (top and bottom, resp.) under certain mutations.



3 Growth-cone dynamics. The image (top) illustrates a growth cone with a geometry representation of the filopodia (green lines and yellow dots). The geometry allows a simplified estimation of length and orientation of the filopodia. The change of length over time is color coded in the diagram (bottom). The plot indicates the life span of a filopodium (y-axis) and the orientation (x-axis).



© Ed Yourdon

ESTIMATING INDIVIDUAL CARTILAGE STATE FROM GAIT DATA AND IMAGING

DIAGNOSTIC SCORES FROM MRI

Osteoarthritis, a degenerative disease of articular cartilage, is prevalent in our aging society, mostly affecting the knee joint. While a strong influence of mechanical loading – in particular of high impacts such as those occurring in extreme sports – on osteoarthritis onset and progression is evident, the mechanisms of cartilage degradation are not well understood, despite decades of clinical research. Within the musculoskeletal research network OVERLOAD/PrevOP, we are working together with Charité Berlin on improving understanding.

In clinical practice, cartilage damage is graded by expert readers employing semiquantitative scoring methods. The cartilage damage is usually assessed in MRI data considering percentages of cartilage thickness loss and cartilage area affected by degradation. We are developing automated methods for localization and grading of cartilage damage. First, using multi-object graph-cutting methods, a region of interest of the femoral and tibial cartilage is automatically determined in the MRI data yielding a 3-D segmentation. Afterwards, using both MRI data and the expert readers'


scores for several hundred data sets, deep neural networks will be trained to automatically classify and localize the cartilage damage. Jointly with Charité Berlin, a prospective study started in February 2016 with up to 240 patients for two years. These patients perform different levels of regular exercise under psychological supervision. Based on image data of this study in combination with deep learning-based methods for automated assessment of cartilage damage, we are aiming to evaluate whether or not regular exercise slows down the rate of osteoarthritis progression.

CARTILAGE LOADING FROM GAIT DATA

While the importance of mechanical loading is generally accepted, correlating disease progression with local stress and strain states of loaded cartilage is hard due to the difficulty of measuring these dynamic quantities in vivo. An approach we pursue is to estimate the local mechanical state of knee cartilage from gait data by coupling a global dynamic multibody model of the lower limb to an elastomechanical contact model of the cartilage. Adjusting the active muscle forces during motion such that the laws of mechanics are satisfied on the patient-specific anatomy and the measurements such as skin-marker positions (see figure 4) and ground reaction forces

are reproduced allows the computation of 3-D distributed stress and strain patterns within the cartilage (see figure 5). These can be correlated to clinical findings, variations of the anatomy, or frequent lesions like loss of the anterior crucial ligament.


Solving the identification problem involves a combination of finite element methods, time-stepping schemes, non-linear solvers, and constrained optimization. In particular, the intricate geometry of cartilage surfaces introduces high nonlinearity into the problem of force equilibration. The focus is therefore on robustness of the solver, as is required for investigating larger patient cohorts.



© Julius Wolff Institute

4 Dynamic multi-body models of the lower extremities relate measurable quantities, such as marker positions, to bone displacements and forces acting on the cartilage.

5 Stress distribution in articular cartilage computed from forces provided by a multi-body gait model. Spatially resolved stresses allow the relation of anatomy and lesions to clinical findings.



© Zuse Institute Berlin

STATISTICAL GAIT AND EXERCISE ANALYSIS

The statistical interpretation of exercise data for different patients is complicated and time-consuming due to its high diversity. To address these problems, we are developing and using interactive and automated mathematical methods on time-dependent gait and exercise data (in vivo and in silico) to extract comparable sets of parameters and features.


INTERPRETING MEDICAL IMAGES

DEEP LEARNING FOR ENDOSCOPIC SURGERIES


With advancement in surgical techniques, modern procedures are driven more towards minimally invasive surgeries (MIS). Nowadays, a surgeon may treat a lesion inside the human body by making small incisions or navigating through natural orifices with the assistance of a video camera and several thin instruments. Despite challenges due to limited view of the surgical site, rotating images and impaired hand-eye coordination, MIS procedures such as laparoscopic cholecystectomy (for gallbladder removal) and endoscopic endonasal (for skull-base surgery) are becoming the de facto surgery standards. A promising way to support MIS procedures is identifying the current position of the endoscope and detecting the surgical tools over the endoscopic video stream. The key benefit is robotic assistance in image-guided surgery and assistance during navigation.

LANDMARK DETECTION

To facilitate a MIS procedure, the therapy planning group of ZIB is working on the development of a multimodal framework for determining the anatomical position of the endoscope by using real-time image, sensor, and process information. It would assist the surgeon in guiding the endoscope through rich visualization and alert them about the risk structures in close proximity. A possible visualization scheme can be seen in figure 6.

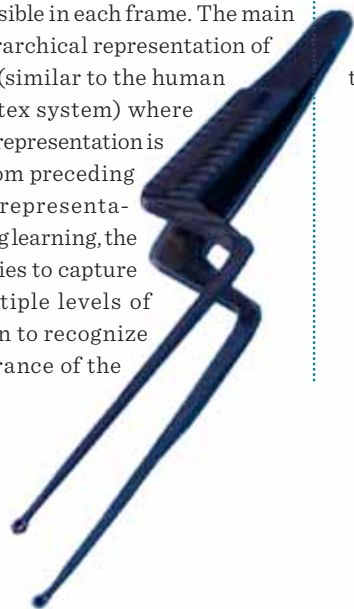


6 Image-based navigation support for endonasal use case scenario. At the top, possible paths inside the nose are represented by a simplified graph, where a green circle represents the current position of the endoscope and green text represents the name of the corresponding landmark. At the bottom, a reference model shows the overall frontal structure of the nose. It is cut to the corresponding position for better visibility.



SURGICAL TOOL DETECTION

A novel Convolution Neural Network architecture – named ZIBNet – has been developed, which is specifically designed for surgical tool detection tasks. It takes endoscopic video frames as an input and learns to detect the presence of the surgical tools visible in each frame. The main idea is hierarchical representation of the image (similar to the human visual cortex system) where high-level representation is derived from preceding low-level representation. During learning, the network tries to capture these multiple levels of abstraction to recognize the appearance of the tools.




MULTI-TOOL PROBLEM

Since multiple surgical tools can be visible in an image, ZIBNet focuses on learning the tool-pair patterns (tools that are commonly used in conjunction with other tools) to map each image to its corresponding tool pair (figure 7). Moreover, ZIBNet employs an enhancement approach (called temporal smoothing) to suppress run-time false detections.

CURRENT STATUS

The utility of ZIBNet has been evaluated for cholecystectomy (gallbladder removal) procedures. In particular, ZIBNet performed second best during the M2CAI16 tool-presence-detection challenge. Currently, ZIBNet is the state-of-the-art method for tool-presence detection in cholecystectomy procedures. A joint tool and landmark-detection method is currently in progress.



7 Chord diagram of second-order tool co-occurrences in gallbladder-removal surgery. Common tool pairs are connected by a ribbon. Rare pairs are marked by red and frequent ones by green.

A PROTOTYPICAL PLATFORM FOR SMART HEALTH MANAGEMENT

Managing your health in the current health-care system is not easy. With changing health plans, multiple and disconnected practitioners, and the wish to focus on prevention, we need to take responsibility for our own health.

Smart Health Management means to change the perspective about how we look at our health and our body. The underlying philosophy is that we can collect all possible information (data) about ourselves and use them to adjust our lifestyle in a preventive way. It means to constantly collect, integrate, and analyze data that our body are producing. Smart Health Management is based on the idea that we can (and want to) see early signals that something might go wrong. And then take appropriate action. Thus – unlike traditional approaches – it is not only focused on diagnosing what what

has already gone wrong (such as catching a bad cold). It also means to recognize when our body needs more sleep, more fitness sessions, or a particular vitamin.

All these findings are based on the data we feed into the analysis system. What we make out of this eventually depends on us. If the data show – based on a metabolomics analysis – that we are lacking some vitamins or other essential nutrients, we can choose to change our food based on the recommendation. If the data show that – based on our genes – we should give favor to a particular drug to treat that nasty headache, it is on us to get that drug or go with the old one.

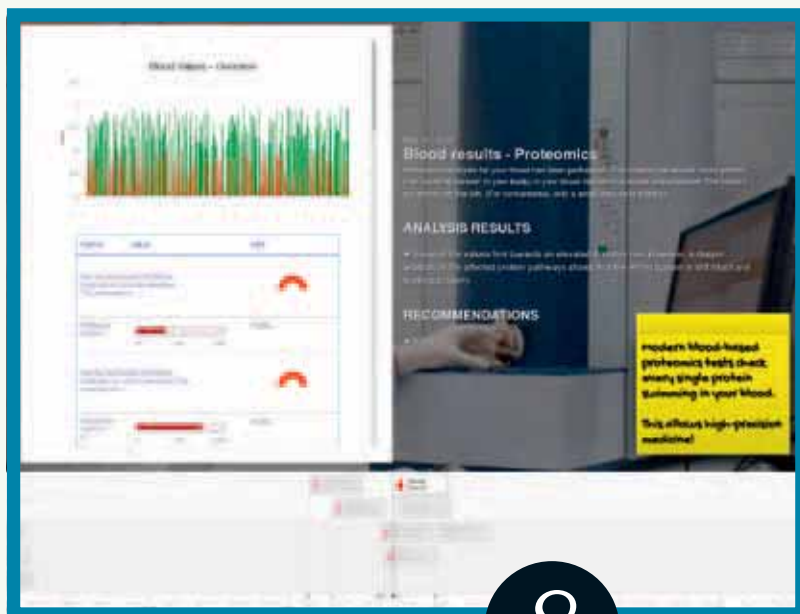
Researchers at ZIB are working on a system to enable Smart Health Management for everyone. This platform allows secure storage of all relevant data and provides smart-analysis services. At the very core, all data sources will be integrated and

put into perspective: information about our genetic setup, type, and concentrations of proteins and metabolites that are floating through our blood, names of drugs we are taking, and tracking information such as step counters or blood pressure throughout the day. Using this variety of complementary information will make it possible to understand and manage our body better. New analysis algorithms that can deal with the size and complexity of this (big) data will be able to find hidden signals and patterns. They will find whether a correlation exists between your low blood pressure, sleep habits, and the weather. Based on this extracted knowledge, smart services will be able to suggest what, where, and how things can be improved and what is the expected effect on us. Most importantly: the system will also alert us when it is better to see an actual human doctor.

More than 70% of the data needed for health management is in your blood. It contains information about your health status, optimal disease treatment, and prevention.

Your blood cells regenerate every 120 days, so you can quickly measure significant improvements from lifestyle changes such as nutrition intake.





8



8 A glimpse into the ZIB Smart Health Management platform. Top: Visualization of the blood proteome and possible diagnoses. Bottom: Integration of step-tracking data over one day.



SHAPE-BASED DATA ANALYSIS – METHODS AND RESULTS

Geometric Shapes in Computers

3-D shapes are all around us and our visual system is very well trained in recognizing, comparing, and categorizing them. But how can we teach this to a computer? What are good shape representations and how can one do statistics on shapes? And can we exploit prior shape knowledge to solve otherwise insoluble problems?



MOTIVATION

We live in a world of 3-D objects, each characterized by its material properties and its geometric form. In this article, we focus on the latter, the geometric shapes, from a computer science perspective. Traditionally, the main topics in computerized processing of geometric shapes were the construction and visualization of shapes. The techniques that were developed later found their way into, for example, CAD and computer animation software. Nowadays, with the availability of techniques for capturing 3-D shapes, the focus has shifted to shape analysis and exploitation of the acquired form knowledge. For example, in archaeology, one aims at quantifying similarities of shapes in order to classify and stratify archaeological finds, or at using shape knowledge to restore damaged and incomplete objects; in biology, one is interested in grouping anatomical shapes to understand relations between anatomical and, for instance, genetic variations; in medicine, one would like to characterize “normal” anatomy including “normal” variation to identify and classify pathologies or malformations, or to utilize the shape knowledge in surgical reconstructions; in autonomous driving, scene recognition from incomplete information becomes possible by utilizing prior information about shapes and appearances of potentially occurring objects. These examples indicate that algorithmic techniques for empirical analysis of shapes and for utilizing shape knowledge have great potential in a broad range of applications, including humanities, natural sciences, medicine, and engineering.

MATHEMATICAL FRAMEWORK

Shapes and shape manifolds.


Considering a physical object in 3-D space, its shape is defined by infinitely many points that constitute its boundary. A boundary surface can be considered as a single point in an infinite-dimensional “configuration space,” comprising the coordinates of its points. The shape is independent of rotation, translation, and sometimes also scaling. Mathematically, this means that a shape is an equivalence class of boundaries with these transformations filtered out; furthermore, the space of all possible shapes, called “shape space,” is a quotient space (Kendall 1989). While the configuration space is a linear space, the shape space is curved and therefore also called “shape manifold.”

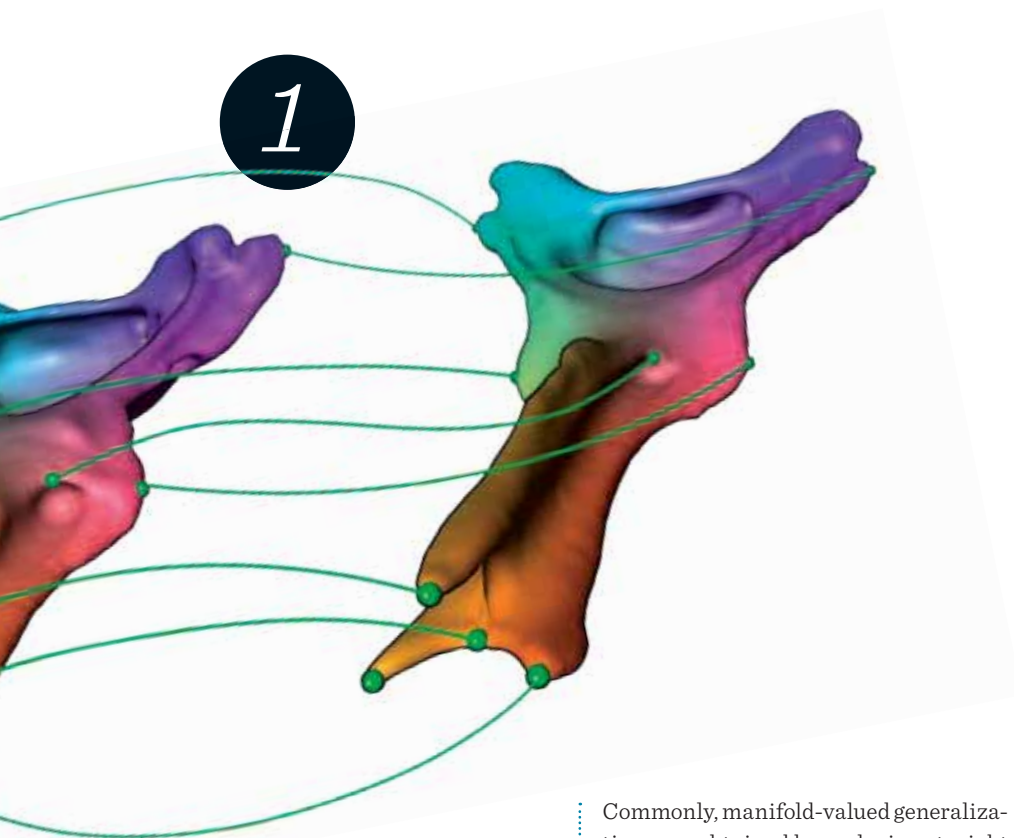
Discretization. To process shapes in the computer, they have to be turned into finite-dimensional objects. There are various ways to do this. In this article we focus on representing them as discrete objects. Examples are representation by finitely many points (point clouds), by a few specific points (landmarks), or by discrete meshes that connect points to form polygons and polyhedra.

Statistical shape models. In many applications, sets of shapes, like collections of clay jugs found in an archaeological field, must be analyzed. Then statistics comes into play. The sets of shapes can be considered as arising from an underlying probability density function that is defined on some shape manifolds. An estimate of this shape manifold and the probability density function, determined from empirically observed occurrences, is called statisti-

cal shape model (SSM). It encodes the probability of occurrence of a certain shape within a given ensemble of shapes. Given a statistical shape model, one can, for example, compute a mean shape or quantify the variability of shapes (Small 2012; Dryden et al. 2016; Srivastava and Klassen 2016).

Correspondence problem. One of the fundamental problems in shape analysis is to find a meaningful relation (or mapping) between their points; this is called the “correspondence problem,” see figure 1. Depending on the application, “correspondence” just means geometric correspondence, defined, for example, by geometric similarities; but often it also means semantic correspondence. Then semantic information has to be considered, too. The task is thus to determine all homologous locations on all shapes in a given set of shapes (Van Kaick et al. 2011). At ZIB, we have developed various approaches to establish highly accurate and dense correspondences for anatomical structures. This ranges from techniques that integrate fundamental expert knowledge using interactive annotation techniques (Lamecker 2008) up to fully-automated techniques (Günther et al. 2013; Grewe and Zachow 2016). Based on this methodology, we have built a software pipeline that allows us to create SSMs for a wide range of applications.





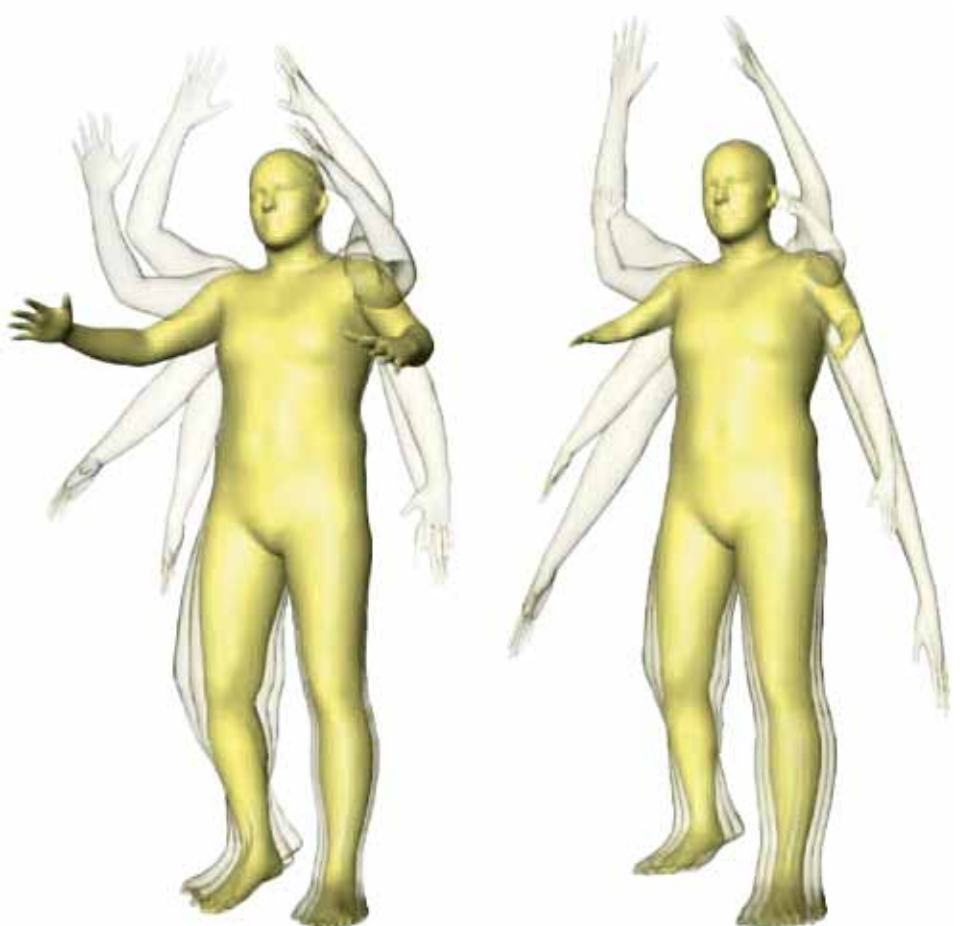
1 Point correspondence between the surfaces of two snake bones. Some corresponding point pairs are shown explicitly by lines. The rest is shown implicitly by color.

2 Mean shapes (opaque) and first principal geodesic curve (transparent at ± 0.75 and ± 1.5 standard deviations) for a data set of 100 human body shapes computed in a curved (left) and in a euclidean shape space (right).

Commonly, manifold-valued generalizations are obtained by replacing straight lines with geodesic curves in the problem formulation, for example, yielding the principal geodesic analysis (see figure 2).

Similarity of shapes. Another fundamental problem is to measure distances between shapes in order to quantify their dissimilarity; this can be considered as measuring distances in the shape manifold, either “intrinsically” along geodesics in the curved manifold, or “extrinsically” along straight lines in the ambient euclidean space in which the shape space is embedded. From an application point of view, however, the similarity and dissimilarity of shapes often is problem-specific; a suitable distance metric thus has to be induced from problem-specific conditions.

Geometric variability and mean. Despite the many methods for capturing the geometric variability in a population, principal component analysis (PCA) and its manifold extensions remain a workhorse for the construction of statistical shape models. PCA, a traditional multivariate statistical tool, determines a hierarchy of major modes explaining the main trends of data variation. As the shape spaces frequently have nontrivial curvature, the statistical analysis needs to account for the nonlinearities in the model.



GEOMETRY RECONSTRUCTION UTILIZING STATISTICAL SHAPE MODELS

Shapes from images. Often, geometric shapes are not captured by geometric measurements but are instead reconstructed from image data. Acquiring highly accurate image data in, for example medicine, biology, or archeology, is often tedious, costly, destructive, or requires patients to be exposed to a significant amount of radiation. To minimize such negative effects during data acquisition, highly sensitive computer-aided approaches are required that extract as much valuable information as possible from the measured data. Challenges for automatic and reliable data processing arise from (a) the high variability in shape and appearance in real-world data, (b) measurement noise due to fast and low-dose acquisition techniques, and (c) the ill-posed reconstruction of 3-D structures from 2-D measurements.

Prior shape knowledge. At ZIB, we develop algorithms that allow the reconstruction of highly accurate geometric representations of medical, biological, and artificial structures as surfaces and volumes from image data. To address the aforementioned challenges, one key strategy is to incorporate prior knowledge on statistical variation of shape and appearance of the measured structures, utilizing statistical shape models and Bayesian inference (Bernard et al. 2017).

In the following, we describe three cases of anatomy reconstruction, where algorithmic utilization of shape knowledge has been instrumental:

Medical image segmentation. One way to create digital shapes of anatomical

structures is to segment the respective anatomy from medical image data and to represent the shapes by surfaces that separate differently labeled volumes. Although the general shape of an organ is known, individual organs differ in shape, either within a normal range of variation or due to a pathology. To fully automate an often labor-intensive process of tissue delineation, advanced algorithms employ SSMs as geometric priors to identify the organ of interest within the image data and to geometrically fit the model to the individual data (for an overview see Lamecker and Zachow 2016).


3-D anatomy from 2-D radiographs.


While MRI is expensive and less suited for imaging bones, CT incurs a significant radiation dose. Using SSMs as strong prior allows the shape of skeletal structures to be reconstructed from few planar radiographs with low doses (Ehlke

et al. 2013, 2015) (see figure 3). The local uncertainty quantification that is possible due to the shape statistics encoded in SSMs does not only give an indication of the reconstruction accuracy, but allows one to utilize methods of experimental design to optimize the imaging setup, defined by number, direction, and dose of radiographs taken, such that a requested geometrical accuracy is achieved with a minimal radiation dose. Since the reconstructed shapes are restricted to the low dimensional space covered by the SSM, pathological cases cannot necessarily be addressed with this method; this requires further research.

Facial shapes and expressions.

For the planning of facial surgery and the development of tools for diagnosis and treatment of psychopathologies, it is of interest to extract the significant morphological patterns in facial shape






ANALYSIS OF SHAPE ENSEMBLES

and appearance from large-scale 3-D face databases. SSMs not only provide a convenient way to analyze the statistical variation of geometric and photometric features; they also allow one to reliably acquire and process large amounts of measured data to capture the great variety in human faces fully automatically. 3-D facial geometry can thus be reconstructed with the highest resolution without the need for manual intervention. Accurate dense correspondence even under high deformations due to facial expressions can be computed up to the level of skin pores by simultaneously exploiting geometric and photometric features (see figure 4). This permits the capture of even the finest details of morphological patterns in human faces and enables statistical examination of the complex deformations of facial tissue on all levels of detail (Grewe and Zachow 2016).

Shape descriptors. Given a SSM, members of a set of shapes can be uniquely encoded within the basis of principal modes of variation. This yields low-dimensional representations that can serve to derive statistical shape descriptors. These shape descriptors can be utilized to perform many shape-analysis tasks, such as investigating the similarity of individual shapes. Shape descriptors not only allow one to compute a distance between individual shapes, but they also allow the changes to be categorized into different modes of variation that can be studied separately, for example, via morphing with respect to selected modes of variation (see figure 5). Thus, shape descriptors enable a detailed study of differences between individual shapes. In the following, we shortly describe two examples of application.

- 3 *Employing statistical shape models allows individual 3-D geometry to be estimated from few planar X-ray images with low radiation exposure.*
- 4 *Matching of a facial surface S to the reference R : parametrizations ϕ_S and ϕ_R are computed and photometric as well as geometric features are mapped to the plane. The dense correspondence mapping $\Psi_{\phi_S \rightarrow \phi_R}$ accurately registers photographic and geometric features from S and R .*
- 5 *Variation of the shapes of a PDM generated from 225 input snake bones along the two first PCA modes.*





6 *Low-dimensional visualization (Sammon projection) of our Riemannian statistical shape descriptor revealing a high degree of separation between healthy and diseased distal femora.*

Evolutionary biology. In collaboration with the Museum für Naturkunde in Berlin, more than 300 data sets of a single head bone of 18 species of the snake genus *Eirenis* were investigated using shape clustering. In particular, hierarchical clustering was applied using the euclidean distance of the shape descriptors resulting from a point distribution model. The clusterings yielded trees comparable to phylogenetic ones, though clear differences were observable indicating that the shape morphologies are not only influenced by genetic but also by functional relations (Baum et al. 2014).

