



18th HPI Research Symposium 2023

Hasso Plattner Institute
Potsdam | April 19 - 21, 2023

Design IT. Create Knowledge.

Agenda

Wednesday, April 19, 2023 International Doctoral Symposium

Room: D-Space (Building H)

- 09:00 – 09:30 **Registration and Welcome Coffee**
- 09:30 – 10:00 **Welcome Note**
- 10:00 – 11:00 **PhD Talks**
 Marcus Pappik, PhD Student, Hasso Plattner Institute, Potsdam
 Discretization-Based Algorithms for Repulsive Gibbs Point Processes
- 11:00 – 11:30 **Invited Talk: "Softwarization in Computer Networks"**
 Prof. Dr. Holger Karl, Internet-Technology and Softwarization, HPI
- 11:30 – 13:00 **Lunch**
- 13:00 – 13:30 **Invited Talk: "Introduction to the KI Service Center"**
 Dr.-Ing. Soner Emec, Administrative Head KISZ
- 13:30 – 14:30 **PhD Talks**
 Shohei Katakura, PhD Student, Hasso Plattner Institute, Potsdam
 Optimizing the Mechanical Fit Between the Parts of Laser Cut Models
- 14:30 – 15:00 **Coffee Break**
- 15:00 – 16:00 **Session on Expanding HPI Collaboration**
- 16:00 – 17:00 **Poster and Ice Cream Session**
- 17:00 – 18:00 **Joint Project Brainstorming**
- 18:00 **Social Dinner**

Thursday, April 20, 2023 Research Symposium

Room: Lecture Hall

13:45 - 14:00 **Registration**

14:00 - 14:15 **Opening Words**

Prof. Dr. Andreas Polze, Head of Operating Systems and Middleware Group and Speaker of the HPI Research School

14:15 - 14:45 **PhD Elevator Pitches**

14:45 - 15:30 **Keynote**

Dr. Alexander Böhm, Principal Engineer, Google
Building Database Management Systems for Modern Enterprise Applications Past, Present, and Future Challenges

15:30 - 16:00 **Coffee Break**

16:00 - 17:20 **Session 1**

Prof. Dr.-Ing. Stefan Schulte, Head of the Institute for Data Engineering, Hamburg University of Technology

A Short History of Service-Oriented Computing

Prof. Dr. Melissa Densmore, University of Cape Town

Jonas Cremerius, PhD Student, Hasso Plattner Institute, Potsdam
Change Patterns in Dynamic Event Attributes

17:20 - 17:50 **Coffee Break**

17:50 - 19:10 **Session 2**

Prof. Dr. Assaf Schuster, Professor of the Computer Science Department Technion, Israel Institute of Technology

Unleashing Affordable Cloud Resources with SpotOS

Dr. Belen Bermejo, Assistant Lecturer & Researcher at the University of the Balearic Islands

Performance and Energy Trade-off in Server Consolidation (STANDARD Project)

Johan Bontes, Department of Computer Science, University of Cape Town

19:10 **Closing**

19:30 **Boat Trip**

Friday, April 21, 2023 Research Symposium

Room: Lecture Hall

- 09:00 - 09:30 **Registration and Welcome Coffee**
- 09:30 - 10:15 **Keynote: "Highlights of HPI: A Retrospective"**
Prof. Dr. Christoph Meinel, Former President and CEO of HPI for Digital Engineering
- 10:15 - 11:30 **Session 3**
Prof. Dr. Ziawasch Abedjan, Leibniz University Hannover
Aikaterini Niklanovits, PhD Student, Hasso Plattner Institute, Potsdam
[Algorithmic Results on Connected Graph Partitioning Problems](#)
Mary-Jane Antia, University of Cape Town
- 11:30 - 12:00 **Coffee Break**
- 12:00 - 13:00 **Session 4**
Dr. Fabio Niephaus, Principal Researcher, Oracle Labs
[Fast and Efficient Java Applications with GraalVM Native Image](#)
Dr. Max Plauth, System Engineer, UltiHash
- 13:00 **