Musculoskeletal medicine. In order to analyze characteristic shape changes incident to radiographic knee osteoarthritis, data of a longitudinal study (performed by the Osteoarthritis Initiative of the NIH, USA) have been analyzed. For these, data shape models and shape descriptors have been derived (see figure 6 for a visualization). Utilizing principal geodesic analysis within a novel Riemannian framework based on Lie groups of differential coordinates, we achieved a highly accurate separation between subpopulations of healthy and diseased patients (von Tycowicz et al. 2016).

A general conclusion from many studies performed at ZIB is that the result of any shape analysis might be severely influenced by the type of statistical shape model being used. This means, one has to take care that the model used is mathematically appropriate and, equally as important, that the metric, which typically is application-specific, has to be chosen with the particular analysis question in mind.



8 The analysis of the evolution of shapes can also be utilized for the stratification of archaeological finds, for example, in application to classification of ancient pottery and different aspects of shape evolution throughout history.

CURRENT AND FUTURE WORK

Major algorithmic advances in shape processing have been achieved over the last few years, allowing us to address many applications. Yet, a number of problems still need to be solved.

Varying topology. Some applications require analyzing sets of shapes with different topologies – which are either intrinsic or emerge as artifacts due to noise and outliers in the acquisition process.

Uncertainties. Another requirement is to consider uncertainties in a coherent way – starting with measurement errors during acquisition, to numerical errors during shape reconstruction and analysis, up to statistical inference.

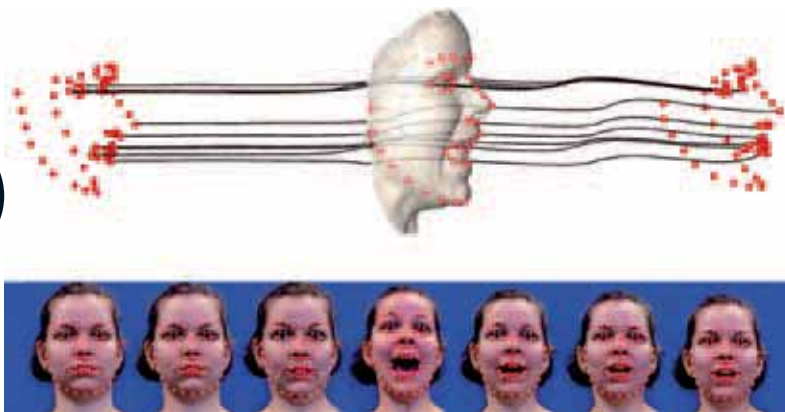
Multiscale representation. A largely unaddressed topic is the multiscale representation of shapes and corresponding techniques of analysis. One such example is the analysis of a collection of shapes representing the outer surfaces of complex insects in high detail.

Problem-adapted descriptors and metrics. Often one is confronted with the task to quantify certain subtle (localized or extended) shape variations while ignoring others. This is related to the particularly important aspect of defining suitable problem-adapted metrics in shape manifolds. Suitable notions of distance (beyond general concepts like elastic deformation energies or norms of flow fields) are hard to conceive. Here, we expect advances from other techniques for nonlinear dimensionality reduction and from interactively controlled metric learning.

Shape trajectories. An important topic is parameter-dependent shapes, especially 1-D shape trajectories, since shapes varying over time occur in many applications. Think, for example, of a beating heart, or transitions between facial expressions (see figure 7). Working with complete shape trajectories instead of single snapshots will improve the segmentation of 4-D medical images with low resolution, or enable certain functional defects to be identified by a classification of the deformation

trajectories. Furthermore, analysis of anatomical changes in longitudinal studies will help to characterize progression of diseases. Time-dependent shape analyzes will become even more interesting when covarying factors are considered. Another example is the analysis of the evolution of shapes, for example in pottery in the ancient world (see figure 8). Currently, we are in close discussion with archaeologists about analysis tasks that could be supported by mathematical shape analysis, such as stratification of archaeological finds. In general, we expect shape-trajectory analysis to bear great potential for all areas of digital humanities that deal with 3-D objects and their changes throughout history.

Computational efficiency. A last but important topic is the reduction of computational effort. The handling of nonlinearities of shape manifolds (“Riemannian computing”) is computationally expensive. Therefore, it is important to transfer computation-saving concepts of numerical mathematics, such as adaptive discretizations and multilevel solvers, to shape computing.



7 As high-resolution facial geometries can only be acquired at high latencies, facial trajectories have to be reconstructed from noisy and low-resolution video data in a manner consistent with physical and dynamical constraints.



Dr. Janina Zittel | zittel@zib.de | +49-30-84185-146

RESEARCH CAMPUS MODAL: REPORT FROM THE GASLAB

DEREGULATION OF THE ENERGY MARKET – NEW CHALLENGES FOR NATURAL GAS TRANSMISSION OPERATORS

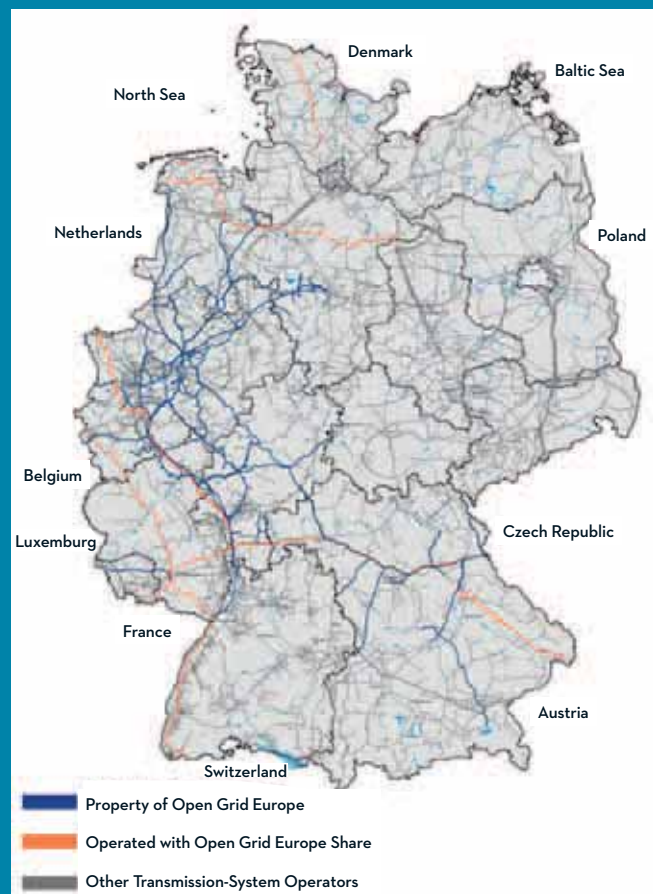
About 20% of the German (and European) energy demand is met by natural gas. Until 2005, gas transport and supply in Europe was provided by a handful of companies, owning and operating the natural gas transmission system to do so. To establish a European gas market, in 2005, the European Union legislated that gas trading and transport had to be done by mutual completely independent compa-

nies to ensure discrimination-free access to the transport network for all traders (Geißler et al. 2015). This changed the operation and business model of the gas transmission system operators (TSOs) that – due to their high investment cost for natural gas pipeline networks – are mostly a natural monopoly. Before this, TSOs were part of an integrated organization and could plan the network operation and expansion together with the

traders. Since then, they are independent and need to plan under the uncertainty regarding the gas-flow situations resulting from trading. This leads to several challenges which require a complex decision support system, including technical operation support and decision support on various nontechnical intervention options.

NATURAL-GAS TRANSMISSION- SYSTEM

The Research Campus Modal industry partner Open Grid Europe GmbH (OGE), located in Essen, Germany, operates about 12,000 kilometers of natural gas pipelines in Germany, a network of pipes roughly as long as the German autobahn network. About two thirds of the natural gas transported in the network is supplied by or consumed in other countries. Therefore, securing the transport requirements in Germany is key for the European cross-border natural-gas supply.



THE VIRTUAL TRADING POINT

In a simplified view, the EU regulations stipulate that the so-called entry/exit model with a virtual trading point is the basis of the capacity market. The virtual trading point is an abstraction of the physical gas network, decoupling the direct local connection between a customer demanding natural gas and a seller supplying natural gas. Both parties, the buyer and the seller, are regarded as transport customers of the TSOs that operate the network represented by the virtual trading point. These transport customers need to buy capacity from the transmission system operator to access the virtual trading point, that is, suppliers need to buy entry capacity, demand customers need to buy exit capacity. Thus, a seller can supply gas at any physical point of the network that a buyer demands at any other physical point of the network.

There are different capacity products available to the transport customers,

guaranteeing the right to supply/extract gas to/from the network with different constraints. The preferred capacity product sold by the TSO is freely allocable, which basically means that the transport customer can trade with any other market participant up to the specified amount of gas. After trading, the amount of capacity used is communicated to the TSO (the nomination). The requirement for the TSO is to be able to transport this gas. In contrast, employing another product, the interruptible capacity, the transport customer buys the right for transport, which, however, can be interrupted by the TSO if the network situation does not allow the respective transport.

A TSO may only sell capacity rights for which it can guarantee that each “likely and realistic” (Gas, 2010, section 9) gas flow complying with the capacity rights booked by all transport customers can technically be realized.



© ZIB

THE MODAL GASLAB FROM SCIENCE TO A DECISION SUPPORT SYSTEM

The MODAL GasLab aims at the realization of fundamentally new possibilities for transmission network operations beyond current feasibilities. The software developed within the GasLab addresses the main challenges of natural-gas dispatching by providing foresighted decision support. To achieve this ambitious goal, fundamental research is needed. The research focuses on:

- Mathematical models integrating the technical network description, such as the transient behavior of the gas transmission network and its technical elements, and nontechnical operation measures, such as capacity limitations.
- Machine learning solutions to predict gas supply and demand in the operation horizon of up to two days.
- Solution algorithms for the respective optimization problems for efficient real-time decision support.

The GasLab does the whole chain from fundamental research to implementation. Therefore, in the framework of the Research Campus MODAL, the GasLab involves several cooperations. First of all, the cooperation with the industrial partner Open Grid Europe (OGE) forms the basis for all research and development in the GasLab. OGE experts provide industry-specific knowledge on the technical and nontechnical operation measures and how these measures are used for efficient operations by the dispatchers. Insights on specific knowledge of technical elements, such as compressor stations, are fundamental for the definition of mathematical models. Furthermore, OGE provides real-world data used as research data sets, and test data for new methods insures the applicability of any developments of the MODAL GasLab. In addition, OGE subcontracted the software-development company Soptim to implement an industry product including interfaces to data source systems of OGE, a graphical user interface for the dispatchers, and a test suite to support the development of the MODAL GasLab.

1 *The GasLab team - the GasLab team includes researchers at different career stages led by Prof. Thorsten Koch, experts from OGE, and software developers at Soptim.*

The MODAL AG acts as a bridge between industrial needs, such as maintainability and stability of software solutions, and priorities of ZIB, such as basic research on fundamental problems and scientific publication of the corresponding results. Likewise, the MODAL AG helps to separate scientific and economic risks, creating a win-win situation for both ZIB and its industrial partners.

The incorporation as one lab of the Research Campus MODAL is fundamental for the success of the project due to the synergies with the other labs creating methodological advances. For example, joint research on machine learning methods together with the MedLab stipulates advances in prediction of gas flows and medical diagnostics alike and can be generalized for multiple purposes. Likewise, joint work with the SynLab on solving the optimization problems of the GasLab creates advances on solvers. On one hand, the GasLab provides challenging problems, driving progress in mathematical programming solvers. On the other hand, advanced solver technologies reflect back on the GasLab.

SOFTWARE SOLUTIONS TO SUPPORT DECISION-MAKING

The changing market environment raises several new challenges for the TSOs from everyday operations up to strategic network planning. Three key challenges for operations are tackled in the MODAL GasLab, providing foresighted decision support on:

- Efficient technical operations of the gas-transport network.
- Ordering additional gas to overcome temporal imbalances in the network.
- Employing a new capacity product for secure supply to power stations of systemic importance.

The aim is to integrate technical and nontechnical operation measures into one tool for efficient network operations. However, solutions for the nontechnical operation measures may be developed separately in a first step. Later on, they will be integrated in one holistic tool.

THE BALANCING TOOL (BILANZTOOL)

The TSO may assume that over a full day the result of market operation is balanced in the sense that the amount nominated at entries is a sum equal to the amount nominated at exits. However, temporal and spatial lacks may appear. For example, gas may be extracted in the morning in the south and injected in the evening in the north. In order to satisfy the transport requirements, the TSO is first obliged to exploit all technical operation measures. Only, when this does not suffice, several nontechnical measures can be taken.


One of the main problems arises from the varying distribution of supply and demand in space and time. Before the deregulation of the European gas market, such problems were solved by storage units controlled by the gas company. Nowadays, to secure the transport requirements in these situations, there are two options for TSOs: to use the line buffer or to order balancing energy.

- Line buffer: To a limited extent, imbalances can be buffered using the gas that is already in the network.
- Balancing energy (Regelenergie): When line buffer is not sufficient to secure network operations, the TSO can buy or sell gas to balance physical differences between supply and demand.

Commercially, this is organized by grouping network nodes into “balancing zones” (Bilanzkreise) for which in- and outflows are supposed to be balanced. An important task of the operator of a balancing zone is to proactively manage the transport of the zone in order to avoid the demand for balancing energy. One possibility is to exchange surpluses with other balancing zones. Any additional and unavoidable balancing energy demand is bought or sold at the entries and exits of the virtual trading point (see box) in a nondiscriminatory manner.

The MODAL GasLab supports these decisions by solving several mathematical problems:

- Forecasting temporal and spatial imbalances of in- and outflows of the network.
- Computing gas flows in the network to tracking down temporal and spatial line-buffer capacities.
- Determining the unavoidable amount of balancing energy needed.



2 The sketch shows a simplified network situation in which the demand (3 units at the power plant) and supply (three entries with 1 unit each) is globally balanced. However, due to construction at one pipe and a capacity bound of 2 units at an alternative pipe, there is no feasible transport solution to the power plant (left image). Once the power plant is restricted to obtain its supply from the Ausgleichsentry, there is a feasible transport solution (right above).

THE KWP TOOL

One particular set of exits of the network corresponds to gas-fired power stations. Some of these stations are system-relevant for the electricity system, especially due to the increasing amount of renewables in electricity generation related to the German “Energiewende”. Gas-fired power stations have high use of gas during their hours of operation, which is highly variable depending on other electricity-generation capacities. In general, the supply of gas to the stations may not be met with freely allocable capacities. One solution to this problem would be to build new pipelines and compressor facilities to ensure supply to this critical infrastructure.

However, to avoid these infrastructure investments, the TSOs will soon have a new regulatory option to ensure the secure operation, related to the so-called “Kraftwerksprodukt” (KWP). This option may be executed only in bottleneck situations. In these situations, the TSO can force the customer corresponding to the power station to buy its gas from one specific entry (Ausgleichsentry) instead of any entry point of the virtual trading point. The appropriate Ausgleichsentry is specific for each power station and specified in the capacity contract. Likewise,

the gas supplier at the Ausgleichsentry is obliged to reserve sufficient capacities. When several of these restrictions would resolve the bottleneck situation, the actions of the TSO have to be nondiscriminatory among the available power plants.

The GasLab aims at the development of a tool – the KWP Tool – that calculates a foresighted, nondiscriminatory decision suggestion for the TSO for which power stations the restriction needs to be executed for the upcoming gas day.

To tackle this challenge, several mathematical problems need to be solved:

- Predicting the gas flow at the entries and exits of the network for planning horizon.
- Identifying bottleneck situations from the spatial and temporal distribution of the gas in the network.
- Analyzing whether or not a subset of power plants exists, whose restrictions admit a feasible network flows.
- Finding a smallest subset of power plants, w.r.t. the cardinality, whose restrictions admit a feasible network flow.

THE NAVI

Gas network operation is a challenging task. In practice, dispatchers operate the network and are responsible for different parts of it. Dispatchers carry great responsibility, as many technical, legal, and regulatory constraints have to be fulfilled. Nowadays, the job requires long experience assisted by limited software solutions mainly based on simulations.

To support the dispatchers at the TSO during everyday operations, a comprehensive decision support system is needed. This tool is supposed to assist the dispatchers by addressing nontechnical measures and technical decisions providing the recommendations for changes in the network-configuration. On one hand, it is expected to support everyday operations; on the other hand, it is a useful tool to decrease the training period of new dispatchers.

In the MODAL GasLab, the dispatcher Navi – like a navigation system for gas network operators – is being developed. Mathematical models serve as the basis for optimization solutions, describing the technical elements of the network and their control options. Some of them, for example the outgoing pressure of a

compressor station, lead to continuous operation options, while some of them, for example opening or closing a valve, lead to discrete operation options. Additionally, rules apply to the order of interventions. First, the dispatcher should use all technical interventions, then, if necessary, balancing energy. If this still is not sufficient to secure operations, capacity limitations may be applied as a last resort, ranging from KWP to cutting supply or demand on certain entries or exits.

For efficient operations, the dispatcher has to meet additional objectives. Just like the driver of a car has secondary goals such as the fastest or shortest route, efficient transmission-system operation involves several goals, such as:

- Operation as far away from technical bounds as possible.
- Robustness to late changes in the transport requirements.
- Minimal fuel use of compressor stations.
- Restricting the number of interventions.

For the decision support system Navi that reflects the technical details, the hierarchical nature of nontechnical interventions, and the additional objective functions, basic research on novel mathematical models and solutions approaches has to be the first step. To face these challenges, several different approaches should be considered and evaluated. The research involves, but is not limited to, the following:

- Predicting future gas flow for foresighted operations.
- Finding a set of technical and nontechnical measures sufficient to fulfill the transport requirements respecting the regulatory terms.
- Computing optimal measures that achieve the different operation objectives.

THE TIP OF THE ICEBERG – REAL-WORLD DATA FOR RESEARCH AND DEVELOPMENT


Research in the MODAL Labs relies on real-world data. In the context of the short-term operations of gas networks, data from very different areas have to be considered and combined. This includes:

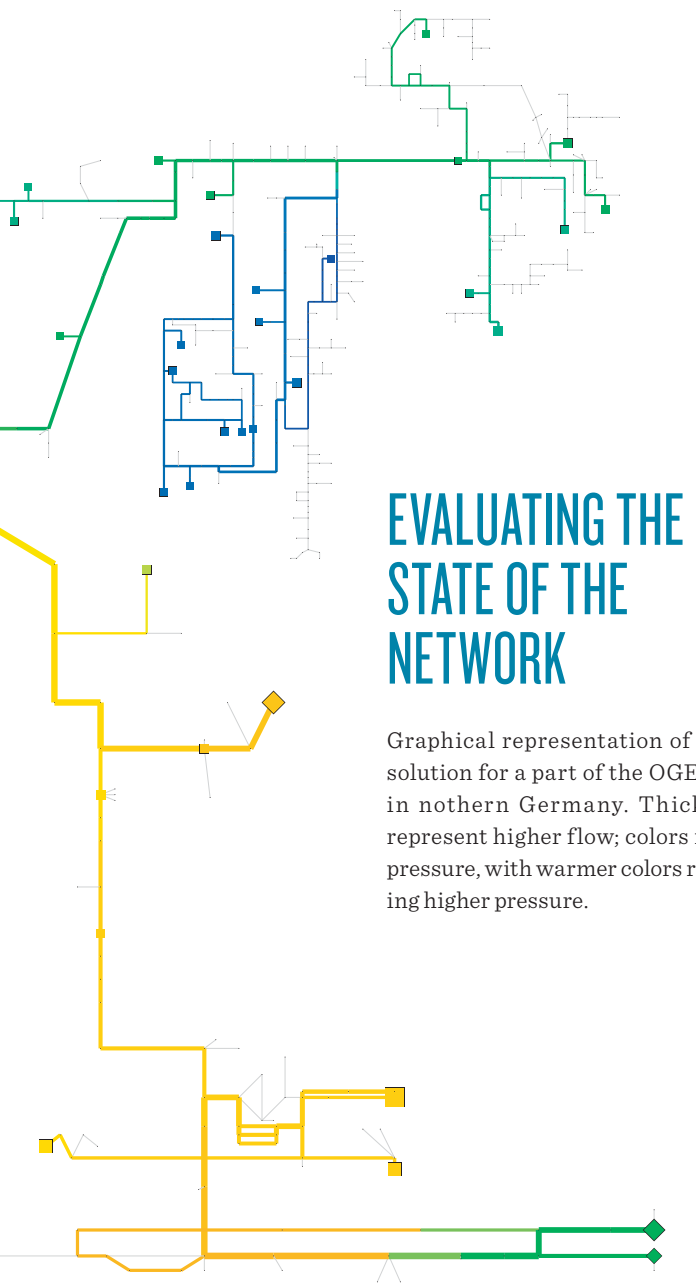
- technical network specifications.
- real-time measurements of physical gas properties.
- real-time network-configuration information.
- various nontechnical intervention options

Gas networks are complex structures consisting mainly of pipelines but containing many special devices such as compressor stations. Physics of pipelines with their relevant parameters is understood in theory. In practice, however, these parameters can often only be estimated or derived from measurements, and vary over time. Compressor stations are typically the most complex facilities dispatchers have to handle. They often not only comprise several machines, but also extensive piping and other special network equipment that is sometimes unique in the network. These stations have been built and extended over time, and thus the equipment within one station often has very different technical

specifications. Moreover, due to aging and modernizations, the technical specification is known at the time of delivery of the hardware, but changes over the years. Taking all these possibilities into account in new data formats is a major challenge that is iteratively approached in the MODAL GasLab.

One major challenge was the integration of data from several sources. These data sources were initially deployed for different purposes in the software landscape of the industrial partner, and therefore use different time scales and abstractions of the network. Within the GasLab coherent structures for both, stationary data, such as the network topology, and transient data, such as the simulated or measured current state of gas properties, have been developed including transformation rules for the different abstractions of the network. A new comprehensive data model is the result. A great effort is needed to integrate this data model into the legacy systems at OGE to get real-world data for the project. Powerful interfaces to the dispatching system of OGE were implemented to generate research data sets, test new methods on live data, and later productive use of our tools.





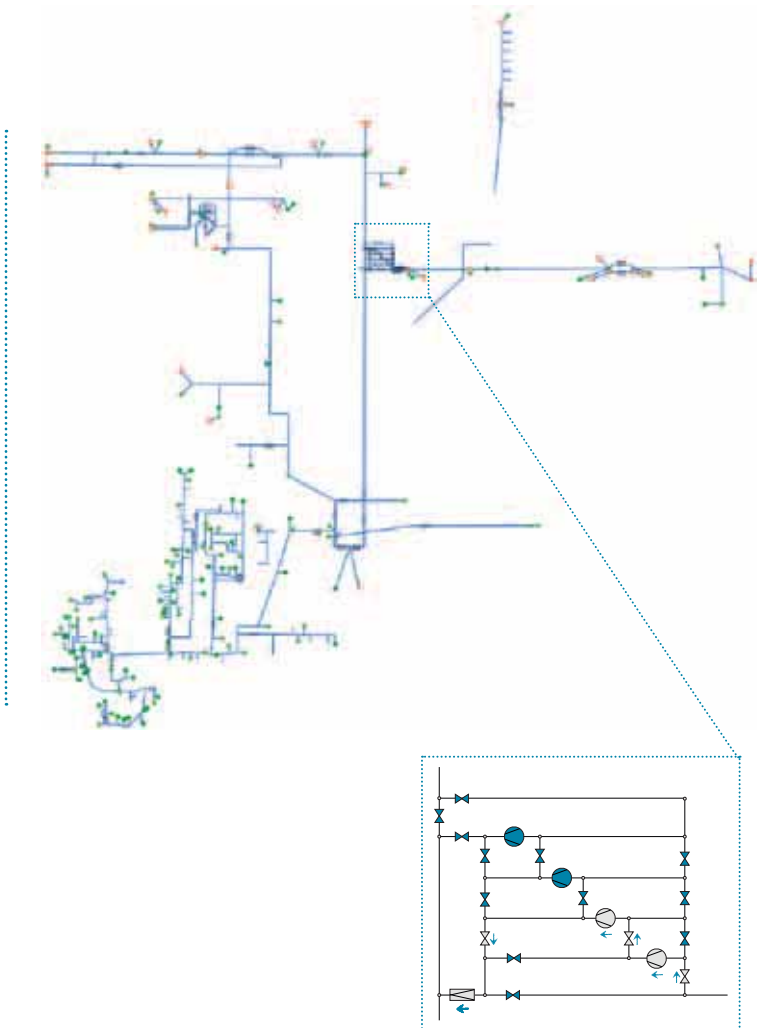
EVALUATING THE STATE OF THE NETWORK

Graphical representation of a sample solution for a part of the OGE network in northern Germany. Thicker lines represent higher flow; colors illustrate pressure, with warmer colors representing higher pressure.

3



3 Compressor station



RESEARCH ADVANCES IN FORECASTING GAS FLOW

The need to forecast gas in- and outflows is common to all tasks tackled in the MODAL GasLab. Expert knowledge and historical data are the basis for such a forecast system. Both support the analysis of the complex market structure, which includes different behavior of the various types of customers. Only about a third of the gas in the network is consumed in Germany for different purposes and with very different demand profiles. Domestic heating depends mainly on outside temperatures. The consumption of industrial complexes varies from high during working hours to low on weekends or public holidays. Gas-fired power plants depend on the availability of other energy sources, in particular renewables. The remaining two thirds of the gas are transported through the network, for example from Eastern Europe to France. These transport customers are obliged to “nominate,” that is declare the capacity needs beforehand, but might “renominate,” that is change the declared amount, during the course of the day. For nominations and renominations, a complex set of rules

applies. All of these different characteristics and constraints have to be taken into account in a forecast model.

Developing the forecast model is in itself a demanding research project, which we have tackled with an international group of researchers. Together we investigated the application of several methods from neural networks over time series analysis to optimization-based multi-regression and combinations of these methods to come up with the best forecast model. For example, Prof. Ying Chen (University of Singapore, Singapore) visited ZIB for a month to apply an adaptive functional autoregressive model, previously used for forecasting electricity prices (Chen and Li, 2015), to the gas flow forecast problem. Several scientists visited ZIB for short-term scientific missions supported by the COST Action TD1207. Dr. Selini Hadjidimitriu (University of Modena and Reggio Emilia, ICOOR, Italy) employed neural networks using external variables such as European temperature data to predict gas demands. Dr. Andrea

Taverna (University of Milan, Italy) applied several outlier detection techniques. Dr. Milena Petkovic (University of Novi Sad, Serbia) used support vector machines to predict days where certain entries and exits are not used at all, a question relevant for, for example, big power plants or cross-border connection points. Currently, we are working on a joint publication comparing the results of the different methods and useful combinations to increase the quality of gas-flow forecasts. The optimization-based multi-regression (see box) is implemented in software integrated at OGE to test its performance on live data. Further development of this method and the integration of additional methods will be researched and implemented in the remainder of the project.

The challenges, described in this feature article, present only some of the aspects tackled in the MODAL GasLab on our way toward a comprehensive decision support system for gas network operations.

OPTIMIZATION-BASED MULTI-REGRESSION

This prognosis method is a regression approach with a rolling horizon. The prediction model is trained using a comparably short historical time period, which moves when we are moving forward in time.

In the training phase, we calculate a function which describes the given historical time series best.

This is done by optimal choice of coefficients (weights) of a function with predefined terms (features), which lead to the smallest error over the training horizon.

For instance, the predefined terms could be the value of a former hour (such as 24 hours ago), the mean of the values of a former day (such as yesterday), or the ratio between different values (such as the ratio between the first value of yesterday and the first value of the day before yesterday).

Then the future values of the time series are predicted using the calculated function.

However, when using too many terms with a small training data set, the function can become overfitted, that is, it perfectly represents the past but is very poor in predicting the future. Since it is difficult to choose the terms a priori, we plan to extend the method by adding more terms and forcing the training algorithm to choose the coefficients in a way such that, at most, k -coefficients are nonzero. Although this makes the training problem more difficult, we hope that it leads to considerably better results.

The optimization-based multi-regression can be used for predicting all kinds of time series.

A special case for the GasLab are time series, which are somehow connected (e.g. time series of demands and supplies at several entries and exits of the gas network). Here, it is promising to also take information of time series of other entries/exits (i.e. boundary nodes of the network) into account when forecasting the time series of a specific node.



Example of gas flow forecasts using optimization-based multi-regression.

Schematic plot (without scale) depicting the forecast (orange) for a specific exit of the gas network supplying an industrial customer compared to measured values (green); hourly forecasts started at the beginning of each day (vertical grey lines); weekends displayed with gray background shade.



Prof. Dr. Ralf Borndörfer | borndoefer@zib.de | +49-30-84185-243




A Man–Machine Search for the Absolute Evidence

Jonad Pulaj



© BenFrantzDale, commons/wiki

COMPUTING THE TRUTH




Using computers to prove results in pure mathematics is not a new endeavor, although it still remains controversial amongst some mathematicians. After all, how are we to trust results if they cannot be directly checked by hand? This question quickly turns into a philosophical one and has little to do with mathematics itself. Suppose that you are reading a proof of a statement, that is, a series of connected logical arguments that begin with some assumptions and end with the desired result.

Is it possible to know with absolute certainty that no mistakes have been made? Arguably, for proofs that are quite complicated and go on for many pages, the answer is simply no. However, it is often enough to be "sufficiently convinced." For example, the theorem which classifies finite simple groups is over ten-thousand pages long. Reading the entire proof would take a long time, and a great deal of optimism is necessary to assume that absolutely no mistake could be made in the writing, reading, and understanding of each crux in the proof. Nearly all mathematicians accept the result as true, although rationally speaking

it is very likely that only a minority has carefully read the whole proof (many may have read portions of it). Indeed, we may safely assume that situations like this are normal, as progress measured in the traditional mathematical way often consists

of results building upon previous results. The Four-Color Theorem is the first major result that was proved via a computer-assisted method in the 1970s by Appel and Haken. Its statement is



deceptively simple: any map in the plane can be colored with four colors such that the regions sharing a common boundary (other than a single point) do not share the same color. Nevertheless, a “simple” proof of the statement is still missing,

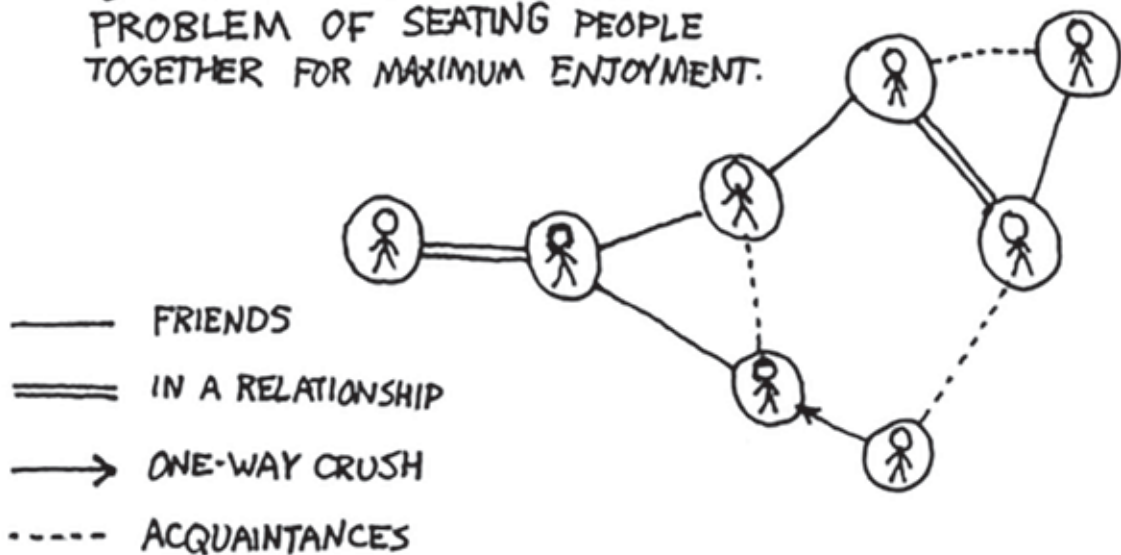


while numerous attempts date back to the 1850s. Indeed, in 1879, Kempe published his first “proof” of the four-color conjecture. As a result, the English lawyer and mathematician was elected a fellow of the Royal Society and even knighted at a later date. However, in 1890 Heawood showed that Kempe’s arguments contained a nontrivial flaw, while at the same time proving that five colors are sufficient for every map. Throughout the decades, many other mathematicians worked on the four-color conjecture, showing that the problem could be reduced to some finite, albeit large, number of cases. Finally, this culminated in the computer-assisted proof of Haken and Appel, which took over 1,200 hours of machine time. The proof was initially met with skepticism, but was eventually (mostly) accepted by the mathematical community. Several other implementations and computer-assisted verifications have reaffirmed the Four-Color Theorem. Along these lines, the final verification in 2005 by Gonthier [5], via the interactive theorem prover Coq, put (nearly) all mathematical fears to rest. Indeed, in some sense, interactive theorem provers such as Coq or Isabelle/HOL represent the highest levels of trusted code for computer-assisted proofs. This is because interactive theorem provers

such as Coq are partly based on type theory invented by Russell at the beginning of the 20th century in order to deal with the logical paradoxes of naive set theory as a foundation for mathematics. Thus, interactive theorem provers serve as a rigorous framework for the formalization of the discipline. Of course, in practical terms, such rigor comes with significant costs that may alienate the average mathematician from this approach in the first place. Formalizing results in interactive theorem provers is incredibly time consuming (note the nearly 30-year gap between the result of Gonthier and that of Haken and Appel) since the implementations in functional programming can be particularly daunting for those not familiar with this paradigm.

Still, for the new generations of mathematicians, Coq and Isabelle/HOL are considered (somehow) well established but rather tedious tools. Interactive theorem provers have shown considerable flexibility, ranging from the formalization of Kepler’s conjecture to results beyond the realm of traditional mathematics. Indeed, the formalization of the ontological proof of God’s existence [1], which can be traced back to Leibniz, is a striking example of computational metaphysics and the far reach of interactive theorem proving.

AT THE MOVIES, I GET FRUSTRATED WHEN WE FILE INTO OUR ROW HAPHAZARDLY, IGNORING THE COMPUTATIONALLY DIFFICULT PROBLEM OF SEATING PEOPLE TOGETHER FOR MAXIMUM ENJOYMENT.



GUYS! THIS IS NOT SOCIALLY OPTIMAL!



Rational Expressions

'ROUNDING ERROR'

BY IVSCHEE



In theory, the methodology of integer programming can also be used to (computationally) prove suitable results of interest. This is due to the flexibility of integer programming as a modeling paradigm. Consider, for example, a challenging sudoku puzzle. Given the nine-by-nine square grid with some fixed values, the key idea is to consider a corresponding cubic nine-by-nine-by-nine array of binary values. We can visualize this as nine square grids stacked on top of each other. The top grid will be assigned a one whenever the solution has a one in the corresponding square. The grid right below it will be assigned a one whenever the solution has a two in the corresponding square, and so on. Thus, we arrive at 0-1 decision variables that correspond to each empty square in each of the nine grids, and it is straightforward to encode all the rules of the game as linear

constraints in the considered decision variables. The objective function is not important, since we are only interested in a feasible solution. By identifying conditions that need to be shown feasibly or infeasibly in a theoretical problem of interest, we may consider similar “abstract” sudoku puzzles. As we will see later on, such conditions are sometimes easy to identify but difficult to compute, as is the case with most Ramsey theoretic numbers. Yet, in other contexts, such conditions are highly nontrivial to identify in the first place, as is the case with questions related to Frankl’s conjecture. Nevertheless, almost any situation of interest can be modeled by an integer program.

Although the use of computational integer programming is well established in applied discrete mathematics

and operations research, its use in pure mathematics is very limited. Given the high burden of proof in establishing theoretical results, the majority of integer programming solvers are simply not good enough without further safeguards or verification routines. This is because most solvers suffer from a dual curse: the possibility of ill-conditioned matrices and floating-point arithmetic can lead to wrong results due to rounding errors, in addition to the ever-present possibility of a programming error.

This situation has led to the extensive use of SAT solvers by researchers in pertinent areas of pure mathematics. SAT solvers provide an even higher level of trust by producing certificate formats that can be checked by interactive theorem provers. Furthermore, in the past few years, all major SAT competitions require participating solvers to produce a certificate of infeasibility which can be checked with DRAT-trim [8]. Perhaps the most spectacular and controversial use of SAT solvers is the recent proof (in the negative) of the boolean Pythagorean triples [4]. The problem asks whether it is possible to color the positive integers red and blue such that no three integers a , b , and c that satisfy the Pythagorean theorem are of the same color. With a proof size of about 200 TB, this result must keep many a pen-and-paper traditionalist up at night!



In the optimization department of Zuse Institute Berlin, researchers have long been interested in safe numerical computations. This led to the development of an exact rational integer programming solver [3] that avoids the trouble that comes with floating-point arithmetic. Yet, the possibility of a programming error remains. In this regard, the recent development of VIPR [2], a tool which verifies the branch and bound tree produced by an integer programming solver, provides solver-independent verification of integer programming results and needed redundancy for theoretical results. All these tools are crucial in the results that we feature in the next few paragraphs.


So what kind of problems in pure mathematics do researchers at ZIB work on? Currently, the area of focus is extremal combinatorics. Roughly speaking, extremal combinatorics is concerned with how large or small a collection of objects can be before a certain pattern of interest appears. The simplest example of this notion can be explained by the

picture above. Suppose we are given the nine boxes above and want to know how many pigeons it takes before a box contains more than one pigeon? Clearly, the answer is ten. This illustrates a well-known mathematical notion, namely the pigeonhole principle. As simple as this principle may sound, its applications are far reaching and far from trivial.

Ramsey theory may be thought of as a generalization of the pigeonhole principle, and van der Waerden numbers are classical objects in Ramsey theory. $W(r,k)$ is the smallest integer M such that any coloring of $\{1,2,\dots,M\}$ with r colors, contains a monochromatic arithmetic progression of k integers. It is easy to see that $W(1,k)$ is simply k , since the only possible coloring is the trivial one. Furthermore, $W(r,2) = r+1$, since if we are to avoid a monochromatic arithmetic progression of two integers, we must color each integer with only one of the given colors. Thus, given a sequence of $r+1$ integers, two of them must have the same color by the pigeonhole principle, where the pigeons are the integers and

the boxes are the colors. For example, given the colors red, blue, green, and yellow, we may color 1 as red, 2 as blue, 3 as green, and finally 4 as yellow. Thus, it is clear that any coloring of the first five integers with four colors cannot avoid a monochromatic arithmetic progression of two integers. Unfortunately, these were already all the nontrivial cases. Only seven other van der Waerden numbers are known, and as other Ramsey theoretic numbers they are notoriously difficult to compute. Here is what the late Paul Erdős had to say about the computation of Ramsey numbers, which also applies to van der Waerden numbers:

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



Using an integer programming approach, researchers at ZIB are able to give the following, previously unknown result: $W(7,3) \geq 258$ (Pulaj 2015). The picture below is a visualization of a coloring that verifies this lower bound. Integer programming for van der Waerden numbers, and extremal combinatorics in general, is beneficial because polyhedral theory is well-developed but not typically used in this context. Thus, the geometry of optimization problems may lend its own structure to questions of interest and shed light on otherwise hidden structures.



Another problem of interest is Frankl's conjecture. Frankl's conjecture concerns union-closed families of sets. A nonempty finite family of distinct finite sets is union-closed (UC) if, and only if, for every two sets in the family, their union is also in the family. Frankl's conjecture states that any UC family containing an element that is in at least half the sets of the given family. The conjecture appears to have little structure and has evaded all attempts at a solution for nearly 40 years. Frankl's conjecture was first noticed by Peter Frankl, a well-known mathematician and juggler that currently has some celebrity status in Japan.

Although, in this article, we lay out a strong case for the effectiveness of computer-assisted theorem proving such an approach is by no means guaranteed to yield answers to all mathematical questions. Indeed, the limits of formal systems with basic arithmetic were already clearly shown by Gödel in 1931. Simply put, Gödel shows that there will always be true statements in mathematics that cannot be proven. Furthermore, Turing showed that it is not the case that all tasks that are mathematically well defined can be performed by a computer – not even in theory.

Although Turing proved that computer-assisted (mathematical) super-optimism is false, it is clear that applications of such methods are still in their infancy. Therefore, it stands to reason that computer-assisted methods will only become more commonplace in theoretical mathematics, and, a few decades from now, the use of computers for either checking correctness or discovering new mathematics will gain wider acceptance in the community. It is not difficult to imagine that one day the use of computers in all mathematics will be as ordinary as using LaTeX. For all the future mathematical purists who believe that computers have little to contribute to the field, we highly recommend a second specialization like that of Peter Frankl, just in case computers prove them wrong.



Florian Schintke | schintke@zib.de | +49-30-84185-306

PERFORMANCE IN DATA SCIENCE



Dig Deep to Uncover Treasures

Data handling is often the main performance bottleneck in big-data applications. Understanding the distributed data-access deep inside the architecture of a data-analysis workflow usually reveals possible improvements of the data layout and distribution. In two practical use cases from satellite imagery and blood spectrography, we illustrate our approach to monitor and improve the data-access patterns, layout, and placement. This provides interesting insights into the data management of analysis frameworks and allows a reduced time to solution.

BIG DATA FROM SPACE

© NASA

DATA-ANALYSIS PIPELINE

Several earth observation missions have been collecting satellite imagery of our planet for many years, such as Landsat, Sentinel, SPOT, RapidEye, or WorldView, to just name a few. Their sensors differ in the spatial resolution, the explored frequency bands, the repeat-pass intervals, and their spherical projection model, making data analysis across multiple sensors a challenging task. In the BMBF-funded project GeoMultiSens, GFZ Potsdam, Humboldt-Universität zu Berlin, and ZIB have built a big data-processing pipeline for automated multisensoral data analysis. The pipeline includes the management of large amounts of data in the tera-to-petabyte scale, the integration of heterogeneous geospatial data in a common reference model, and the parallel analysis of large geospatial data. It is framed by visual exploration tools for the available input data and analysis results to facilitate the correct detection and assessment of spatial and temporal changes of the earth's surface.

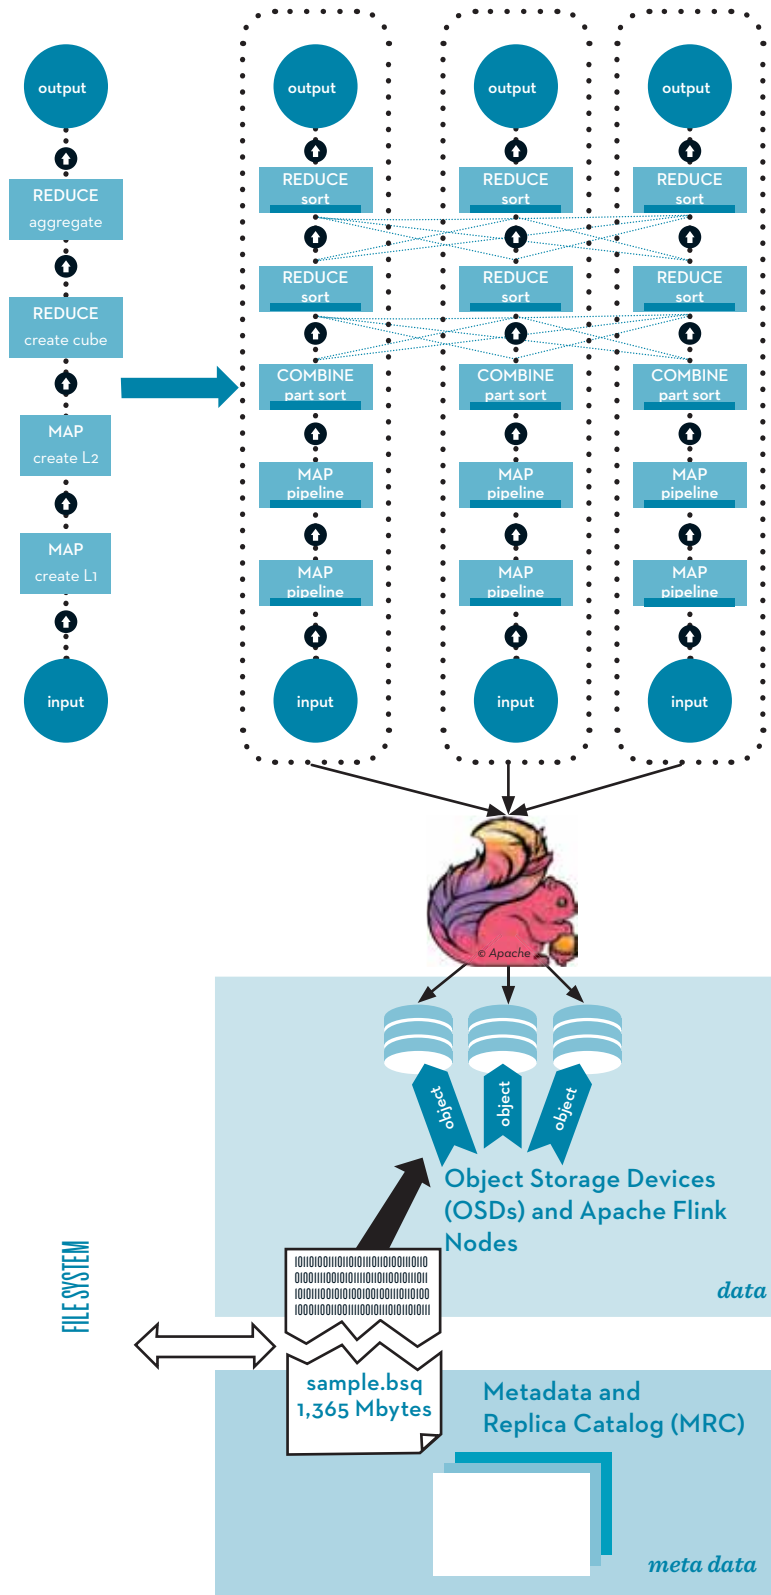
The data analysis in GeoMultiSens is based on the Apache Flink data-analytics framework, which is used at different levels throughout the pipeline. When the user has selected the region and data sources of interest in a Web interface, the raw imagery data are downloaded from the data servers of the particular missions, if it has not already been done for a previous analysis job. Then, the data are homogenized and prepared for the analysis by a first multistep Flink job that geometrically aligns and stitches the tiles, corrects atmospheric distortion, unifies lighting

conditions, eliminates shadows, detects clouds, takes observation angle into account, and so on. Instead of choosing the coarsest resolution of all involved sensors, extrapolation with a separate error-probability layer is used. This results in a homogeneous data cube that represents a virtual sensor consisting of the data of multiple input sensors over time. It can be used for the actual data analysis, such as the classification of land usage and its change over time – which is another Flink job. Finally, the results can be explored in a visual-data-analytics Web interface.



SOFTWARE ARCHITECTURE


1



The analysis tasks are user-defined functions written in Python and Flink jobs that express the required data flow. This description is independent of the number of nodes and enables horizontal scalability of the system. The execution environment is a Flink compute cluster where each node consists of some CPUs, RAM, local disks, and a network interface to the other nodes. Typically, such clusters use the Hadoop File System (HDFS) that spreads the data across the disks of the cluster nodes. In the case of GeoMultiSens, we use ZIB's distributed file system XtreamFS at the lowest level. As XtreamFS offers [2] both a POSIX-compliant interface and the HDFS interface, the input data do not have to be staged in before starting the job nor do the output data have to be staged out after the job is finished.

1 Data-storage scheme and software architecture of GeoMultiSens based on Apache Flink.

2



LOCALITY PARADIGM

Clusters built for big-data applications typically use commodity hardware; single nodes serve as storage and compute nodes. The questions “Where to store file x?” and “Where to process file x?” come up. As the network layer is often the main bottleneck, the data-affinity paradigm is fundamental in the design of efficient data-processing pipelines: files should be processed where the data are stored. In general, data should be transferred between nodes as little as possible.

INPUT SPLITS

In data-analysis frameworks, an input split denotes the smallest portion of data whose processing cannot be parallelized, that is, must be processed on the same node. Typically, one such split is one file but must be written once and never written again.

In GeoMultiSens, however, satellite images are grouped into input splits that correspond to the same geographic region. Each image is cut into multiple files; more images may be added later. To achieve the aforementioned data affinity, it is required that all files associated with a certain region are located on the same node. Together with the possibility to add further data later on, the proper placement of the input data becomes a challenging task that requires new techniques.

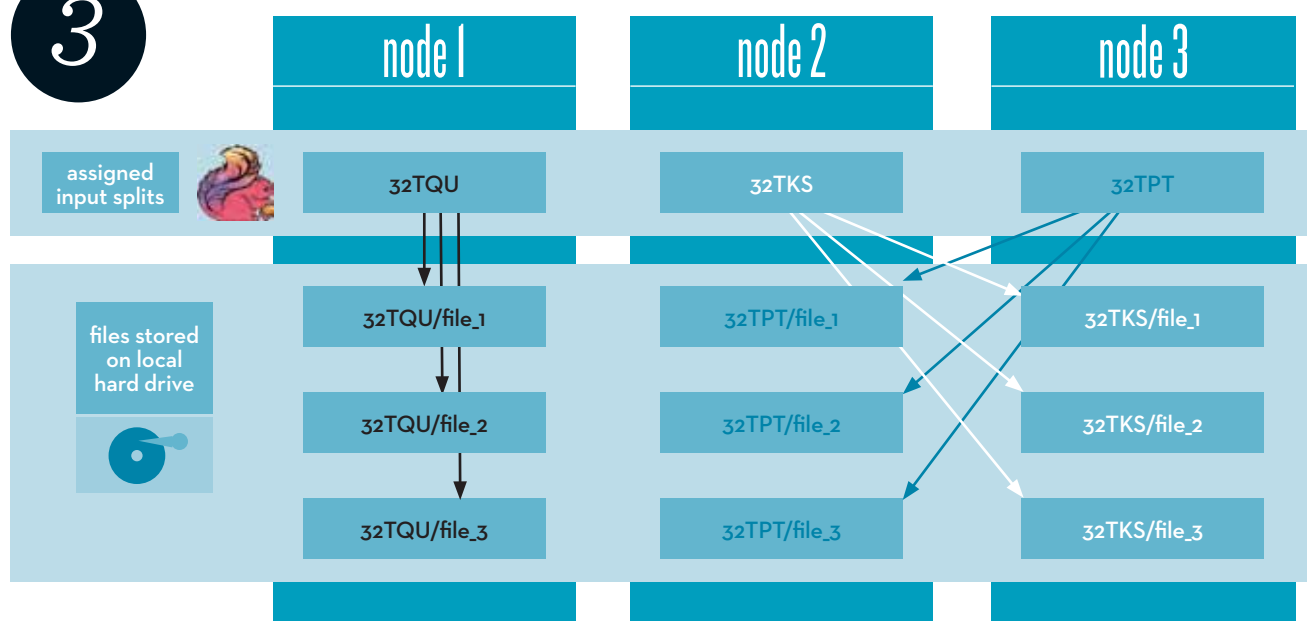
POSIX VS. HDFS INTERFACE

The HDFS interface is a file-system interface especially designed for large-scale distributed data processing and compatible with most MapReduce implementations. Our file system XtreamFS recently received support for this interface. Unfortunately, the libraries to work with satellite data used in GeoMultiSens require byte-level access to files, which is only supported by the POSIX interfaces and not the HDFS interface. Hence, our solution has to take this into account.

- 2 Random layout shows file accesses as provided by a naive XtreamFS setup.
- 3 Good layout depicts the current state of development.

DATA PLACEMENT FOR BETTER PERFORMANCE

3



REALIZATION IN XTREEMFS


We have enhanced the XtremFS servers with a new data-placement mechanism based on path prefixes. This way, one can ensure that all files in a directory and its subdirectories are stored on the same node. For the client side, we have developed a placement-management tool that allows automatic load balancing for new

input splits and automatic rebalancing of existing data.

With the described grouping of data files, we were able to improve the overall job run time from 80 minutes to 72 minutes for around 600 GB of input data distributed over 16 nodes. But this was just the

first step of improvement. The next step is to schedule the Flink tasks near the node storing the corresponding data, which will further improve the speedup. The presented grouping of data was a necessary preparatory step to exploit the full performance potential.

SHARE THE WORK: THE NEED FOR DISTRIBUTED MEDICAL DATA ANALYSIS



It started only recently that hospitals use acquired medical patient data not only for directly treating their patients, for example, using it for diagnostics or for therapy monitoring. If a patient consents, all the data is also stored in an electronic central storage infrastructure, for exam-

ple, for research-oriented data-mining purposes. Most of this clinical information, such as demographic data (age, gender, or ethnicity), medication lists, or simple parameters (e.g. blood type), is the so-called structured information. Structured information is easy to col-

lect, to exchange, and to query because it is standardized, (typically) stored in a common database system, and thus computer understandable. However, although these pieces of information certainly help medical doctors during their daily routine, there is yet another



NEW DATA SOURCES DELIVER BIG DATA

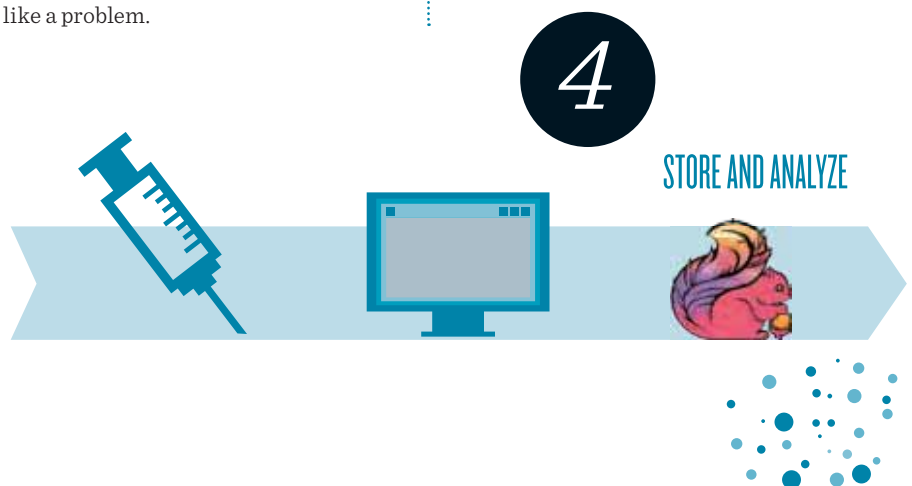
DNA sequencer and mass spectrometer are modern medical devices that can extract information about genes and proteins from a biological sample. A single drop of blood contains information about the full genetic setup (the building blocks) and all active proteins (the working horses in a cell) of a person. Detailed knowledge about all of a person's genes and proteins allows the status to be characterized in ways that have not been possible before. Was the person digesting, stressed, or carrying a disease? What are the odds of developing a major disease in the next years? However, to even start developing new medical tests based on these data, one has to first understand the raw information generated by these machines. One of the first problems is the sheer size: a raw data set can easily consume up to 1 TB (one terabyte) of disk space. This is the size of about 300,000 images or 40 days of video material. This seems like a lot of data, but somehow still manageable. However, for a meaningful scientific study, not one single data set is needed, but several hundred, better thousand participants. Now, storing (and watching) 300 million images or 40,000 days of video seems much more like a problem.

WHAT TO DO WITH ALL THESE DATA?

The point of collecting all the data is of course to later analyze them. This has at least two parts: First, the data need to be copied to the compute device that is assigned for the actual analysis. And second, the computationally heavy analysis has to be performed on some compute infrastructure. Both tasks are quite well understood and several frameworks are available to do the job. However, to really do this in a modern IT infrastructure environment takes quite some time: just to transfer 1 TB of patients' data takes about 2.5 hours on a standard gigabit-Ethernet (1,000 Mbit/s) intranet network. The standard analysis (remember: going from unstructured to structured) of that data takes another ten hours on a normal desktop PC. Now, multiplying that with the 1,000 study-participants' data sets gives some uncomfortable numbers: one would have to wait about 1.5 years until the results are available.

For more details about new analysis methods developed by ZIB, see the feature article "Bigger Data, Better Health."

type of information that is at least as important. This other type is often called unstructured information. The main feature of unstructured data is that it is not computer understandable. Simply put, this means that a human is needed for proper interpretation and – ultimately – translation into (computer-understandable) structured data. Examples for medical unstructured data are images (think of x-ray or MRT images), documents, or hand-written reports. Patient attributes (structured information), medical history reports, and x-ray images (unstructured information) certainly already provide a medical doctor with many pieces of the puzzle that characterize a patient. Nevertheless, modern medical devices can deliver even more.



4 Medical pipeline.

SCALABILITY

CALL THE CAVALRY

A simple yet very effective way of dealing with this kind of scaling problem is known: divide and conquer. This essentially means to split up the problem into smaller independent pieces. These small tasks can then be distributed over a whole army of compute nodes and solved in parallel. Obviously, this is needed twice: once for the storage problem and again for the analysis part. Divide-and-conquer strategies are well known in computer science. However, the hard part is to adjust the quite general approach to the actual storage and analysis problem at hand. However, once this is done, the idea needs to be implemented such that it can be executed on a distributed compute cluster. Researchers at ZIB together with partners from the Berlin Big Data Center have found very efficient ways to implement algorithms for the distributed analysis of genomics and proteomics data. This effort contributed to the development of the software frameworks Apache Flink and XtreamFS. The new approach now allows the distributed analysis of very large data sets in significantly less time than before, thus enabling the analysis of clinically relevant amounts of data with cohorts of thousands of patients.

SIMPLE BIG-DATA ANALYSIS MADE IN BERLIN



The development and ubiquitous employment of sensors of all kinds, collection of usage statistics, and refinement of scientific equipment – such as DNA sequences and mass spectrometers – are only a few recent developments that contribute to the ever-growing amount of data that is being collected for analysis. Managing the three Vs of big data – large volume, great variety, and high arriving velocity – call for specific application-domain expertise, knowledge of machine learning techniques, and scaling their implementations to systems of thousands of nodes. Funded by the BMBF, researchers from Beuth Hochschule, DFKI, Fritz Haber Institute, Technische Universität Berlin, and ZIB are joining efforts within the Berlin Big Data Center (BBDC) to develop “Technology X,”

which aims at opening up machine-learning-based, scalable data analytics for a large audience of data scientists through a concise declarative API. The intricacies of distributed systems development are hidden from the programmer by providing an execution engine that transparently scales the analysis from local development to cluster deployment. Queries are expressed declaratively using well-understood concepts of sets and transformations. With applications in information market places, materials science, and medical imaging, we enable data scientists to formulate scalable and effective programs using approaches known from SQL without being experts in distributed systems development or machine learning.

Y AT ALL LEVELS

FOLLOWING THE BYTE

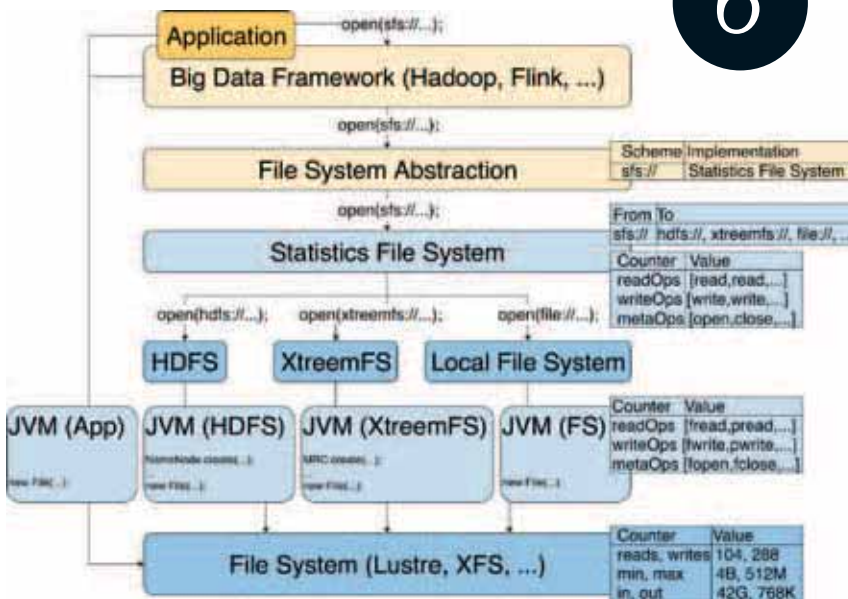
Most of these applications require terabytes of input data, which usually do not fit into a single machine. Therefore, distributed file systems (such as Hadoop's Distributed File System HDFS) are used to incorporate many machines and their disks into one logical file system. When an application reads from HDFS, it incurs read and write operations to potentially many disks in many machines. It is therefore vital to

understand and trace data-access at all abstraction levels: the application itself, the execution framework it is run on, the framework's distributed-file-system abstraction and the actual low-level file system. To this end, we have built the Statistics File System SFS, which can be plugged into any HDFS-compatible big-data-analytics engine at the highest and lowest levels to collect various I/O operation statistics.

STATISTICS FILE SYSTEM

At the highest level, SFS wraps around HDFS by logging every call and passing the request on to HDFS, measuring execution time, relevant parameters, and number of bytes accessed, among others. At the lowest level, SFS uses bytecode instrumentation to intercept all native I/O calls leaving the Java Virtual Machine JVM, logging similar statistics. This way, low-level disk access can be correlated with high-level HDFS interactions, and unintentional I/O can be identified. Cross-checking the SFS statistics with the underlying file system's statistics – as found in Lustre, XFS, and others – allows for precise determination of the application's share on total I/O on the disk.

6



5 Construction and spirit of the Berlin Big Data Center.

6 The Statistics File System wraps distributed file systems on different abstraction levels, gathering vital I/O usage and performance data.

	XFS Total	SFS Yarn	SFS HDFS	SFS Flink	SFS Total
Write	417	3	202	207	415
Read	545	0	101	332	534

(All numbers in gigabytes)

PEEK INTO PEAK IDENTIFICATION

We used SFS on an Apache Flink application analyzing 205 GB worth of blood-mass-spectrography data to obtain detailed I/O information. First, all data are staged into HDFS. Next, it is analyzed for specific peaks to identify diseases. These peaks are then written back to HDFS (just a few kilobytes). The application is run on YARN and HDFS on 16 nodes. Without SFS, the only available I/O statistics are those of the underlying file system’s counters, XFS in this case: 417 GB of writes to and 545 GB of reads from the file system (see table 7). However, we would expect 205 GB of writes during staging, 205 GB of reads during processing, and a few kilobytes of writes for the output – where does the I/O come from?

Because we instrument each JVM separately, with SFS, we are able to report each component’s share of the overall I/O. For HDFS, the numbers reported by SFS meet our expectations. Motivated by the unexpected I/O for Flink, we identified the cause to be a mass-spectrometry-data-processing library which does not support HDFS directly, and therefore downloads all data to local storage first before (repeatedly) reading it. The gap between the total XFS and total SFS numbers is largely accounted for by loading of libraries during JVM startup, which we have not explicitly reported in this case, as there is also unpredictable file-system caching involved.



ANALYSIS OF (SOME) SORT

Depending on the job, however, I/O can become a lot more involved. We have analyzed TeraSort as a common big-data benchmark, using Hadoop and Flink on YARN and HDFS on 16 nodes. First, 1 TB worth of 10-byte-key/90-byte-value pairs of random numbers is generated into HDFS using TeraGen. Then the numbers are globally sorted on the keys.

Ideally, we expect 2 TB of writes – 1 TB during input generation, 1 TB during output generation – and 1 TB of reads during sorting. For Hadoop, 1 TB of reads and writes each occurs during the shuffle phase. Again, our expectations are not met: XFS reports 4.1 TB of writes and 6.4 TB of reads for Hadoop. Using SFS, we are able to match these numbers to each component involved over time (see figure 8).

- yarn: jvm: read
- yarn: jvm: write
- hdfs: jvm: read
- hdfs: jvm: write
- map: sfs: read
- map: sfs: write
- map: jvm: read
- map: jvm: write
- reduce: jvm: read
- reduce: jvm: write

- 7 File I/O recorded by XFS and detailed I/O per component recorded by SFS.
- 8 I/O profiles for the TeraSort benchmark.

In the first 20 minutes, TeraGen generates the data using map tasks, and the writes in HDFS at JVM level correlate with the writes issued by these map tasks at SFS level, accumulating to 1 TB. Following that, the TeraSort map tasks read the data at SFS level, and HDFS serves these requests with 1 TB of reads at JVM level. The map tasks write all their data to local disk for the shuffle phase, causing 1 TB of writes in their JVM. Next, the reducers read this data from disk, and write the output to HDFS, incurring 1 TB of writes at SFS and JVM level.

So far, the analysis agrees with our expectations for TeraSort on Hadoop. To explain the gap between expectations and XFS counters, we point out three facts: 1) over the entire reduce phase, YARN

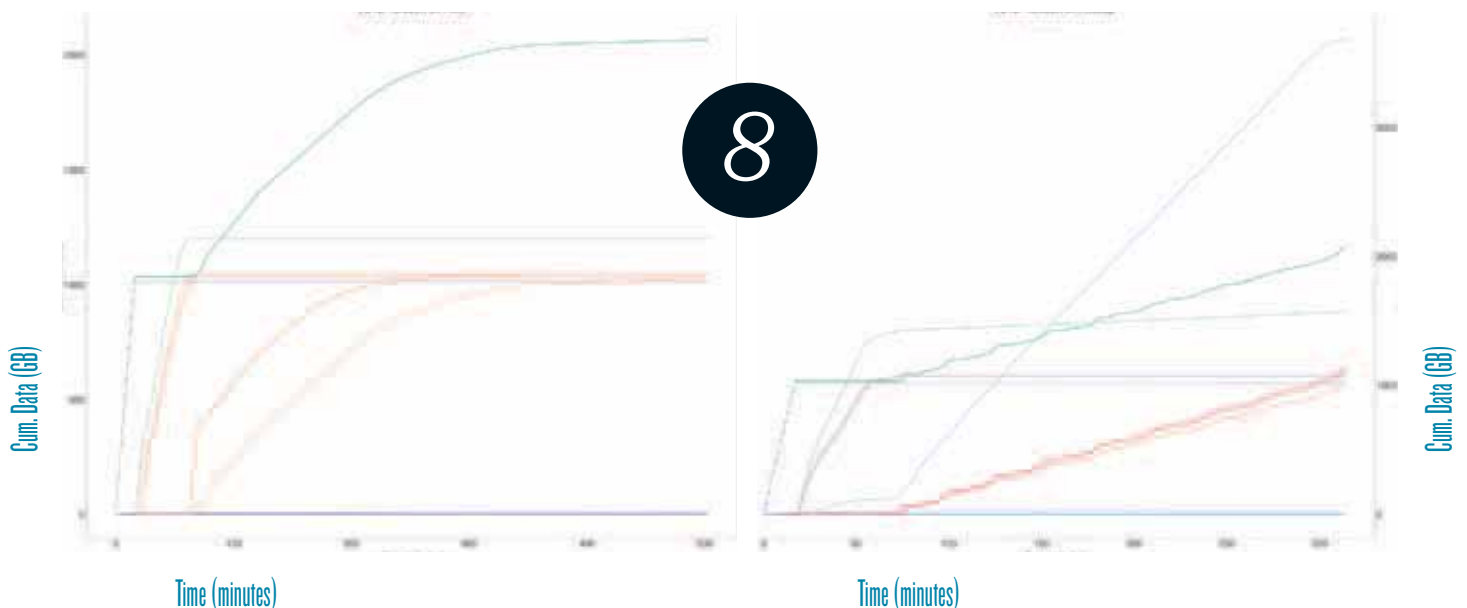
incurs more than 3.6 TB of reads; 2) the reducer JVMs write 1 TB of data without corresponding SFS requests; and 3) over the entire TeraSort run, the HDFS JVMs read 600 GB without corresponding SFS requests. For Flink, XFS reports 3 TB of writes and 3.2 TB of reads. The TeraGen step is the same as before, and all components except two behave as expected: again, the HDFS JVMs read 200 GB more data than are requested through SFS, and Flink adds 1 TB of reads and writes each (figure 8). Note that YARN incurs almost no overhead because instead of scheduling thousands of mappers and reducers, only few components need to be deployed. To close the gap between the SFS and XFS counters, 1 TB of reads is still unaccounted for, which we are currently investigating.

MIND THE DISK

Tracing file system calls from the highest to the lowest abstraction level allows the requests to be matched to I/O, which allows the detection of bottlenecks in disk access, especially in highly concurrent contexts. We have learned about unexpected I/O incurred by commonly used big-data-processing engines, only part of which can be explained by spilling because of memory exhaustion. Given that thousands of terabytes of data are being processed worldwide using these frameworks, surprisingly little is published about their underlying disk access. With our work, we are able to shed some light in the dark, and, using the entry-points we have at the highest and lowest levels, can make these frameworks more disk-aware.

I/O for TeraSort on Flink

I/O for TeraSort on Hadoop






SUPERCOMPUTING AT THE LIMIT

How Technology Innovations Impact Software Designs


Innovative technologies are changing the way of supercomputing. Many-core CPUs and accelerators are the answer to the ever-growing performance demands while staying within power budget constraints. High-bandwidth memory of processor devices extends the memory hierarchy, and high-capacity, persistent storage is seen near the processors. But how can we cope with the resulting challenges for software design?

1



NG

1 *Waver with CPUs manufactured with Intel's improved 14nm+ process.*




TECHNOLOGY INNOVATIONS CHALLENGES THE CODE DEVELOPERS ON SUPERCOMPUTERS

Moore's law from 1975, which proclaims the exponential growth of processor complexity and thus compute power, is still valid today. Consequently, hardware is outdated rather fast and the typical lifetime of a supercomputer is about five years. After this time, it is no longer economical to continue running the system, as the performance per energy invested is no longer competitive with what current hardware would deliver. Another reason is an increased component failure rate due to aging effects and thus higher maintenance cost.

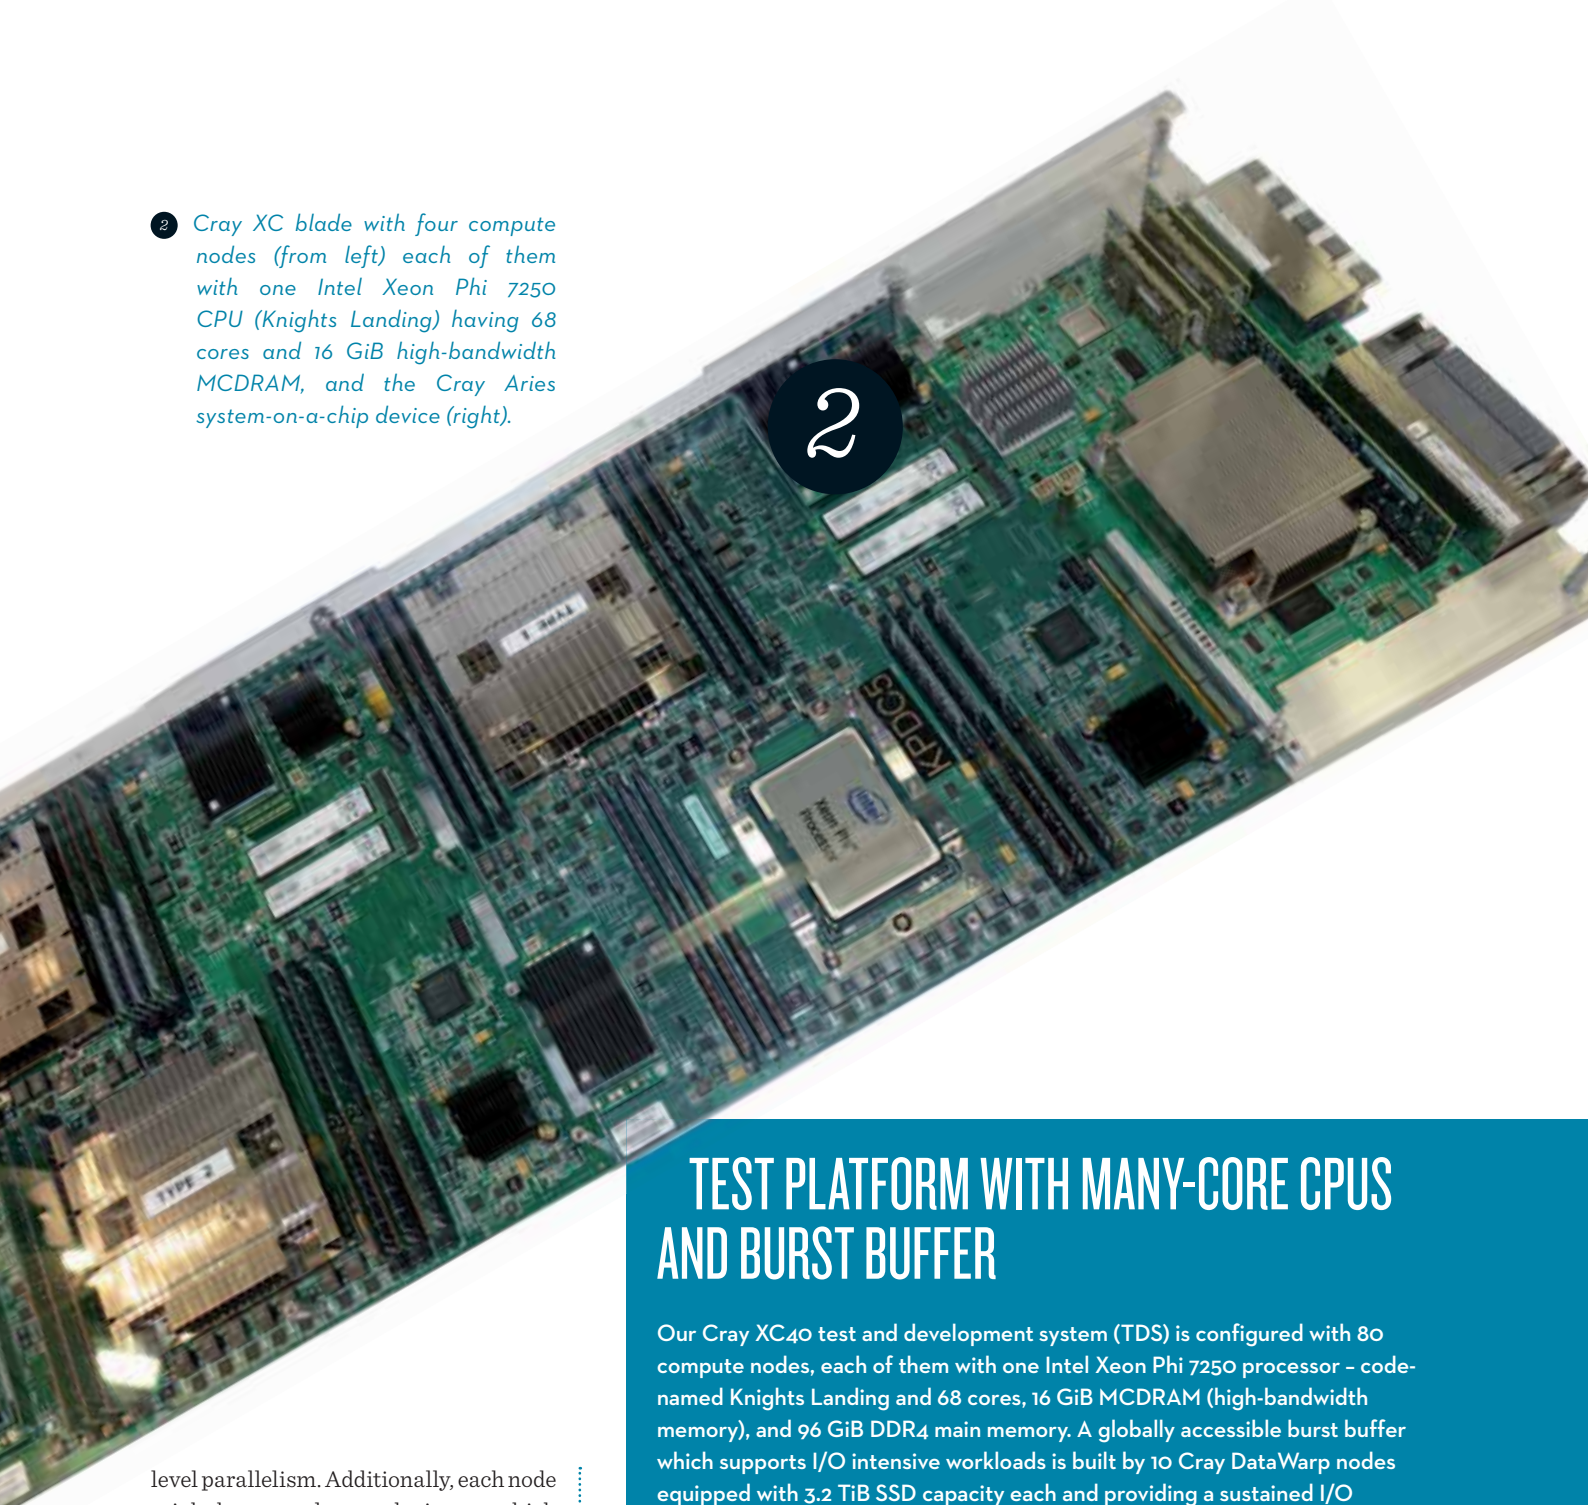
On the other hand, the application and library software running on these systems is rather long lived. Some codes date back to the 1970s and beyond. While software engineering in general can be fast paced, the scientific computing community is rather conservative in adapting the merits of modern programming techniques. This is a consequence of scientific diligence, but also of code being written mostly by noncomputer scientists whose focus is on computational results rather than modern code. Many codes are still written in FORTRAN, a programming language whose development started in the mid-1950s that is not found outside the domain of scientific computing anymore. This builds on a huge amount of

legacy code still in use that needs to be modernized for every new hardware architecture.

Modernization in this context means to adapt the code such that it uses the new hardware efficiently. The earliest computers would sequentially execute a stream of instructions provided by the programmer on a single processor core. Increasing the frequency at which these processors work would increase the application performance. Today's supercomputers are networks of workstation-like computers (nodes) that solve problems in parallel. Each node has one or more processors that has tens of compute cores working in parallel, each of which has multiple forms of instruction



2 Cray XC blade with four compute nodes (from left) each of them with one Intel Xeon Phi 7250 CPU (Knights Landing) having 68 cores and 16 GiB high-bandwidth MCDRAM, and the Cray Aries system-on-a-chip device (right).



level parallelism. Additionally, each node might have accelerator devices to which parts of the computation can be outsourced for faster, more energy-efficient processing – in parallel, of course. For each of the components, there exist multiple vendors, architectures, programming models, and software tools. Writing a correct scientific simulation program for a given mathematical model, that keeps all these resources busy in parallel, while handling communication is a hard problem. Transforming an existing code, that has been developed for decades, and uses a bunch of software libraries with an even longer history, to a new supercomputer system can be even harder.

TEST PLATFORM WITH MANY-CORE CPUS AND BURST BUFFER

Our Cray XC40 test and development system (TDS) is configured with 80 compute nodes, each of them with one Intel Xeon Phi 7250 processor – code-named Knights Landing and 68 cores, 16 GiB MCDRAM (high-bandwidth memory), and 96 GiB DDR4 main memory. A globally accessible burst buffer which supports I/O intensive workloads is built by 10 Cray DataWarp nodes equipped with 3.2 TiB SSD capacity each and providing a sustained I/O bandwidth of 3 GiB/s locally. All nodes are connected through the Cray Aries interconnect. In summary, the Cray TDS features the following total performance characteristics:

- ▶ Total theoretical peak performance: 244 TFLOPS.
- ▶ Total DRAM memory capacity: 7.6 TiB DDR4.
- ▶ Total high-bandwidth memory capacity: 1.2 TiB MCDRAM.
- ▶ Total burst-buffer capacity: 32 TiB SSD.

CODE MODERNIZATION FOR MANY CORES

Code modernization is one of the main tasks for a provider of supercomputing resources like ZIB. It is a collaborative effort that should be addressed hand in hand with the code developers and industry partners. It starts long before the call for bids for a new system procurement, as the codes should ideally be ready when a new production system goes online. At ZIB, one such effort is the Cray XC40 test-and-development system (TDS) that features Intel's newest Xeon Phi processors (Knights Landing) and Cray's recent DataWarp technology. It serves as a platform for code-modernization efforts within the HLRN and other European compute center through the Intel Parallel Compute Center collaboration network.





3 Intel Xeon Phi processor of the 7200 series (Knight Landing) as installed in our Cray XC40 test-and-development system.

MODERNIZING A LEGACY CODE:

THE PALM SIMULATION PROGRAM



4 Simulation of a densely built-up artificial island off the coast of Macau.



5 Simulation of dust devil with PALM.

THE PALM CODE


PALM – short for “A Parallelized Large-Eddy Simulation Model for Atmospheric and Oceanic Flows” – is one of the main production codes on the HRLN-III, and thus a priority target for modernization. It computes oceanic and atmospheric flows, and has a variety of applications in research. They range from studying natural phenomena such as dust devils or cloud physics, over technical applications such as wind energy or the interaction between aircraft and the atmosphere, to modeling and simulating air flows within whole cities. For instance, figure 4 shows the simulation of a densely built-up artificial island off the coast of Macau. The results allow studying the effects of such a mega-project on the air flows within the existing city in the early planning phase.

The PALM code has been being developed by the Siegfried Raasch’s research group at the Institute for Meteorology and Climate at Leibniz Universität Hannover since 1997. In that year, the fastest computer in the world provided a computational power of around one trillion floating-point operations per second (1 TFLOPS). The current HRLN-III Konrad system at ZIB provides roughly a thousandfold of that, or one quadril-

lion operations per second (1 PFLOPS). PALM is written in Fortran using the 95 and 2003 language standard, with around 140 thousand lines of code, structured in 79 modules and 171 source files. Parallelization and communication build on the MPI and OpenMP standards. It has proven to be highly scalable, for up to 43,200 cores.

In the following, we will describe the process and challenges of porting this code to the Intel Many Integrated Cores (MIC) architecture. This architecture is the foundation of the Intel Xeon Phi processor that powers recent supercomputer installations in the USA, Asia, and Europe, respectively. At ZIB, we are working on our Cray XC40 test-and-development system (TDS) with 80 Intel Xeon Phi (KNL) nodes.

A three-day hackathon meeting with members of the PALM development team and the Algorithms for Innovative Architectures Group of ZIB’s Supercomputing department are getting the process of code modernization started. The schema on top of the next page illustrates the main steps performed.



GETTING OPERATIONAL

The first step is to get operational on the new system. This means the code has to be successfully compiled and executed, not yet looking at performance numbers. Since the Intel Xeon Phi many-core chip has a very similar programming model compared with the existing multi-core Xeon processors, the same tool chain can be used with different settings. This renders building an executable program from the source code a relatively easy task. However, executing the program for the first time, and validating the results, revealed some bugs in the code that had not surface on any other architecture before. This is a common observation, as complex software cannot be formally validated to prove its correctness in a mathematical sense. It can only be tested within reasonable effort to minimize code defects causing unexpected behavior at runtime (i.e. bugs). Fixing these problems almost took the entire first day of the hackathon, at the end of which, we had a running program that produced the expected results for a simple test case.

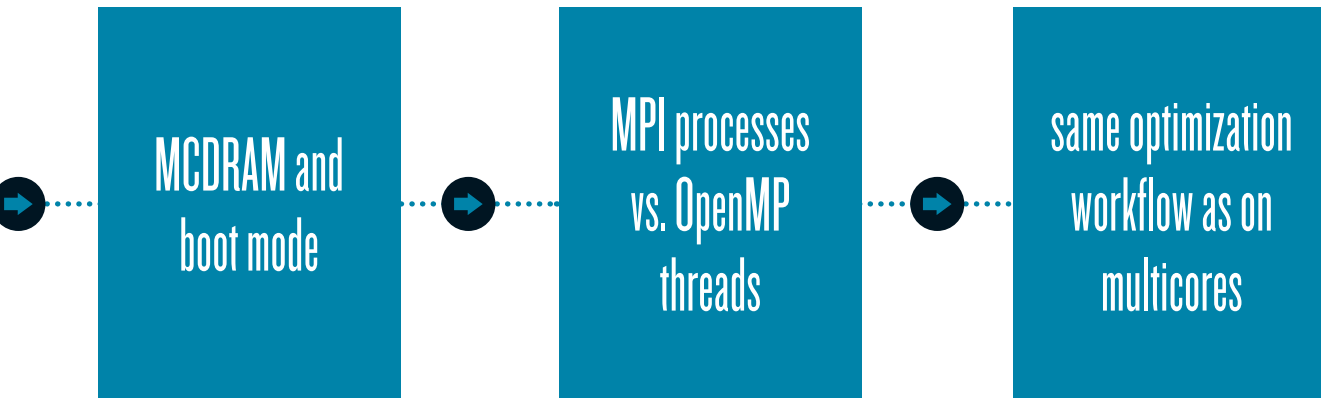
DEFINING THE BENCHMARKS

The next step in the process is to define a whole series of benchmarks covering different problems and problem sizes that represent typical use cases of PALM. These benchmarks can then be run to determine optimization success and compare the effects of different changes on the program's performance. At this stage, it is important to not focus just on a single test case to avoid overoptimizing for just a single execution path through the program's modules that might not represent the typical production workload. We selected three differently sized benchmarks using 4, 8 and 16 compute nodes of the Cray XC40 TDS.

SETTING REALISTIC EXPECTATIONS

Before starting to look at the performance impact of different system settings and code changes, it is beneficial to define a baseline to compare with later, and to build some expectations of what can be achieved. As a baseline, we measured the runtime of the defined benchmarks on the HLRN-III production system. A maximum speedup over that on the new system can be estimated by looking at the theoretical performance of both, the production, and the test-and-development system.

One compute node of the HLRN-III system (phase two), with its two 12-core Xeon processors (E5-2680v3) provides 960 GFLOP of computational performance, and 136 GiB/s of memory bandwidth between the processor and main memory. A node on the TDS has a single 68-core Xeon Phi processor (Intel Xeon Phi 7250) that has a peak of



MULTIPLE OPTIONS FOR ACCESSING THE HIGH-BANDWIDTH MEMORY ON INTEL KNL

2611 GFLOPS in compute performance, and 115 GiB/s of bandwidth to the main memory. To balance this huge compute power with the rather low main-memory bandwidth, it has an additional 16 GiB of MCDRAM (Multi-Channel DRAM), which is a high-bandwidth memory located on the processor package showing a sustained bandwidth of 490 GiB/s.

Based on these numbers, an application whose performance is limited by the computational throughput (compute bound), can run up 2.7 times faster than on the production system, while an application limited by data transfers (memory bound) can be sped up by a maximum of 3.6 times – given the newly introduced MCDRAM is used. In practice, these upper bounds are not likely to be reached, but they provide a means to assess what ever is measured later.

To support a wide range of applications, the processors of the Intel Xeon Phi 7200 series can be booted in a series of modes, which affect how the main memory, MCDRAM, and caches are configured and exposed to the software. The most important decision here is how to configure the aforementioned MCDRAM.

There are three options:

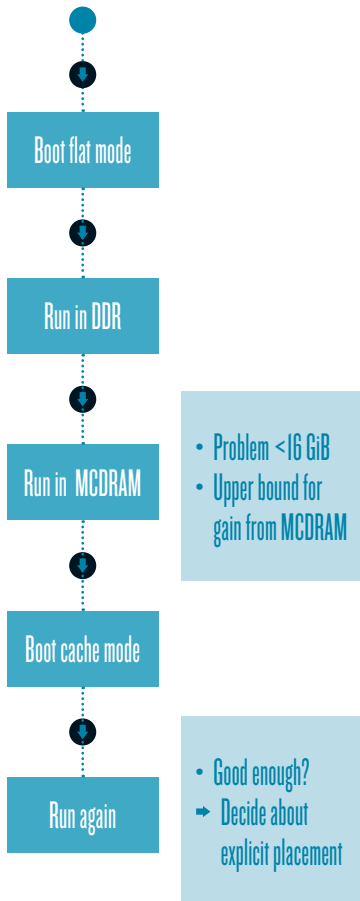
- It can be used as a last-level cache (cache mode), meaning it will not be exposed to the software at all, and all data in the main memory go through it. Thus, subsequent accesses to the same data are sped up by the higher memory bandwidth of the MCDRAM over the main memory.
- Another mode is to expose it as a separate region of memory that has to be explicitly used (flat mode).
- The third option is a hybrid mode, configuring part of the memory as a cache, and the other part for explicit use.

The cache mode has the advantage of not requiring any further action to make use of the MCDRAM. However, the cache protocol causes overhead. Since the MCDRAM is much smaller than the main memory it caches, that is, old data are constantly moved out to make room for new data, and data that are highly reused might get evicted for data that are only used once and thus does not benefit from the caching.

The explicit-usage model allows placing data that are known to be frequently accessed into the limited MCDRAM. This, of course, requires carefully analyzing the application and adapting the code to allocate certain parts of the memory in MCDRAM.

6

THE STRATEGY FOR DETERMINING MCDRAM PERFORMANCE IMPACT



6 The strategy for determining MCDRAM performance impact.

First, the flat mode is used. Running the benchmarks in this mode provides timings for only using the slow DRAM main memory. In a second step, still in flat mode, the application can be run using a special tool (numactl) to make it use MCDRAM only without changing the code. This requires problem sizes that entirely fit into the MCDRAM. Now we have two timings completely run in the main memory (DRAM), and completely run in MCDRAM. The ratio provides a maximum speedup due to MCDRAM usage for that application. Now the system is booted up in cache mode, and the benchmarks are run again. The ratio

between these numbers and the DRAM-only numbers cause the speedup due to using the MCDRAM as a cache – also for larger problem sizes that fit into the MCDRAM. If the difference between the MCDRAM-only and cache-mode speedups are small, the cache mode is a viable option.

If not, the code needs to be adapted for explicitly using the MCDRAM.

Figure 7 shows the data for PALM. The gain due to MCDRAM ranges from 25% to 41%, while the cache mode is less than 3% slower than the MCDRAM only. So, the cache mode works well for PALM.

BALANCING PROCESSES AND THREADS IN HYBRID WORKLOADS

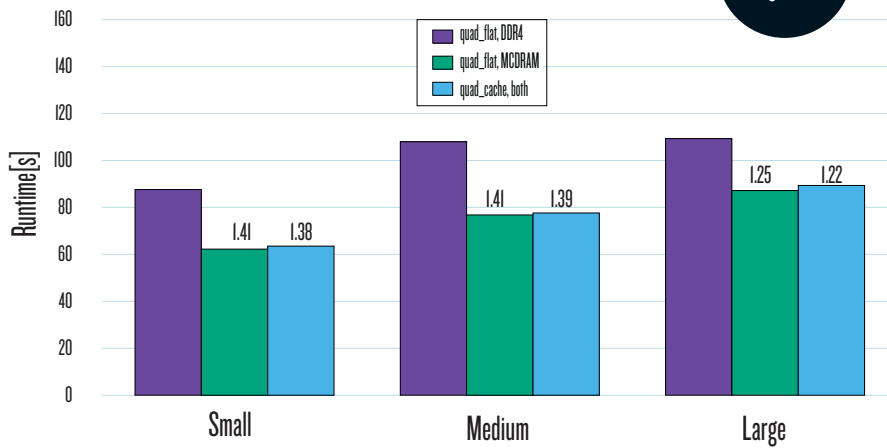
Since PALM is a hybrid code, using OpenMP for thread parallelism and MPI for communication, the “right” number of threads and processes needs to be figured out. A process is one instance of the program being run in an executable fashion. It has its own memory space and can use multiple parallel threads of execution to utilize multiple processor cores. MPI enables communication between such processes, either over the network between compute nodes or between processes running on the same node by copying memory between them. OpenMP handles the threads within each process, which share the same memory and thus do not need to communicate data between each other, but instead only synchronize concurrent accesses to the same memory region.

While naturally, there would be one MPI process per compute node that uses as many threads as there are processor cores available, it is often beneficial to have multiple MPI processes on each compute node that use less OpenMP threads each. One reason is that the fast, internal network of a supercomputer is easier to saturate by having multiple MPI processes feeding it with data. Another reason is the historic development of MPI being there before OpenMP, such that MPI parallelism is often better implemented than OpenMP thread parallelism – even though it causes avoidable memory copies.

Assuming we want to use 64 cores per Xeon Phi compute node, we could start a single MPI process per node using 64 OpenMP threads each, or 64 MPI processes per node using a single OpenMP thread each – or anything in between.



MCDRAM Usage Comparison



7 Impact of different MCDRAM usage modes on runtimes and speedups for the three PALM benchmarks “small,” “medium,” and “large.” Shown are data for main memory only (DDR4, purple bars) as baseline vs. MCDRAM only (green bars), and cache mode (blue bars) usage.

THE INTEL XEON PHI USER’S GROUP (IXPUG)

The Intel Xeon Phi User’s Group (www.ixpug.org) is an independent nonprofit organization whose mission is to provide a forum for the free exchange of information that enhances the usability and efficiency of scientific and technical applications running on large High Performance Computing (HPC) systems using the Intel Xeon Phi processor. IXPUG is administered by representatives of member sites that operate large Xeon Phi-based HPC systems. Thomas Steinke from ZIB is cofounder of IXPUG and served as its vice president from 2014 to 2017. He now helps as senior advisor on the Steering Group of IXPUG.

Which setting works best highly depends on the application at hand, and is hard to predict. The easiest way to figure it out is to run a parametric study for a series of settings to see what works best.

Figure 8 shows the results for PALM. While the fastest setting is a different one for each benchmark, there is one setting that performs reasonably well for all of them (i.e. 16 MPI processes [often called ranks] per node with four OpenMP threads each). Comparing the fastest with the slowest tested configuration reveals a factor of 2.3 between the two, underlining how crucial this step is on the way to optimizing for performance on that architecture.



8 Parametric study for the impact of the number of MPI processes (ranks) vs. OpenMP threads per compute node. Shown are runtime results of the three PALM benchmark cases as a heat map.

MPI Ranks vs. OpenMP Threads per KNL Node

Ranks	64	32	16	8	4	2	1
Threads	1	2	4	8	16	32	64
Small	Orange	Orange	Blue with X	Orange	Red	Red	Red
Medium	Orange	Orange	Blue	Blue with X	Red	Red	Red
Large	Blue	Blue with X	Blue	Blue	Orange	Red	Red

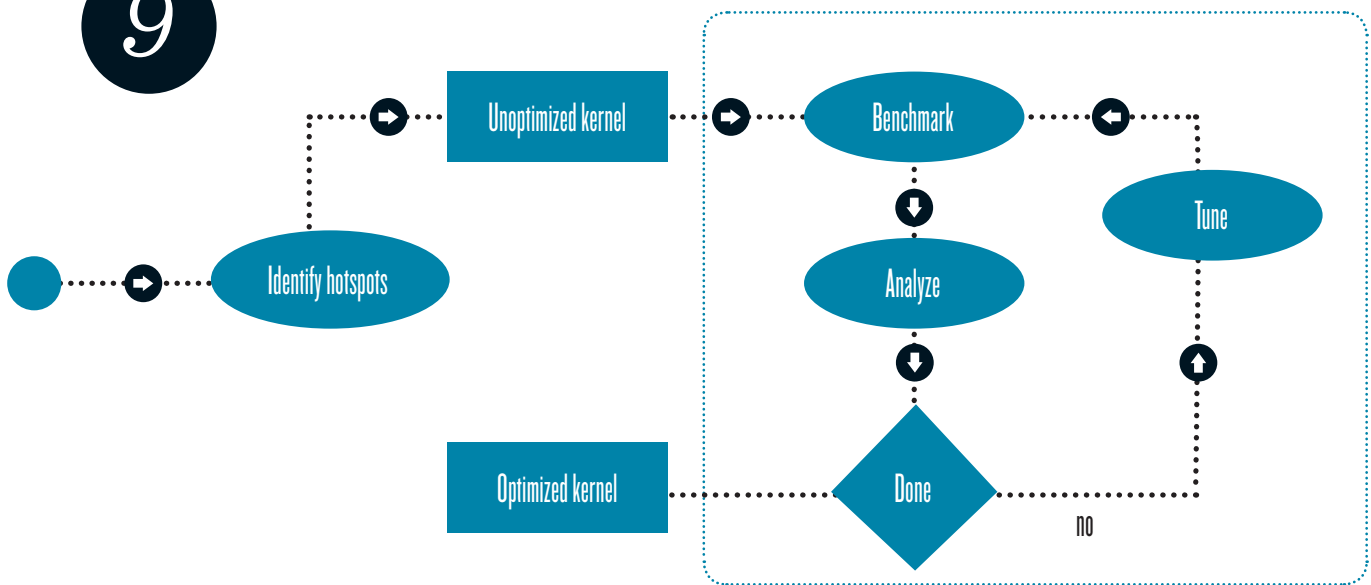


THE CODE OPTIMIZATION WORKFLOW

Now that the basics are determined (i.e. the MCDRAM-usage mode and MPI/OpenMP configuration), the actual work with the code can start. The corresponding workflow is shown in figure 9. It is basically the same as on multi-core architectures such as the Xeon processors of the HLRN-III production system.

At first, for the whole application, a runtime profile is generated using a tool like Intel VTune Amplifier. This profile reveals which parts, or kernels, of the code consume the most runtime, and thus are the best target for optimization efforts. This kernel is then benchmarked, analyzed, and modified in a loop until the performance matches expectations. Which changes to perform are guided by experience, compiler optimization reports, and special analysis tools.

9



9 General code-optimization workflow.

On the Xeon Phi, the most important optimization is to make sure the code is vectorized, as most of the computational power comes from so-called SIMD (single-instruction multiple data) units. They can perform arithmetic operations, like multiplying two inputs into an output, on vectors of data of given length in one step instead of using single scalar data arguments requiring multiple steps. The Intel Xeon Phi works with 512-bit vectors, which can fit eight double precision floating-point numbers. On the hardware side, this means producing eight results at once without introducing additional transistors for the control logic, as compared with eight scalar multiplication units. This design results in more computational performance per

chip area and per energy consumed. Not using the SIMD units costs a potential factor eight in application performance.

For PALM, the compiler report revealed the currently implemented way of numeric error handling completely prevents vectorization and requires adapting the application design. The currently used exception handling needs to be disabled, and instead regular, manual checks on the data need to be introduced, for example, when checkpointing the application state. With that change, the compiler still did not generate significantly vectorized code. Further analysis revealed that the current memory layout needs to be adapted for efficient vectorization too. This is necessary to avoid gathering vector elements from eight

different memory locations and instead to load them as on contiguous piece of memory. However, this is a bigger change and part of the ongoing efforts with that code.

From the different things we have tested so far, introducing the CONTIGUOUS keyword of the Fortran 2008 language to tell the compiler explicitly that an array is contiguously allocated in memory has had the biggest impact. It prevents the compiler from making the pessimistic assumption of noncontiguous memory. Beside the Intel Fortran compiler, we also have tested the Cray Fortran compiler which produced faster machine code due to better performing optimization techniques for this application.

10 Projected speedups, that is, without one-time initialization cost, for PALM on Cray XC40 with Intel Xeon Phi 7250 (KNL) over the original code on the HLRN-III, a Cray XC40 with Intel Haswell nodes. For the Intel Xeon Haswell CPU, the code is generated by the Intel compiler. For the Intel Knights Landing CPU, the Intel and Cray compiler were used.

MANY-CORE RÉSUMÉ

Figure 10 shows the results obtained after the three-day hackathon. The speedup over the production system is up to 1.45.

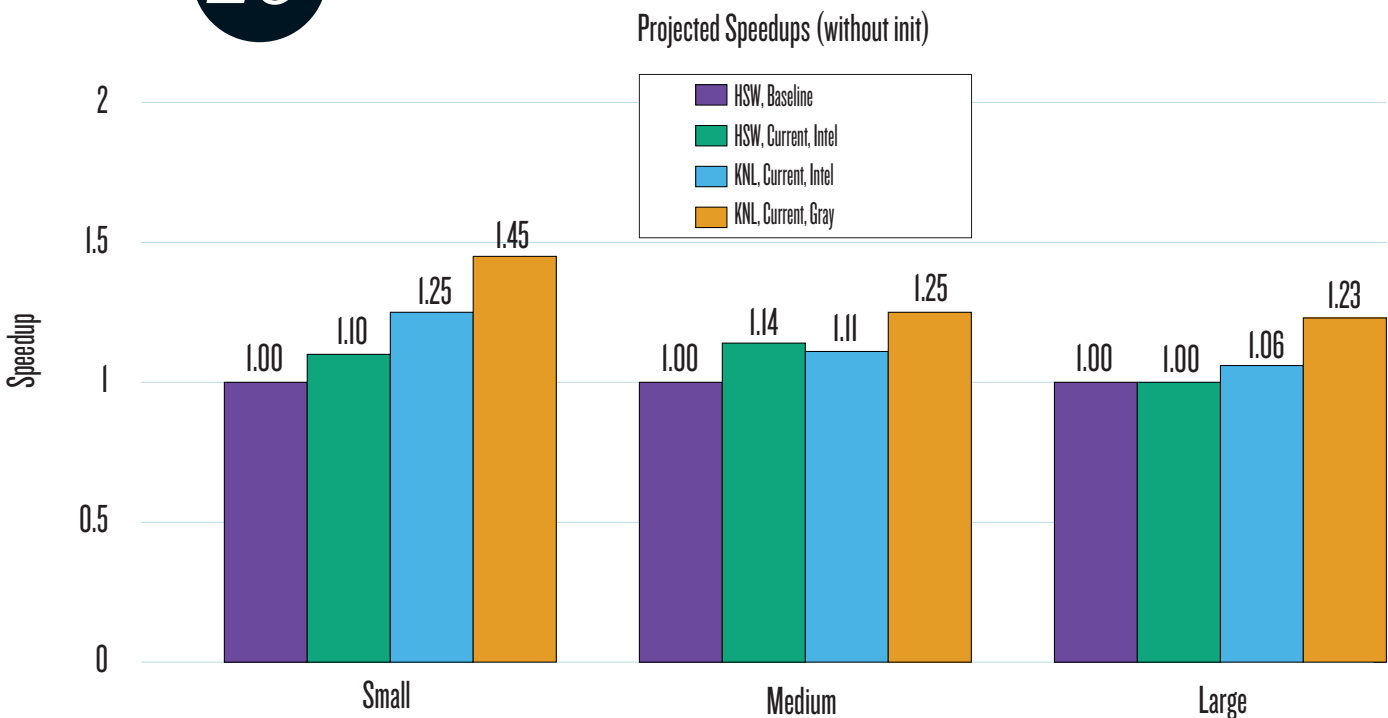
Compared to accelerator architectures such as GPUs (Graphics Processing Units) or the accelerator version of the Xeon Phi, porting code from multi-core CPUs like the Intel Xeon to the current generation many-core Xeon Phi is relatively easy. Code does not need to be offloaded from the host processor to the accelerator, and no data have to be transferred between their memories. Also, the tool chain remains the same.

However, getting performance out of the Xeon Phi can be challenging due to their new architecture features: the increased level of parallelism and slower single-core performance, the MCDRAM



as a new layer in the memory hierarchy, and the mandatory use of the SIMD units. Fortunately, optimizations for the Xeon Phi are usually beneficial on Xeon as well because the latter has SIMD units (currently 256-bit vectors) too. The first results for the PALM code are very promising, but getting the code to better use the processor's SIMD units is a larger effort still in progress.

Pushing the limits of scientific high-performance computing means not just buying faster computers every few years, but, most of all, continuously modernizing the applications to make best use of them.

10



STORING DATA BECOMES MULTIFACETED


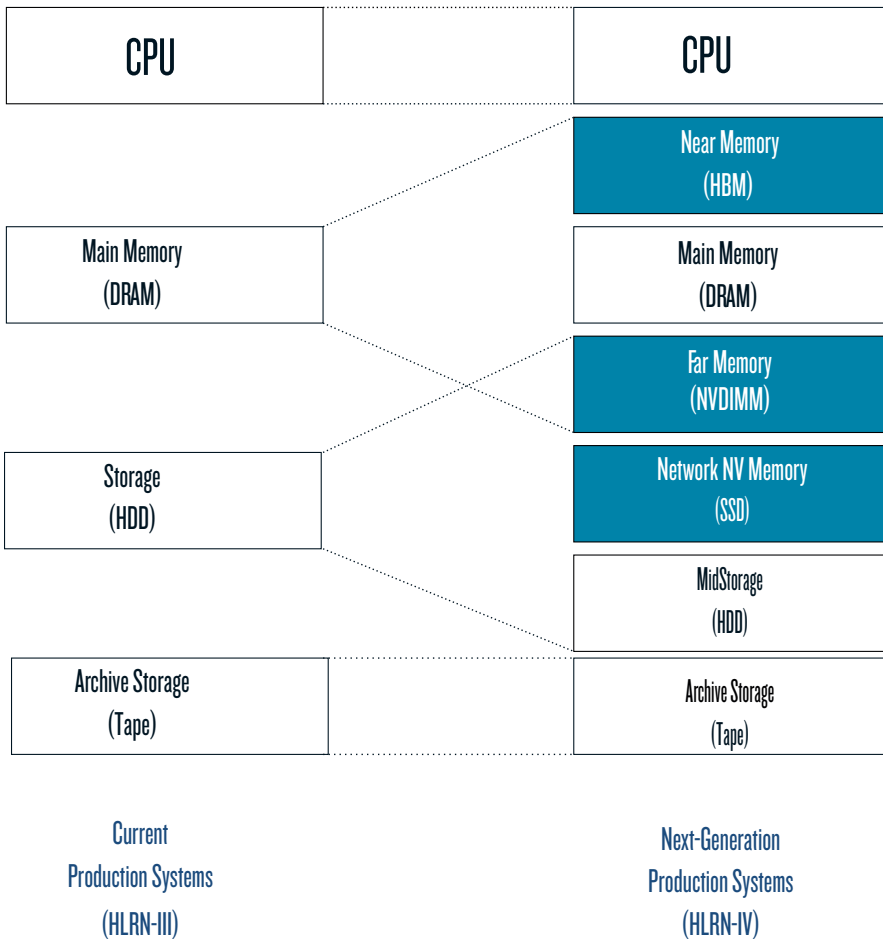


11 Cray DataWarp node **a** with two SSD slots **b** to integrate a burst buffer into a Cray XC system. **c** State-of-the-art SSD storage solutions for new data tiers providing low-latency, high throughput and ultraendurance in data centers. Shown is an Intel Optane SSD DC P4800X series.

Beginning in the 1950s, the memory and storage hierarchy in computing developed into a hierarchy of three tiers: main memory for fast access to temporary data; secondary memory for storing data persistently on disks; and tape-based archive storage for long-living data archives. This traditional three-level storage hierarchy is now going to change into a more complex one. Solid-state drives (SSD) are increasingly integrated in any computer system to replace and enhance the traditional disk-based storage solution. Recently, in the new generations of many-core CPUs and GPUs, high-bandwidth memory (HBM) is stacked onto the CPU or GPU device substrate and thereby integrated into the device package. With that, HBM deepens the memory hierarchy as it is functionally placed between the on-chip cache and the external DRAM (main) memory. HBM is not faster, but is wider and tackles today's increasingly critical "memory wall" – the restricted compute performance due to the limited bandwidth to data residing in the main memory.



11



The next technology innovation in the memory segment is visible on the horizon: nonvolatile dual in-line memory modules (NVDIMM). It is expected that NVDIMM and HBM become widely available in the next years and complement existing technologies such as DRAM and SSD. A possible scenario for, but not limited to, supercomputer configurations is a memory and storage hierarchy consisting of six tiers as shown in figure 12.

Complex HPC workflows are orchestrated as multiple-data-processing and data-management steps. Each access to high-volume or nontrivially structured data must be adapted to take advantage of these new storage technologies and their different characteristics, such as bandwidth, access time, and capacity. At ZIB, we are building the “I/O Test Bed” platform helping us to explore future technologies for data management in supercomputer systems.

12 *The memory and storage hierarchy of the current HLRN-III production system (left), consists of the three tiers of main memory in DRAM: persistent storage using hard-disk drives (HDD), and an archive using tapes. For the next-generation HPC system HLRN-IV (left), we expect high-bandwidth memory (HBM) and nonvolatile DIMMs (NVDIMM) extending the main memory by near memory and far memory, respectively. NVDIMMs and SSDs as network-attached nonvolatile memory can extend disk-based file-system functionality.*

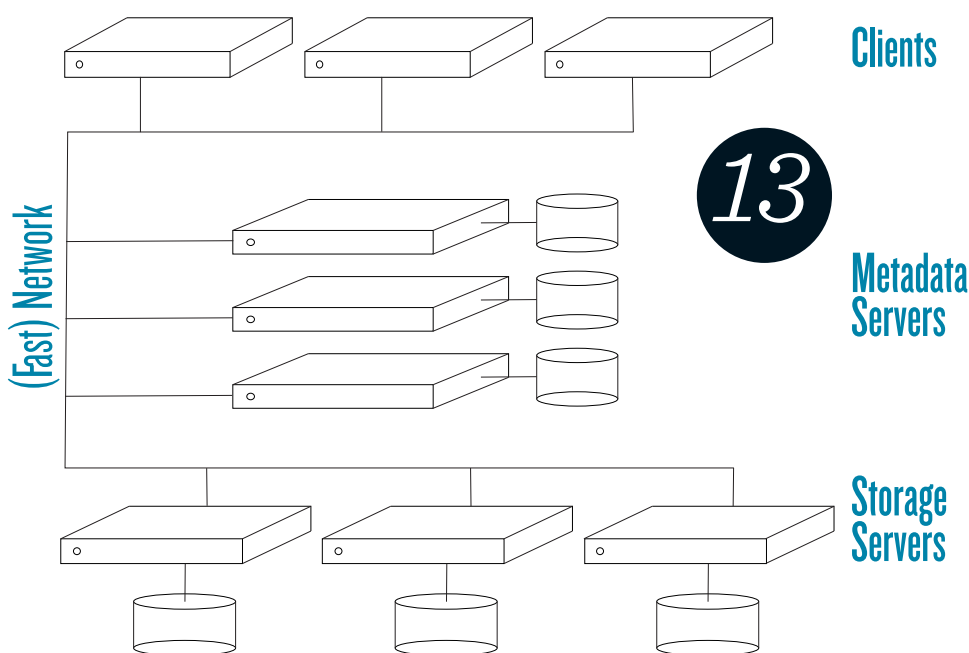
ARCHITECTURE OF PARALLEL HPC FILE SYSTEMS

Parallel file systems are designed for storing data persistently to meet the demands of supercomputing workloads: access with high bandwidth and reliability. To circumvent the limited I/O performance of a single storage server, the principle of parallelism is also applied to the storage system in today's super-computer installation. Multiple storage servers are combined to offer parallel data paths in a network such that the load is distributed over many storage servers (see figure below). Many clients, for

example, thousands of compute nodes, communicate with multiple servers simultaneously, enabling parallel access to data by the workloads running concurrently on the supercomputer. Ideally, the I/O performance and the file-system capacity scales with the number of installed storage servers.

Scalability is achieved by separating file-system metadata such as file names and directory information from the content of a file. So-called metadata servers

manage the metadata of the file system, whereas the storage servers save and serve stripes of user files contents. The file system is mounted on all compute nodes as a network file system. Current parallel file-system installations have thousands of clients, a capacity in the petabyte range, and deliver an I/O streaming throughput in the range from hundreds of gigabytes per second up to multiple terabytes per second. Commonly used parallel file systems are Lustre and BeeGFS.



13 The basic structure of a parallel file system: clients (e.g. compute nodes), metadata and storage servers are connected through a high-bandwidth network to provide multiple data paths to data in a file system.



EXPLORING THE CAPABILITIES OF PARALLEL FILE SYSTEMS

Operating a large supercomputer production system like the current HLRN-III requires a detailed understanding of the configuration of the installed storage system. Parallel file systems support a variety of installation requirements. This desired flexibility is addressed by many parameters which control the file-system configuration. Primarily, parallel file systems can be tuned for best performance to minimize unnecessary slowdown of the production runs with their often complicated data flows.

Data on a parallel file system is globally accessible from all compute nodes, and usually only one or a few instances of such a file systems are part of a supercomputer installation. Therefore, a high availability is needed, that is, data on the parallel file system is expected to be accessible all the time. An unreliable storage system would become the weak spot of any supercomputer system, rendering the whole system unusable during downtimes.

We explore the configuration space of state-of-the-art parallel file systems for high-performance computing by means of synthetic I/O benchmarks and real-world, I/O-intensive workloads from different scientific fields on our storage testbed.

OUR STORAGE TESTBED

First, our storage testbed acts as a test infrastructure for different configuration and tuning settings, optimizations, and file-system updates. New file-system features, such as multiple metadata servers, file system snapshots, access-control mechanisms, and resource-usage policies, are tested before their deployment on the production system.

Second, the testbed serves as a platform for the development of tools for the I/O profiling of workloads, and for file-system monitoring. Successfully tested, the tools can be subsequently rolled out on the larger production supercomputer system with hundreds of users.

METADATA PERFORMANCE IN COMPARISON

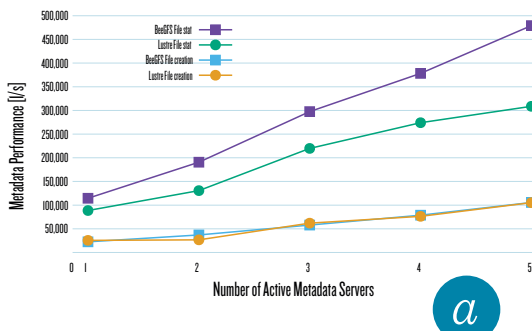
For system configurations with thousands of compute nodes, the metadata performance of a parallel file system becomes a bottleneck for trivial I/O operations, such as creating a new file or retrieving the metadata of existing files. Recent versions of parallel file systems support multiple metadata services to distribute the load. We measured the scaling of the metadata performance of the BeeGFS and Lustre file system as a function of the number of active metadata servers. For this benchmark, we used the widely known “mdtest” tool. As shown in figure [a](#), the retrieving of metadata (“file stat”) scales well with the number of metadata server for both BeeGFS and Lustre with advantages for the former. In contrast, the “file creation” performance scales weakly, and is similar for both file systems.

THE I/O TESTBED

The I/O testbed consists of five file servers each equipped with one Intel Xeon Haswell processor E5-1620v3, 64 GB RAM, a hardware RAID controller with two 120 GB SSDs and 10 SAS drives of 2 TiB capacity each for the parallel file system, and two FDR InfiniBand cards for fast network connections. The testbed also contains eight compute nodes with two Intel Xeon Haswell E5-2620v3 processors, 64 GB RAM, and one InfiniBand card. The operating system is CentOS 7. The testbed is integrated into the InfiniBand fabric of the Cray TDS and HLRN-III production system “Konrad” so that scaling experiments with a large number of nodes can be realized.

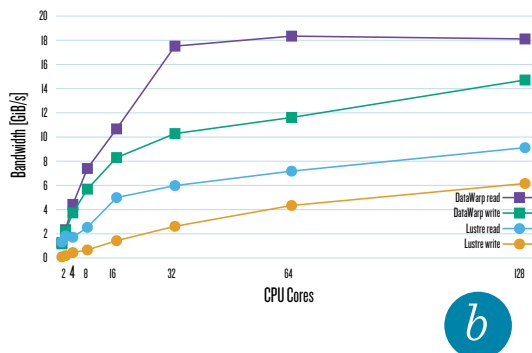
14

Scaling of Metadata Performance on I/O Test Bed




a

Synthetic IOR Benchmark Lustre versus Cray DataWarp



b

Application Benchmark TURBOMOLE Lustre versus Cray DataWarp



c

IMPROVING I/O PERFORMANCE BY BURST BUFFER

A burst buffer is an additional storage layer between the node memory and the disk-based parallel file system that accelerates I/O operations of the applications on the supercomputer system. Nowadays, implemented as nonvolatile storage using SSDs, it provides a significant performance boost for a broad range of file-access patterns – from small-sized block I/O up to streaming high-volume data sets.

We used the Cray DataWarp technology which implements a burst buffer as network-attached memory. In the Cray XC system, a pool of fast SSDs configured in dedicated I/O nodes are made accessible on each compute node via the Aries network as one additional SSD-based tier with a high aggregated bandwidth.

We compared the performance of the Cray DataWarp technology with our disk-based Lustre file system. We used a synthetic benchmark and a typical HPC workload.

First, we measured the scaling of the streaming bandwidth using a synthetic and widely known IOR benchmark. With 128 processor cores, the stream-

ing bandwidth for writing and reading data on Cray DataWarp is twice as much compared to the Lustre performance (see figure **b**). But can real-world I/O-intensive workloads benefit from a burst buffer?

To answer that question, we ran TURBOMOLE (reference 3), a well-known ab-initio quantum chemistry program package. We selected the TURBOMOLE implementation of the RI-MP2 method which is often used by scientists on the HLRN system. This method includes I/O-intensive operations on temporary data. Within our benchmark, almost 80 GiB were written to the Lustre file system and/or the burst buffer, respectively, and read back. The measured wall times clearly indicated the advantage of the burst-buffer technology for the I/O pattern in the workload as shown in figure **c**. For all test cases, the run time with scratch I/O on SSD-based burst buffers is significantly lower than on disk-based Lustre file systems. The performance gain is more pronounced for higher node counts, which is important for scalable workloads.

- 14** **a** Scaling of metadata performance for the “file stat” and “file creation” operation as function of the number of active metadata servers.
- b** Streaming performance for read and write access to the disk-based Lustre file system and Cray DataWarp burst buffer.
- c** Total runtimes of the RI-MP2 method as implemented in TURBOMOLE using up to 16 nodes by performing the scratch I/O on the Lustre file system or on the burst buffer (Cray DataWarp), respectively.

A FUTURE WITH OPEN POSSIBILITIES

Innovative data-processing technologies such as many-core CPUs and GPUs provide the capabilities to match demands for computational power and lower energy consumption. The key architectural features of CPUs are wider SIMD units – up to 512 bit – or wide SIMT (Single Instruction Multiple Threads) in GPUs. To exploit this performance potential, existing legacy codes need to be adapted in code-modernization efforts. Developing best-practice solutions in collaboration with a strong community of early adopters (e.g. the Intel Xeon Phi User's Group), are essential so that domain scientists can design state-of-the-art simulations codes.





DATA DIRECT
NETWORKS
DataDirect

DATA DIRECT
NETWORKS
DataDirect

DATA DIRECT
NETWORKS
DataDirect

ZIB PUBLICATIONS

PEER-REVIEWED (116)

Aimi Abass, Philipp Gutsche, Bjorn Maes, Carsten Rockstuhl, Emiliano R Martins (2016). Insights into directional scattering: from coupled dipoles to asymmetric dimer nanoantennas. *Opt. Express*, 24(17):19638-19650.

Carlo Barth, Jürgen Probst, Sven Herrmann, Martin Hammerschmidt, Christiane Becker (2016). Numerical characterization of symmetry properties for photonic crystals with hexagonal lattice. *Proc. SPIE*, 9885:988506.

Carlo Barth, Sven Burger, Christiane Becker (2016). Symmetry-Dependency of Anticrossing Phenomena in Slab-Type Photonic Crystals. *Opt. Express*, 24:10931.

Gernot Bauer, Nadezhda Gribova, Alexander Lange, Christian Holm, Joachim Gross (2016). Three-body effects in triplets of capped gold nanocrystals. *Molecular Physics*, 1-10.

Daniel Baum, Jürgen Titschack (2016). Cavity and Pore Segmentation in 3-D Images with Ambient Occlusion. In *EuroVis 2016 - Short Papers*.

Isabel Beckenbach, Ralf Borndörfer (2016). An Approximation Result for Matchings in Partitioned Hypergraphs. In *Operations Research Proceedings 2014*, 31-36.

Kaspar Bienefeld, Fred Zautke, Pooja Gupta (2016). A novel method for undisturbed long-term observation of the honey bee (*Apis mellifera*) behaviour - illustrated by the hygienic behaviour towards *Varroa infestation*. *Journal of Apicultural Research*, 54(5):541-547.

Marco Blanco, Ralf Borndörfer, Nam Dung Hoang, Anton Kaier, Adam Schienle, Thomas Schlechte, Swen Schlobach (2016). Solving

Time Dependent Shortest Path Problems on Airway Networks Using Super-Optimal Wind. In *16th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2016)*, volume 54 of *OpenAccess Series in Informatics (OASISs)*.

Ralf Borndörfer, Armin Fügenschuh, Torsten Klug, Thilo Schang, Thomas Schlechte, Hanno Schülldorf (2016). The Freight Train Routing Problem for Congested Railway Networks with Mixed Traffic. *Transportation Science*.

Ralf Borndörfer, Guillaume Sagnol, Thomas Schlechte, Elmar Swarat (2016). Optimal duty rostering for toll enforcement inspectors. *Annals of Operations Research*.

Ralf Borndörfer, Heide Hoppmann, Marika Karbstein (2016). Passenger routing for periodic timetable optimization. *Public Transport*.

Ralf Borndörfer, Heide Hoppmann, Marika Karbstein (2016). Separation of Cycle Inequalities for the Periodic Timetabling Problem. In *24th Annual European Symposium on Algorithms (ESA 2016)*, volume 57 of *Leibniz International Proceedings in Informatics (LIPIcs)*.

Ralf Borndörfer, Guillaume Sagnol, Stephan Schwartz (2016). An Extended Network Interdiction Problem for Optimal Toll Control. In *INOC 2015 - 7th International Network Optimization Conference*, volume 52 of *Electronic Notes in Discrete Mathematics*, pages 301-308.

Ralf Borndörfer, Marika Karbstein, Julika Mehrgahrtdt, Markus Reuther, Thomas Schlechte (2016). The Cycle Embedding Problem. In *Operations Research Proceedings 2014*, 465-472.

Ralf Borndörfer, Markus Reuther, Thomas Schlechte, Kerstin Waas, Steffen Weider (2016). Integrated Optimization of Rolling Stock Rotations for Intercity Railways. *Transportation Science*, 50(3):863-877.

Ralf Borndörfer, Sebastian Schenker, Martin Skutella, Timo Strunk (2016). PolySCIP. In G.-M. Greuel, Thorsten Koch, Peter Paule, Andrew Sommese, editors, *Mathematical Software - ICMS 2016*, 5th International Conference, Berlin, Germany, July 11-14, 2016, Proceedings, volume 9725 of *Lecture Notes in Computer Science*, pages 259-264.

Ralf Borndörfer, Heide Hoppmann, Marika Karbstein, Fabian Löbel (2016). The Modulo Network Simplex with Integrated Passenger Routing. In *Operations Research Proceedings 2016*. (accepted for publication on 2016-10-21)

Christopher Brandt, Christoph von Tycowicz, Klaus Hildebrandt (2016). Geometric Flows of Curves in Shape Space for Processing Motion of Deformable Objects. *Computer Graphics Forum*, 35(2).

Karl-Kiên Cao, Ambros Gleixner, Matthias Miltenberger (2016). Methoden zur Reduktion der Rechenzeit linearer Optimierungsmodelle in der Energiewirtschaft - Eine Performance-Analyse. In *EnInnov 2016: 14. Symposium Energieinnovation 2016*.

Tim Conrad, Xintian You (2016). Acfs: accurate circRNA identification and quantification from NGS data. *Nature Scientific Reports*, 6.

Natasa Djurdjevic Conrad, Marcus Weber, Christof Schütte (2016). Finding dominant structures of nonreversible Markov processes. *Multiscale Modeling and Simulation*, 14(4):1319-1340.

Marta Costa, James D. Manton, Aaron D. Ostrovsky, Steffen Prohaska, Gregory S. X. E. Jefferis (2016). NBLAST: Rapid, Sensitive Comparison of Neuronal Structure and Construction of Neuron Family Databases. *Neuron*, 91(2):293-311.

Fabio D'Andreagiovanni, Ambros Gleixner (2016). Towards an accurate solution of wireless network design problems. *Proceedings of the 4th International Symposium on Combinatorial Optimization (ISCO)*.

Fabio D'Andreagiovanni, Fabian Mett, Jonad Pulaj (2016). An (MI)LP-based Primal Heuristic for 3-Architecture Connected Facility Location in Urban Access Network Design. In *EvoApplications: European Conference on the Applications of Evolutionary Computation. Applications of Evolutionary Computation. 19th European Conference, EvoApplications 2016, Porto, Portugal, March 30 - April 1, 2016, Proceedings, Part I, volume 9597 of Lecture Notes in Computer Science*, pages 283-298.

Fabio D'Andreagiovanni, Giovanni Felici, Fabrizio Lacalandra (2016). Multiband Robust Optimization for optimal energy offering under price uncertainty. In *Proc. of ROADEF 2016*.

M. N. Dean, R. Seidel, D. Knötel, K. Lyons, D. Baum, J. C. Weaver, P. Fratzl (2016). To build a shark: 3-D tiling laws of tessellated cartilage. In *Abstract in Integrative and Comparative Biology; conference Society of Integrative and Comparative Biology annual meeting, January 3-7, 2016, Portland, USA*, 56 (suppl 1).

Jason Dunlop, Dmitry Apanaskevich, Jens Lehmann, Rene Hoffmann, Florian Füsseis, Moritz Ehlke, Stefan Zachow, Xianghui Xiao (2016). Microtomography of the Baltic amber tick *Ixodes succineus* reveals affinities with the modern Asian disease vector *Ixodes ovatus*. *BMC Evolutionary Biology*, 16(1).

Rainald Ehrig, Thomas Dierkes, Stefan Schäfer, Susanna Röblitz, Enrico Tronci, Toni Mancini, Ivano Salvo, Vadim Alimguzhin, Federico Mari, Igor Melatti, Annalisa Massini, Tillmann H. C. Krüger, Marcel Egli, Fabian Ille, Brigitte Leeners (2016). An integrative approach for model driven computation of treatments in reproductive medicine. In *BIOMAT - Proceedings of the 15th International Symposium on Mathematical and Computational Biology, Rorkee, India*.

David Eisenhauer, Klaus Jäger, Grit Köppel, Bernd Rech, Christiane Becker (2016). Optical Properties of Smooth Anti-reflective Three-dimensional Textures for Silicon Thin-film Solar Cells. *Energy Procedia*, 102:27-35.

J. Fajerski, M. Noack, A. Reinefeld, F. Schintke, T. Schütt, T. Steinke (2016). Fast In-Memory Checkpointing with POSIX API for Legacy Exascale-Applications. In *SPPEXA Symposium 2016*.

Bálint Fodor, Peter Kozma, Sven Burger, Miklos Fried, Peter Petrik (2016). Effective medium approximation of ellipsometric response from random surface roughness simulated by finite-element method. *Thin Solid Films*, 617:20.

Claudia Färber, Jürgen Titschack, Christine H. L. Schönberg, Karsten Ehrig, Karin Boos, Daniel Baum, Bernd Illerhaus, Ulla Asgaard, Richard G. Bromley, André Freiwald, Max Wisshak (2016). Long-term macrobioerosion in the Mediterranean Sea assessed by micro-computed tomography. *Biogeosciences*, 13(11):3461-3474.

Gerald Gamrath, Thorsten Koch, Stephen Maher, Daniel Rehfeldt, Yuji Shinano (2016). SCIP-Jack - A solver for STP and variants with parallelization extensions. *Mathematical Programming Computation*, 1-66.

Patrick Gelß, Sebastian Matera, Christof Schütte (2016). Solving the master equation without kinetic Monte Carlo: Tensor train approximations for a CO oxidation model. *Journal of Computational Physics*, 314:489-502.

Ambros M. Gleixner, Timo Berthold, Benjamin Müller, Stefan Weltge (2016). Three Enhancements for Optimization-Based Bound Tightening. *Journal of Global Optimization*.

Ambros M. Gleixner, Daniel E. Steffy, Kati Wolter (2016). Iterative Refinement for Linear Programming. *INFORMS Journal on Computing*, 28(3):449-464.

Leonid Goubergrits, Jan Osman, Ricardo Mevert, Ulrich Kertzscher, Kai Pöthkow, Hans-Christian Hege (2016). Turbulence in blood damage modeling. *The International Journal of Artificial Organs*, 39(4):147-210.

Carl Martin Grewe, Stefan Zachow (2016). Fully Automated and Highly Accurate Dense Correspondence for Facial Surfaces. *Computer Vision - ECCV 2016 Workshops*, 9914:552-568.

Boris Grimm, Ralf Borndörfer, Markus Reuther, Thomas Schlechte, Stanley Schade (2016). Regularity patterns for rolling stock rotation optimization. In 8th International Conference on Applied Operational Research, *Proceedings, volume 8 of Lecture Notes in Management Science*, pages 28-32.

Philipp Gutsche, Raquel Mäusle, Sven Burger (2016). Locally Enhanced and Tunable Optical Chirality in Helical Metamaterials. *Photonics*, 3:60.

Philipp Gutsche, Raquel Mäusle, Sven Burger (2016). Tailoring local optical chirality in helical metamaterials. In 2016 10th International Congress on Advanced Electromagnetic Materials in Microwaves and Optics, *IEEE*, pages 73-75.

- Philipp Gutsche, Lisa V. Poulikakos, Martin Hammerschmidt, Sven Burger, Frank Schmidt (2016). Time-harmonic optical chirality in inhomogeneous space. *Proc. SPIE*, 9756:97560X.
- Tobias Günther, Alexander Kuhn, Hans-Christian Hege, Markus Gross, Holger Theisel (2016). Progressive Monte-Carlo Rendering of Atmospheric Flow Features Across Scales. In *69th Annual Meeting of the APS Division of Fluid Dynamics, Gallery of Fluid Motion, Nov 20-22, 2016, Portland, OR, USA*.
- Tobias Günther, Alexander Kuhn, Hans-Christian Hege, Holger Theisel (2016). MCFTLE: Monte Carlo Rendering of Finite-Time Lyapunov Exponent Fields. *Computer Graphics Forum*, 35(3):381-390.
- Martin Hammerschmidt, Sven Herrmann, Jan Pomplun, Sven Burger, Frank Schmidt (2016). Model order reduction for the time-harmonic Maxwell equation applied to complex nanostructures. *Proc. SPIE 9742, Physics and Simulation of Optoelectronic Devices XXIV*, 9742M.
- Martin Hammerschmidt, Sven Herrmann, Jan Pomplun, Sven Burger, Frank Schmidt (2016). Reduced basis method for electromagnetic scattering problem: a case study for FinFETs. *Optical and Quantum Electronics*, 48:250.
- Martin Hammerschmidt, Carlo Barth, Jan Pomplun, Sven Burger, Christiane Becker, Frank Schmidt (2016). Reconstruction of photonic crystal geometries using a reduced basis method for nonlinear outputs. *Proc. SPIE 9756, Photonic and Phononic Properties of Engineered Nanostructures VI*, 97561R.
- Carsten Hartmann, Christof Schütte, Wei Zhang (2016). Model reduction algorithms for optimal control and importance sampling of diffusions. *Nonlinearity*, 29(8):2298-2326.
- Gregor Hendel (2016). Exploiting Solving Phases for Mixed-Integer Programs. In *Operations Research Proceedings 2015*, 3-9.
- Gunter Hermann, Vincent Pohl, Jean Christophe Tremblay, Beate Paulus, Hans-Christian Hege, Axel Schild (2016). ORBKIT - A modular Python toolbox for cross-platform post-processing of quantum chemical wavefunction data. *Journal of Computational Chemistry*, 37(16):1511-1520.
- Katharina Heye, Dennis Becker, Christian Lütke-Eversloh, Vedat Durmaz, Thomas A. Ternes, Matthias Oetken, Jörg Oehlmann (2016). Effects of carbamazepine and two of its metabolites on the non-biting midge *Chironomus riparius* in a sediment full life cycle toxicity test. *Water Research*, 98:19-27.
- C. Hoppe, P. Obermeier, S. Mehlhans, M. Alchikh, L. Seeber, F. Tief, K. Karsch, X. Chen, S. Boettcher, S. Diedrich, T. Conrad (2016). Innovative Digital Tools and Surveillance Systems for the Timely Detection of Adverse Events at the Point of Care: A Proof-of-Concept Study. *Drug Safety*, 39(10):977-988.
- Heide Hoppmann (2016). An Extended Formulation for the Line Planning Problem. In *Operations Research Proceedings 2015*, 11-17.
- Jesco Humpola, Felipe Serrano (2016). Sufficient pruning conditions for MINLP in gas network design. *EURO Journal on Computational Optimization*.
- Nikolay Ledentsov, Jr., Vitaly A. Shchukin, Nikolay N. Ledentsov, Jörg-R. Kropp, Sven Burger, Frank Schmidt (2016). Direct Evidence of the Leaky Emission in Oxide-Confined Vertical Cavity Lasers. *IEEE J. Quant. Electron.*, 52(3):2400207.
- Nikolay Ledentsov, Jr., Vitaly A. Shchukin, Jörg-R. Kropp, Sven Burger, Frank Schmidt, Nikolay N. Ledentsov (2016). Direct visualization of the in-plane leakage of high-order transverse modes in vertical-cavity surface-emitting lasers mediated by oxide-aperture engineering. *Proc. SPIE*, 9766:976608.
- Leonardo Agudo Jácome, Hans-Christian Hege, Olaf Paetsch, Kai Pöthkow (2016). 3-D Reconstruction, Visualization and Quantification of Dislocations from Transmission Electron Microscopy Stereo-Pairs. In *Microscopy and Microanalysis 2016, July 24-28 Columbus, Ohio*.
- Klaus Jäger, Martin Hammerschmidt, Grit Köppel, Sven Burger, Christiane Becker (2016). On Accurate Simulations of Thin-Film Solar Cells With a Thick Glass Superstrate. In *Light, Energy and the Environment 2016, OSA Technical Digest*, page PM3B.5.
- Klaus Jäger, Carlo Barth, Martin Hammerschmidt, Sven Herrmann, Sven Burger, Frank Schmidt, Christiane Becker (2016). Simulations of sinusoidal nanotextures for coupling light into c-Si thin-film solar cells. *Opt. Express*, 24:A569.
- Klaus Jäger, Grit Köppel, Carlo Barth, Martin Hammerschmidt, Sven Herrmann, Sven Burger, Frank Schmidt, Christiane Becker (2016). Sinusoidal gratings for optimized light management in c-Si thin-film solar cells. *Proc. SPIE*, 9898:989808.
- Marika Karbstein (2016). Integrated Line Planning and Passenger Routing: Connectivity and Transfers. In *Operations Research Proceedings 2014*, 263-269.
- Jens Kasten, Jan Reininghaus, Ingrid Hotz, Hans-Christian Hege, Bernd R. Noack, Guillaume Daviller, Marek Morzyński (2016). Acceleration feature points of unsteady shear flows. *Archives of Mechanics*, 68(1):55-80.

Stefan Klus, Peter Koltai, Christof Schütte (2016). On the numerical approximation of the Perron-Frobenius and Koopman operator. *Journal of Computational Dynamics*, 3(1):51-77.

Stefan Klus, Christof Schütte (2016). Towards tensor-based methods for the numerical approximation of the Perron-Frobenius and Koopman operator. *Journal of Computational Dynamics*.

James C Knight, Philip J Tully, Bernhard A Kaplan, Anders Lansner, Steve Furber (2016). Large-scale simulations of plastic neural networks on neuromorphic hardware. *Frontiers in Neuroanatomy*, 10:37.

Peter Koltai, Giovanni Ciccotti, Christof Schütte (2016). On metastability and Markov state models for non-stationary molecular dynamics. *The Journal of Chemical Physics*, 145(174103).

Barbora Kozlíková, Michael Krone, Martin Falk, Norbert Lindow, Marc Baaden, Daniel Baum, Ivan Viola, Julius Parulek, Hans-Christian Hege (2016). Visualization of Biomolecular Structures: State of the Art Revisited. *Computer Graphics Forum*.

Tobias Kramer, Christoph Kreisbeck, Christian Riha, Olivio Chiatti, Sven Buchholz, Andreas Wieck, Dirk Reuter, Saskia Fischer (2016). Thermal energy and charge currents in multi-terminal nanorings. *AIP Advances*, 6:065306.

Tobias Kramer, Matthias Noack (2016). On the origin of inner coma structures observed by Rosetta during a diurnal rotation of comet 67P/Churyumov-Gerasimenko. *The Astrophysical Journal Letters*, 823(1):L11.

Tobias Kramer, Mirta Rodriguez, Yaroslav Zelinskyi (2016). Modeling of Transient Absorption Spectra in Exciton Charge-Transfer Systems. *Journal of Physical Chemistry B*.

Michael Krone, Barbora Kozlíková, Norbert Lindow, Marc Baaden, Daniel Baum, Julius Parulek, Hans-Christian Hege, Ivan Viola (2016). Visual Analysis of Biomolecular Cavities: State of the Art. *Computer Graphics Forum*, 35(3):527-551.

Olaf Krzikalla, Florian Wende, Markus Höhnertbach (2016). Dynamic SIMD Vector Lane Scheduling. In *High Performance Computing, ISC High Performance 2016 International Workshops, ExaComm, E-MuCoCoS, HPC-IODC, IXPUG, IWOPH, P*3MA, VHPC, WOPSSS, volume 9945 of LNCS*, pages 354-365.

Hans Lamecker, Stefan Zachow (2016). Statistical Shape Modeling of Musculoskeletal Structures and Its Applications. In *Computational Radiology for Orthopaedic Interventions, volume 23 of Lecture Notes in Computational Vision and Biomechanics*, pages 1-23, Springer.

Itamar D. Landau, Robert Egger, Vincent J. Dercksen, Marcel Oberlaender, Haim Sompolinsky (2016). The Impact of Structural Heterogeneity on Excitation-Inhibition Balance in Cortical Networks. *Neuron*, 92(5):1106-1121.

Alexander Lange (2016). Reconstruction of disease transmission rates: applications to measles, dengue, and influenza. *Journal of Theoretical Biology*, 400:138-153.

Alexander Lange, Julia Plöntzke, Stefan Schäfer, Susanna Röblitz (2016). Follicular maturation in cows: mathematical models and data. In *10. European Conference on Mathematical and Theoretical Biology*.

Markus Leitner, Ivana Ljubic, Markus Sinnl, Axel Werner (2016). ILP heuristics and a new exact method for bi-objective 0/1 ILPs: Application to FTTx-network design. *Computers & Operations Research*.

Robert Lemanis, Dieter Korn, Stefan Zachow, Erik Rybacki, René Hoffmann (2016). The Evolution and Development of Cephalopod Chambers and Their Shape. *PLOS ONE*, 11(3).

Robert Lemanis, Stefan Zachow, René Hoffmann (2016). Comparative cephalopod shell strength and the role of septum morphology on stress distribution. *PeerJ*, 4:e2434.

Lars Lubkoll, Anton Schiela, Martin Weiser (2016). An affine covariant composite step method for optimization with PDEs as equality constraints. *Optimization Methods and Software*.

Stephen J. Maher, John M. Murray (2016). The unrooted set covering connected subgraph problem differentiating between HIV envelope sequences. *European Journal of Operational Research*, 248(2):668-680.

Stephen Maher, Matthias Miltenberger, Joao Pedro Pedrosa, Daniel Rehfeldt, Robert Schwarz, Felipe Serrano (2016). PySCIPOpt: Mathematical Programming in Python with the SCIP Optimization Suite. In *Mathematical Software - ICMS 2016, volume 9725 of Lecture Notes in Computer Science*, pages 301-307.

Anirban Mukhopadhyay, Arun Kumar, Suchendra Bhandarkar (2016). Joint Geometric Graph Embedding for Partial Shape Matching in Images. *IEEE Winter Conference on Applications of Computer Vision*, 1-9.

Anirban Mukhopadhyay, Fatih Porikli, Suchendra Bhandarkar (2016). Detection and Characterization of Intrinsic Symmetry of 3-D Shapes. *Proceedings of IEEE International Conference on Pattern Recognition*.

Anirban Mukhopadhyay (2016). Total Variation Random Forest: Fully automatic MRI segmentation in congenital heart disease. *RAMBO 2016, HVSMR 2016: Reconstruction, Segmentation, and Analysis of Medical Images, LNCS 10129:165-171*.

Anirban Mukhopadhyay, Oscar Morillo, Stefan Zachow, Hans Lamecker (2016). Robust and Accurate Appearance Models Based on Joint Dictionary Learning Data from the Osteoarthritis Initiative. *Lecture Notes in Computer Science, Patch-Based Techniques in Medical Imaging. Patch-MI 2016*, 9993:25-33.

Adam Nielsen (2016). The Monte Carlo Computation Error of Transition Probabilities. *Statistics & Probability Letters*, 118:163-170.

P. Obermeier, S. Muehlhans, Ch. Hoppe, K. Karsch, F. Tief, L. Seeber, X. Chen, T. Conrad, S. Boettcher, S. Diedrich, B. Rath (2016). Enabling Precision Medicine With Digital Case Classification at the Point-of-Care. *EBio-Medicine*, 4:191-196.

Jan Pomplun, Sven Burger, Lin Zschiedrich, Philipp Gutsche, Frank Schmidt (2016). Method for fast computation of angular light scattering spectra from 2-D periodic arrays. *Proc. SPIE*, 9778:977839.

Lisa Poulidakos, Philipp Gutsche, Kevin McPeak, Sven Burger, Jens Niegemann, Christian Hafner, David Norris (2016). The Optical Chirality Flux as a Useful Far-Field Probe of Chiral Near Fields. *ACS Photonics*, 3:1619.

J. Reimers, M. Biczysko, D. Bruce, D.F. Coker, T.J. Frankcombe, H. Hashimoto, J. Hauer, R. Jankowiak, Tobias Kramer, J. Linnanto, F. Mamedov, F. Müh, M. Rätsep, T. Renger, S. Styring, J. Wan, Z. Wang, Z.-Y. Wang-Otomo, Y.-X. Weng, C. Yang, J.-P. Zhang, A. Freiberg, E. Krausz (2016). Challenges facing an understanding of the nature of low-energy excited states in photosynthesis. *BBA Bioenergetics*, 1857(9):1627-1640.

Beate Rusch, Markus Putnigs (2016). DeepGreen - Entwicklung eines rechtssicheren Workflows zur effizienten Umsetzung der Open-Access-Komponente in den Allianz-Lizenzen für die Wissenschaft. In Klaus-Rainer Brintzinger, Ulrich Hohoff, Ulrike Scholle, Thomas Stäcker, Helge Steenweg, Heidrun Wiesenmüller, editors, *o-bib. Das offene Bibliotheksjournal / herausgegeben vom VDB*, 4/2016, volume 4/2016 of *o-bib*, pages 110-118.

Manish Sahu, Daniil Moerman, Philip Mewes, Peter Mountney, Georg Rose (2016). Instrument State Recognition and Tracking for Effective Control of Robotized Laparoscopic Systems. *International Journal of Mechanical Engineering and Robotics Research*, 5(1):33-38.

Stanley Schade, Martin Strehler (2016). The Maximum Flow Problem for Oriented Flows. In Marc Goerigk, Renato Werneck, editors, *16th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2016)*, volume 54 of *Open Access Series in Informatics (OASIS)*, pages 1-13.

Robert Schmidtke, Guido Laubender, Thomas Steinke (2016). Big Data Analytics on Cray XC Series DataWarp using Hadoop, Spark and Flink. In *CUG Proceedings*.

Peter Schnauber, Alexander Thoma, Christoph V. Heine, Alexander Schlehahn, Liron Gantz, Manuel Gschrey, Ronny Schmidt, Caspar Hopfmann, Benjamin Wohlfeil, Jan-Hindrik Schulze, Andre Strittmatter, Tobias Heindel, Sven Rodt, Ulrike Woggon, David Gershoni, Stephan Reitzenstein (2016). Bright Single-Photon Sources Based on Anti-Reflection Coated Deterministic Quantum Dot Microlenses. *Technologies*, 4(1):1.

Stefan Schäfer, Julia Plöntzke, Susanna Röblitz (2016). Mathematical Modelling of Follicular Maturation in Cows and Women. In *49. Jahrestagung der Physiologie und Pathologie der Fortpflanzung und gleichzeitig 41. Veterinär-Humanmedizinische Gemeinschaftstagung 2016*, Leipzig.

Borong Shao, Tim Conrad (2016). Epithelial Mesenchymal Transition Regulatory Network-based Feature Selection in Lung Cancer Prognosis Prediction. *Lecture Notes in Computer Science (LNCS)*, 9656:1235-146.

Yuji Shinano, Tobias Achterberg, Timo Berthold, Stefan Heinz, Thorsten Koch, Michael Winkler (2016). Solving Open MIP Instances with ParaSCIP on Supercomputers using up to 80,000 Cores. In *Proc. of 30th IEEE International Parallel & Distributed Processing Symposium*.

Victor Soltwisch, Anton Haase, Jan Wernecke, Jürgen Probst, Max Schoengen, Sven Burger, Michael Krümrey, Frank Scholze (2016). Correlated diffuse x-ray scattering from periodically nanostructured surfaces. *Phys. Rev. B*, 94:035419.

Gunther Sprösser, Sebastian Schenker, Andreas Pittner, Ralf Borndörfer, Michael Rethmeier, Ya-Ju Chang, Matthias Finkbeiner (2016). Sustainable Welding Process Selection based on Weight Space Partitions. In *Procedia CIRP*, volume 40 of 13th Global Conference on Sustainable Manufacturing - Decoupling Growth from Resource Use, pages 127-132.

Andreas Spötter, Pooja Gupta, Manfred Mayer, Norbert Reinsch, Kaspar Bienefeld (2016). Genome-wide association study of a Varroa-specific defense behavior in honeybees (*Apis mellifera*). *Journal of Heredity*, 107 (3):220-227.

Alexander Tesch (2016). A Nearly Exact Propagation Algorithm for Energetic Reasoning in $O(n^2 \log n)$. In *International Conference on Principles and Practice of Constraint Programming (CP 2016)*, volume 22 of Lecture Notes in Computer Science, pages 493-519.

F. Tief, Ch. Hoppe, L. Seeber, P. Obermeier, X. Chen, K. Karsch, S. Muehlhans, E. Adamou, T. Conrad, B. Schweiger, T. Adam, B. Rath (2016). An inception cohort study assessing the role of bacterial co-infections in children with influenza and ILI and a clinical decision model for stringent antibiotic use. *Antiviral Therapy*, 21:413-424.

Jürgen Titschack, Hiske G. Fink, Daniel Baum, Claudia Wienberg, Dierk Hebbeln, André Freiwald (2016). Mediterranean cold-water corals - an important regional carbonate factory? *The Depositional Record*, 2(1):74-96.

Iliusi Vega, Christof Schütte, Tim Conrad (2016). Finding metastable states in real-world time series with recurrence networks. *Physica A: Statistical Mechanics and its Applications*, 445:1-17.

José Villatoro, Martin Zülke, Daniel Riebe, Toralf Beitz, Marcus Weber, Jens Riedel, Hans-Gerd Löhmannsröben (2016). IR-MALDI ion mobility spectrometry: physical source characterization and application as HPLC detector. *International Journal for Ion Mobility Spectrometry*, 197-297.

C. Weinhold, A. Lackorzynski, J. Bierbaum, M. Küttler, M. Planeta, H. Härtig, A. Shiloh, E. Levy, T. Ben-Nun, A. Barak, T. Steinke, T. Schütt, J. Fajerski, A. Reinefeld, M. Lieber, W. E. Nagel (2016). FFMK: A Fast and Fault-tolerant Microkernel-based System for Exascale Computing. In *SPPEXA Symposium 2016*.

Florian Wende, Martijn Marsman, Thomas Steinke (2016). On Enhancing 3-D-FFT Performance in VASP. In *CUG Proceedings*.

Florian Wende, Matthias Noack, Thomas Steinke, Michael Klemm, Georg Zitzlsberger, Chris J. Newburn (2016). Portable SIMD Performance with OpenMP* 4.x Compiler Directives. In Pierre-Francois Dutot, Denis Trystam, editors, , volume Euro-Par 2016: Parallel Processing; 22nd International Conference on Parallel and Distributed Computing of LNCS.

Stefanie Winkelmann, Christof Schütte (2016). The Spatiotemporal Master Equation: Approximation of Reaction-Diffusion Dynamics via Markov State Modeling. *Journal of Chemical Physics*, 145(21).

Stefanie Winkelmann (2016). Markov Control with Rare State Observation: Average Optimality. *Markov Processes and Related Fields*. (accepted for publication on 2016-08-29)

Benjamin Wohlfeil, Georg Rademacher, Christos Stamatiadis, Karsten Voigt, Lars Zimmermann, Klaus Petermann (2016). A Two Dimensional Fiber Grating Coupler on SOI for Mode Division Multiplexing. *IEEE Photon. Technol. Lett.*, 28:1241.

Robert Zahn, Sarah Grotjohann, Heiko Ramm, Stefan Zachow, Matthias Pumberger, Michael Putzier, Carsten Perka, Stephan Tohtz (2016). Influence of pelvic tilt on functional acetabular orientation. *Technology and Health Care*, 1-9.

Thomas Zander, Marcel Dreischarf, Anne-Katrin Timm, Wolfgang Baumann, Hendrik Schmidt (2016). Impact of material and morphological parameters on the mechanical response of the lumbar spine - A finite element sensitivity study. *Journal of Biomechanics*, 53:185-190.

L. Krebek, von, A. Achazi, M. Solleder, M. Weber, B. Paulus, C. Schalley (2016). Allosteric and Chelate Cooperativity in Divalent Crown Ether-Ammonium Complexes with Strong Binding Enhancements. *Chem. Eur. J.*, 22(43):15475-15484.

NOT PEER-REVIEWED (17)

Pascal-Nicolas Becker, Roland Bertelmann, Klaus Ceynowa, Jürgen Christof, Thomas Dierkes, Julia Alexandra Goltz, Matthias Groß, Regina Heidrich, Tobias Höhnow, Michael Kasube, Thorsten Koch, Monika Kuberek, Lilian Landes, Heinz Pampel, Markus Putnings, Beate Rusch, Hildegard Schäffler, Dagmar Schobert, Oliver Schwab, Jens Schwidder, Konstanze Söllner, Tonka Stoyanova, Paul Vierkant (2016). Questionnaire for effective exchange of bibliographic metadata – current status of publishing houses.

Florian Bernard, Luis Salamanca, Johan Thunberg, Alexander Tack, Dennis Jentsch, Hans Lamecker, Stefan Zachow, Frank Hertel, Jorge Goncalves, Peter Gemmar (2016). Shape-aware Surface Reconstruction from Sparse Data. *arXiv*, 1602.08425v1.

Timo Berthold, James Farmer, Stefan Heinz, Michael Perregaard (2016). Parallelization of the FICO Xpress-Optimizer. In *Mathematical Software – ICMS 2016, 5th International Conference Berlin, Germany, July 11-14, 2016 Proceedings*, 251-258.

Hagen Chrapary, Yue Ren (2016). The Software Portal swMATH: A State of the Art Report and Next Steps. In *Mathematical Software – ICMS 2016, 5th International Conference Berlin, Germany, July 11-14, 2016 Proceedings*, 397-402.

Carl Martin Grewe, Lisa Schreiber (2016). Digitale Bildarchive. Archivierung und Codierung der Gefühle. In *ultra gestaltung schafft wissen*, 285-290, Seemann Henschel.

Sebastian Götschel, Christian Höhne, Sanjeevareddy Kolkooori, Steffen Mitzscherling, Jens Prager, Martin Weiser (2016). Ray Tracing Boundary Value Problems: Simulation and SAFT Reconstruction for Ultrasonic Testing. In *Proceedings 19th World Conference on Non-Destructive Testing (WCNDT 2016)*.

Sebastian Götschel, Christiane Maierhofer, Jan Müller, Nick Rothbart, Martin Weiser (2016). Quantitative Defect Reconstruction in Active Thermography for Fiber-Reinforced Composites. In *Proceedings 19th World Conference on Non-Destructive Testing (WCNDT 2016)*.

Helene Hahn (2016). Kooperativ in die digitale Zeit – wie öffentliche Kulturinstitutionen Cultural Commons fördern. digiS – Servicestelle Digitalisierung Berlin, editor.

Klaus Jäger, Carlo Barth, Martin Hamerschmidt, Sven Herrmann, Sven Burger, Frank Schmidt, Christiane Becker (2016). Sinusoidal Nanotextures for Enhanced Light Management in Thin-Film Solar Cells. In *European Society for Quantum Solar Energy Conversion*, editor, *28th Workshop on Quantum Solar Energy Conversion – (QUANTSOL)*.

Han Cheng Lie, Tim Sullivan (2016). Cameron-Martin theorems for sequences of Cauchy-distributed random variables. *arXiv*, 1608.03784.

Stefan Lohrum, Mathias Kratzer, Uwe Risch, Peter Thiessen (2016). Zum Stand des Projektes “Cloudbasierte Infrastruktur für Bibliotheksdaten” (CIB). *Zeitschrift für Bibliotheks- und Bibliographie*, Jahrgang 63 (Heft 5-6):250-257.

Anja Müller, Beate Rusch (2016). Was in der Zwischenzeit geschah – vier Jahre und 49 Digitalisierungsprojekte später: Förderprogramm Digitalisierung und Servicestelle Digitalisierung Berlin (digiS) 2012 bis 2016. In *Föderale Vielfalt – Globale Vernetzung. Strategien der Bundesländer für das kulturelle Erbe in der digitalen Welt*.

Beate Rusch, Roland Bertelmann, Heinz Pampel, Pamela Aust, Kerstin Helbig, Anja Doreen Müller, Steffi Conrad-Rempel, Signe Weihe, Thomas Dierkes, Rainer Kuhlen, Dagmar Schobert, Michaela Voigt, Paul Klimpel, Christina Riesenweber, Andreas Kennecke, Mario Kowalak, Birgit Schlegel, Niels Taubert, Andreas Hübner, Katja Mruck (2016). KOBV-Sonderedition zur International Open Access Week 2016 “Open in Action”.

Beate Rusch, Rita Albrecht, Gabriele Meßmer, Peter Thiessen (2016). Katalogisierung in der Datenwolke. *Zeitschrift für Bibliotheks- und Bibliographie*, Jahrgang 63 (Heft 5-6):258-264.

Guillaume Sagnol, Hans-Christian Hege, Martin Weiser (2016). Using sparse kernels to design computer experiments with tunable precision. In 22nd Intern. Conf. on Computational Statistics – COMSTAT 2016, Oviedo, Spain, 23-26 August 2016, *Proceedings* ISBN 978-90-73592-36-0, 397-408.

Yuji Shinano, Timo Berthold, Stefan Heinz (2016). A First Implementation of ParaXpress: Combining Internal and External Parallelization to Solve MIPs on Supercomputers. In *Mathematical Software – ICMS 2016, 5th International Conference Berlin, Germany, July 11-14, 2016 Proceedings*, 308-316.

Alexander Tesch (2016). Kompakte MIP Modelle für das Ressourcenbeschränkte Projektplanungsproblem. *OR News*, 58:19-21.

BOOKS (2)

Martin Weiser (2016). Inside Finite Elements. De Gruyter.

Martin Greuel, Thorsten Koch, Peter Paule, Andrew Sommese, editors (2016). Mathematical Software - ICMS 2016, 5th Int. Conf. Berlin, Germany, July 11-14, 2016, Proceedings. Springer.

DISSERTATIONS (8)

Animesh Agarwal (2016). Path Integral Techniques in Molecular Dynamics Simulations of Open Boundary Systems, Freie Universität Berlin.

Tomasz Badowski (2016). Adaptive importance sampling via minimization of estimators of cross-entropy, mean square and inefficiency constants, Freie Universität Berlin.

Adman Daragmeh (2016). Model Order Reduction of Linear Control Systems: Comparison of Balance Truncation and Singular Perturbation Approximation with Application to Optimal Control, Freie Universität Berlin.

Raheem Gul (2016). Mathematical Modeling and Sensitivity Analysis of Lumped Parameter Model of the Human Cardiovascular System, Freie Universität Berlin.

Han Cheng Lie (2016). On a strongly convex approximation of a stochastic optimal control problem for importance sampling of metastable diffusions, Freie Universität Berlin.

Adam Nielsen (2016). Computation Schemes for Transfer Operators. Freie Universität Berlin.

Nisara Sriwattanaworachai (2016). Spectral approach to metastability of non-reversible complex processes, Freie Universität Berlin.

Illiusi Vega (2016). Reconstruction and analysis of the state space for the identification of dynamical states in real-world time-series, Freie Universität Berlin.

ZIB REPORTS (73)

Tobias Achterberg, Robert E. Bixby, Zonghao Gu, Edward Rothberg, Dieter Weninger (2016). Presolve Reductions in Mixed Integer Programming. ZIB-Report 16-44.

Kilian Amrhein, Tim Hasler, Marco Klindt, Elias Oltmanns, Wolfgang Peters-Kottig (2016). Digitale Langzeitarchivierung: Mustervorlage für Übereinkommenvereinbarungen. ZIB-Report 16-27.

Daniel Baum, Jürgen Titschack (2016). Cavity and Pore Segmentation in 3-D Images with Ambient Occlusion. ZIB-Report 16-17.

Isabel Beckenbach, Robert Scheidweiler (2016). Perfect f-Matchings and f-Factors in Hypergraphs - A Combinatorial Approach. ZIB-Report 16-22.

Isabel Beckenbach, Leon Eifler, Konstantin Fackeldey, Ambros Gleixner, Andreas Grever, Marcus Weber, Jakob Witzig (2016). Mixed-Integer Programming for Cycle Detection in Non-reversible Markov Processes. ZIB-Report 16-39.

Pascal-Nicolas Becker, Roland Bertelmann, Klaus Ceynowa, Jürgen Christof, Thomas Dierkes, Julia Alexandra Goltz, Matthias Groß, Regina Heidrich, Tobias Höhnnow, Michael Kasube, Thorsten Koch, Monika Kuberek, Lilian Landes, Heinz Pampel, Markus Putnings, Beate Rusch, Hildegard Schäffler, Dagmar Schobert, Oliver Schwab, Jens Schwidder, Konstanze Söllner, Tonka Stoyanova, Paul Vierkant (2016). DeepGreen - Metadata Schema for the exchange of publications between publishers and open-access repositories. Version 1.1. June 2016. ZIB-Report 16-32.

Timo Berthold, Gregor Hendel, Thorsten Koch (2016). The Three Phases of MIP Solving. ZIB-Report 16-78.

Marco Blanco, Ralf Borndörfer, Nam Dung Hoang, Anton Kaier, Thomas Schlechte, Swen Schlobach (2016). The Shortest Path Problem with Crossing Costs. ZIB-Report 16-70.

Ralf Borndörfer, Oytun Arslan, Ziena Eljazyfer, Hakan Güler, Malte Renken, Güvenc Sahin, Thomas Schlechte (2016). Line Planning on Path Networks with Application to the Istanbul Metrobüs. ZIB-Report 16-38.

Ralf Borndörfer, Heide Hoppmann, Marika Karbstein, Fabian Löbel (2016). The Modulo Network Simplex with Integrated Passenger Routing. ZIB-Report 16-43.

Ralf Borndörfer, Boris Grimm, Markus Reuther, Thomas Schlechte (2016). Optimization of Handouts for Rolling Stock Rotations Visualization. ZIB-Report ZR-16-73.

Christopher Brandt, Christoph von Tycowicz, Klaus Hildebrandt (2016). Geometric Flows of Curves in Shape Space for Processing Motion of Deformable Objects. ZIB-Report 16-29.

Kevin K. H. Cheung, Ambros Gleixner, Daniel E. Steffy (2016). Verifying Integer Programming Results. ZIB-Report 16-58.

Jon Cockayne, Chris Oates, Tim Sullivan, Mark Girolami (2016). Probabilistic Meshless Methods for Partial Differential Equations and Bayesian Inverse Problems. ZIB-Report 16-31.

Marta Costa, James D. Manton, Aaron D. Ostrovsky, Steffen Prohaska, Gregory S. X. E. Jefferis (2016). NBLAST: Rapid, sensitive comparison of neuronal structure and construction of neuron family databases. ZIB-Report 16-34.

Fabio D'Andreagiovanni (2016). Revisiting wireless network jamming by SIR-based considerations and Multiband Robust Optimization. ZIB-Report 15-12.

Fabio D'Andreagiovanni, Ambros Gleixner (2016). Towards an accurate solution of wireless network design problems. ZIB-Report 16-12.

Rainald Ehrig, Thomas Dierkes, Stefan Schäfer, Susanna Röblitz, Enrico Tronci, Toni Mancini, Ivano Salvo, Vadim Alinguzhin, Federico Mari, Igor Melatti, Annalisa Massini, Tillmann H. C. Krüger, Marcel Egli, Fabian Ille, Brigitte Leeners (2016). An Integrative Approach for Model Driven Computation of Treatments in Reproductive Medicine. ZIB-Report 16-04.

Frank Fischer, Boris Grimm, Torsten Klug, Thomas Schlechte (2016). A Re-optimization Approach for Train Dispatching. ZIB-Report 16-49.

Gerald Gamrath, Ambros Gleixner, Thorsten Koch, Matthias Miltenberger, Dimitri Kniasev, Dominik Schlögel, Alexander Martin, Dieter Weninger (2016). Tackling Industrial-Scale Supply Chain Problems by Mixed-Integer Programming. ZIB-Report 16-45.

Gerald Gamrath, Tobias Fischer, Tristan Gally, Ambros M. Gleixner, Gregor Hendel, Thorsten Koch, Stephen J. Maher, Matthias Miltenberger, Benjamin Müller, Marc E. Pfetsch, Christian Puchert, Daniel Rehfeldt, Sebastian Schenker, Robert Schwarz, Felipe Serrano, Yuji Shinano, Stefan Vigerske, Dieter Weninger, Michael Winkler, Jonas T. Witt, Jakob Witzig (2016). The SCIP Optimization Suite 3.2. ZIB-Report 15-60.

Gerald Gamrath, Thorsten Koch, Stephen J. Maher, Daniel Rehfeldt, Yuji Shinano (2016). SCIP-Jack - A solver for STP and variants with parallelization extensions. ZIB-Report 16-41.

Ambros M. Gleixner, Timo Berthold, Benjamin Müller, Stefan Weltge (2016). Three Enhancements for Optimization-Based Bound Tightening. ZIB-Report 15-16.

Robert Lion Gottwald, Stephen J. Maher, Yuji Shinano (2016). Distributed domain propagation. ZIB-Report 16-71.

Andreas Griewank, Tom Streubel, Richard Hasenfelder, Manuel Radons (2016). Piecewise linear secant approximation via Algorithmic Piecewise Differentiation. ZIB-Report 16-54.

Philipp Gutsche, Matthias Läuter, Frank Schmidt (2016). Parameter-dependent Parallel Block Sparse Arnoldi and Döhler Algorithms on Distributed Systems. ZIB-Report 16-15.

Sebastian Götschel, Christiane Maierhofer, Jan P. Müller, Nick Rothbart, Martin Weiser (2016). Quantitative Defect Reconstruction in Active Thermography for Fiber-Reinforced Composites. ZIB-Report 16-13.

Sebastian Götschel, Christian Höhne, Sanjeevareddy Kolkoori, Steffen Mitzscherling, Jens Prager, Martin Weiser (2016). Ray Tracing Boundary Value Problems: Simulation and SAFT Reconstruction for Ultrasonic Testing. ZIB-Report 16-14.

Tobias Günther, Alexander Kuhn, Hans-Christian Hege, Holger Theisel (2016). MCFTLE: Monte Carlo Rendering of Finite-Time Lyapunov Exponent Fields. ZIB-Report 16-21.

Martin Hammerschmidt, Sven Herrmann, Jan Pomplun, Sven Burger, Frank Schmidt (2016). Model order reduction for the time-harmonic Maxwell equation applied to complex nanostructures. ZIB-Report 16-05.

Martin Hammerschmidt, Sven Herrmann, Sven Burger, Jan Pomplun, Frank Schmidt (2016). Reduced basis method for the optimization of nano-photonics devices. ZIB-Report 16-10.

Martin Hammerschmidt, Carlo Barth, Jan Pomplun, Sven Burger, Christiane Becker, Frank Schmidt (2016). Reconstruction of photonic crystal geometries using a reduced basis method for nonlinear outputs. ZIB-Report 16-06.

Kai Hennig, Robert Schwarz (2016). Using Bilevel Optimization to find Severe Transport Situations in Gas Transmission Networks. ZIB-Report 16-68.

Benjamin Hiller, René Saitenmacher, Tom Walther (2016). Analysis of operating modes of complex compressor stations. ZIB-Report 16-61.

Heide Hoppmann (2016). An Extended Formulation for the Line Planning Problem. ZIB-Report 16-08.

Ilja Klebanov, Alexander Sikorski, Christof Schütte, Susanna Röblitz (2016). Prior estimation and Bayesian inference from large cohort data sets. ZIB-Report 16-09.

Ilja Klebanov, Alexander Sikorski, Christof Schütte, Susanna Röblitz (2016). Empirical Bayes Methods, Reference Priors, Cross Entropy and the EM Algorithm. ZIB-Report 16-56.

Ilja Klebanov, Alexander Sikorski, Christof Schütte, Susanna Röblitz (2016). Empirical Bayes Methods for Prior Estimation in Systems Medicine. ZIB-Report 16-57.

Thorsten Koch, Rolf Griebel, Konstanze Söllner, Jürgen Christof, Roland Bertelmann (2016). DeepGreen - Entwicklung eines rechtssicheren Workflows zur effizienten Umsetzung der Open-Access-Komponente in den Allianz-Lizenzen für die Wissenschaft. ZIB-Report 15-58.

Peter Koltai, Giovanni Ciccotti, Schütte Christof (2016). On metastability and Markov state models for non-stationary molecular dynamics. ZIB-Report 16-11.

Tobias Kramer, Matthias Noack (2016). On the origin of inner coma structures observed by Rosetta during a diurnal rotation of comet 67P/Churyumov-Gerasimenko. ZIB-Report 16-26.

Michael Krone, Barbora Kozlikova, Norbert Lindow, Marc Baaden, Daniel Baum, Julius Parulek, Hans-Christian Hege, Ivan Viola (2016). Visual Analysis of Biomolecular Cavities: State of the Art. ZIB-Report 16-42.

Alexander Kuhn, Wito Engelke, Markus Flatken, Hans-Christian Hege, Ingrid Hotz (2016). Topology-based Analysis for Multimodal Atmospheric Data of Volcano Eruptions. ZIB-Report 16-03.

Ralf Lenz, Robert Schwarz (2016). Optimal Looping of Pipelines in Gas Networks. ZIB-Report 16-67.

Han Cheng Lie, T. J. Sullivan (2016). Cameron-Martin theorems for sequences of Cauchy-distributed random variables. ZIB-Report 16-40.

Stephen Maher, Matthias Miltenberger, Joao Pedro Pedroso, Daniel Rehfeldt, Robert Schwarz, Felipe Serrano (2016). PySCIPOpt: Mathematical Programming in Python with the SCIP Optimization Suite. ZIB-Report 16-64.

Adam Nielsen (2016). The Monte Carlo Computation Error of Transition Probabilities. ZIB-Report 16-37.

Matthias Noack, Florian Wende, Georg Zitlsberger, Michael Klemm, Thomas Steinke (2016). KART - A Runtime Compilation Library for Improving HPC Application Performance. ZIB-Report 16-48.

Jonad Pulaj (2016). Cutting Planes for Families Implying Frankl's Conjecture. ZIB-Report 16-51.

Ted Ralphs, Yuji Shinano, Timo Berthold, Thorsten Koch (2016). Parallel Solvers for Mixed Integer Linear Programming. ZIB-Report 16-74.

Daniel Rehfeldt, Thorsten Koch (2016). Transformations for the Prize-Collecting Steiner Tree Problem and the Maximum-Weight Connected Subgraph Problem to SAP. ZIB-Report 16-36.

Daniel Rehfeldt, Thorsten Koch, Stephen Maher (2016). Reduction Techniques for the Prize-Collecting Steiner Tree Problem and the Maximum-Weight Connected Subgraph Problem. ZIB-Report 16-47.

Guillaume Sagnol, Christoph Barner, Ralf Borndörfer, Mickaël Grima, Mathees Seeling, Claudia Spies, Klaus Wernecke (2016). Robust Allocation of Operating Rooms: a Cutting Plane Approach to handle Lognormal Case Durations and Emergency Arrivals. ZIB-Report 16-18.

Guillaume Sagnol, Felix Balzer, Ralf Borndörfer, Claudia Spies, Falk von Dincklage (2016). Makespan and Tardiness in Activity Networks with Lognormal Activity Durations. ZIB-Report 16-23.

Guillaume Sagnol, Ralf Borndörfer, Mickaël Grima, Mathees Seeling, Claudia Spies (2016). Robust Allocation of Operating Rooms with Lognormal case Durations. ZIB-Report 16-16.

Guillaume Sagnol, Hans-Christian Hege, Martin Weiser (2016). Using sparse kernels to design computer experiments with tunable precision. ZIB-Report 16-33.

Adrian Sali (2016). Coupling of Monodomain and Eikonal Models for Cardiac Electrophysiology. ZIB-Report 16-50.

Sebastian Schenker, Ralf Borndörfer, Martin Skutella (2016). A novel partitioning of the set of non-dominated points. ZIB-Report 16-55.

Stephan Schwartz, Ralf Borndörfer, Gerald Bartz (2016). The Graph Segmentation Problem. ZIB-Report 16-53.

Stephan Schwartz, Thomas Schlechte, Elmar Swarat (2016). Designing Inspector Rosters with Optimal Strategies. ZIB-Report 16-65.

Claudia Stötzel, Rainald Ehrig, H. Marike T. Boer, Julia Plöntzke, Susanna Röblitz (2016). Exploration of different wave patterns in a model of the bovine estrous cycle by Fourier analysis. ZIB-Report 16-02.

Tim Sullivan (2016). Well-posed Bayesian inverse problems and heavy-tailed stable Banach space priors. ZIB-Report 16-30.

Alexander Tesch (2016). A Nearly Exact Propagation Algorithm for Energetic Reasoning in $O(n^2 \log n)$. ZIB-Report 16-25.

Alexander Tesch (2016). Exact Energetic Reasoning in $O(n^2 \log^2 n)$. ZIB-Report 16-46.

Alexander Tesch (2016). Improved Compact Models for the Resource-Constrained Project Scheduling Problem. ZIB-Report 16-76.

Christoph von Tycowicz, Felix Ambellan, Anirban Mukhopadhyay, Stefan Zachow (2016). A Riemannian Statistical Shape Model using Differential Coordinates. ZIB-Report 16-69.

Stefan Vigerske, Ambros Gleixner (2016). SCIP: Global Optimization of Mixed-Integer Nonlinear Programs in a Branch-and-Cut Framework. ZIB-Report 16-24.

Martin Weiser, Sunayana Ghosh (2016). Theoretically optimal inexact SDC methods. ZIB-Report 16-52.

Christian Tobias Willenbockel, Christof Schütte (2016). Variational Bayesian Inference and Model Selection for the Stochastic Block Model with Irrelevant Vertices. ZIB-Report 16-01.

Stefanie Winkelmann (2016). Markov Control with Rare State Observation: Average Optimality. ZIB-Report 16-59.

Stefanie Winkelmann, Christof Schütte (2016). The spatiotemporal master equation: approximation of reaction-diffusion dynamics via Markov state modeling. ZIB-Report 16-60.

Jakob Witzig, Timo Berthold, Stefan Heinz (2016). Experiments with Conflict Analysis in Mixed Integer Programming. ZIB-Report 16-63.

Wei Zhang, Carsten Hartmann, Christof Schütte (2016). Effective Dynamics Along Given Reaction Coordinates, and Reaction Rate Theory. ZIB-Report 16-35.

REFERENCES

BIGGER DATA, BETTER HEALTH

[1] M. Sahu, A. Mukhopadhyay, A. Szengel, S. Zachow, Addressing multilabel imbalance problem of Surgical Tool Detection using CNN, International Journal on Computer Assisted Radiology and Surgery (IJCARS) - Spl. Issue IPCAI, 2017.

SHAPE-BASED ANALYSIS – METHODS AND RESULTS

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

[2] Small CG. *The Statistical Theory of Shape*. Springer, 2012

[3] Dryden IL, Mardia KV. *Statistical Shape Analysis: With Applications in R*. John Wiley & Sons, 2016

[4] Srivastava A, Klassen EP, *Functional and Shape Data Analysis*, Springer, 2016

[5] von Tycowicz C, Ambellan F, Mukhopadhyay A, Zachow S. "A Riemannian Statistical Shape Model using Differential Coordinates". ZIB Report 16-69, Zuse Institute Berlin, 2016

[6] Günther A, Lamecker H, Weiser M. "Flexible Shape Matching with Finite Element Based LDDMM." *International Journal of Computer Vision* 105:2 (2013), pp. 128-143

[7] Lamecker H. "Variational and Statistical Shape Modeling for 3-D Geometry Reconstruction". Doctoral thesis, Freie Universität Berlin, Fachbereich Mathematik und Informatik, 2008

[8] Grewe CM, Zachow S. "Fully Automated and Highly Accurate Dense Correspondence for Facial Surfaces." *European Conference on Computer Vision ECCV 2016*, pp. 552-568. Springer, 2016.

[9] Baum D, Mahlow K, Lamecker H, Zachow D, Müller J, Hege HC. "The Potential of Surface-based Geometric Morphometrics for Evolutionary Studies: an Example using Dwarf Snakes (eirenis)". *Digital Specimen 2014*, Sept. 8-12, Berlin, Germany, 2014, pp. 1-2.

[10] Bernard F, Salamanca L, Thunberg J, Tack A, Jentsch A, Lamecker H, Zachow S, Hertel F, Goncalves J, Gemmar P. Shape-aware Surface Reconstruction from Sparse 3-D Point-Clouds. *Medical Image Analysis* 38 (2017), pp. 77-89

[11] Van Kaick O, Zhang H, Hamarneh G, Cohen-Or D. A Survey on Shape Correspondence. *Computer Graphics Forum*, 30:6 (2011), pp. 1681-1707.

[12] Ehlke M, Ramm H, Lamecker H, Hege HC, Zachow, S. "Fast Generation of Virtual X-Ray Images for Reconstruction of 3-D Anatomy." *IEEE Transactions on Visualization and Computer Graphics*, 19:12 (2013), pp. 2673-2682.

[13] Ehlke M, Frenzel T, Ramm H, Shandiz MA, Anglin C & Zachow S. "Towards Robust Measurement of Pelvic Parameters from AP Radiographs Using Articulated 3-D Models." *Int J CARS* 10 (Suppl 1) (2015), pp. 25-26.

[14] Lamecker H, Zachow S, "Statistical Shape Modeling of Musculoskeletal Structures and Its Applications", in: Zheng, G., *Computational Radiology for Orthopaedic Interventions*, pp. 1-23, Springer, 2016.

RESEARCH CAMPUS MODAL: REPORT FROM THE GASLAB

[1] Ying Chen and Bo Li (2015): An Adaptive Functional Autoregressive Forecast Model to Predict Electricity Price Curves, *Journal of Business and Economic Statistics*

[2] Gas (2010): Verordnung über den Zugang von Gasversorgungsnetzen (Gasnetz-zugangsverordnung – GasNZV), 2010. Version released Sep. 03, 2010.

[3] Nina Geißler, Uwe Gotzes, Benjamin Hiller, Jessica Rövekamp, and Thorsten Koch (2015): Regulatory rules for gas markets in Germany and Europe. In *Evaluating Gas Network Capacities* [Thorsten Koch, Benjamin Hiller, Marc Pfetsch, and Lars Schewe, editors]. MOS-SIAM Series on Optimization. SIAM, 2015.

COMPUTING THE TRUTH

[1] Christopher Benzmüller and Brune Wolzenlogel Paleo. The inconsistency in Gödel's ontological argument: A success story for AI in metaphysics. In *Proceedings IJCAI*, pages 936–942, 2016.

[2] Kevin K.H. Cheung, Ambros Gleixner, and Daniel E. Steffy. Verifying integer programming results. ZIB-Report 16-58, Zuse Institute Berlin, November 2016.

[3] William Cook, Thorsten Koch, Daniel E Steffy, and Kati Wolter. An exact rational mixed-integer programming solver. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 104–116. Springer, 2011.

[4] Marijn JH Heule, Oliver Kullmann, and Victor W Marek. Solving and verifying the boolean pythagorean triples problem via cube-and-conquer. arXiv preprint arXiv:1605.00723, 2016.

[5] George Gonthier. Formal Proof – the four-color theorem. *Notices Amer. Math. Soc.* 55 (2008), no. 11, 1383–1393.

[6] Jonad Pulaj, Annie Raymond, and Dirk Theis. New conjectures for union-closed families. ZIB-Report 15-57, Zuse Institute Berlin, December 2015.

[7] Jonad Pulaj. Cutting planes for Families Implying Frankl's Conjecture. ZIB-Report 16-51, Zuse Institute Berlin, November 2016.

[8] Nathan Wetzler, Marijn JH Heule, and Warren A Hunt Jr. Drat-trim: Efficient checking and trimming using expressive clausal proofs. In *International Conference on Theory and Applications of Satisfiability Testing*, pages 422–429. Springer, 2014.

PERFORMANCE IN DATA SCIENCE

[1] Robert Schmidtke, Guido Laubender, Thomas Steinke: Big Data Analytics on Cray XC Series DataWarp using Hadoop, Spark and Flink. CUG Proceedings (2016)

[2] Felix Hupfeld, Toni Cortes, Björn Kolbeck, Jan Stender, Erich Focht, Matthias Hess, Jesus Malo, Jonathan Marti, Eugenio Cesario: The XtreamFS architecture—a case for object-based file systems in Grids (2008)

[3] Alexander Alexandrov, Rico Bergmann, Stephan Ewen, Johann-Christoph Freytag, Fabian Hueske, Arvid Heise, Odej Kao, Marcus Leich, Ulf Leser, Volker Markl, Felix Naumann · Mathias Peters · Astrid Rheinländer, Matthias J. Sax, Sebastian Schelter, Mareike Höger, Kostas Tzoumas, Daniel Warneke: The Stratosphere platform for big data analytics. *VLDB Journal* (2014)


SUPERCOMPUTING AT THE LIMIT

[1] www.lustre.org


[2] www.beegfs.com

[3] TURBOMOLE V6.6 (June 26, 2014), a development of University of Karlsruhe and Forschungszentrum Karlsruhe GmbH, 1989-2007, TURBOMOLE GmbH, since 2007; available from <http://turbomole.com>

[4] www.ixpug.org



NOTES



IMPRINT

Published by

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

Phone +49-30-841 85-0

Fax +49-30-841 85-125

E-mail info@zib.de

www.zib.de

Layout

SCHIEBE PREIL BAYER

werbung und kommunikation – Berlin

ISSN

0934-5892

©2017 Konrad-Zuse-Zentrum für Informationstechnik Berlin

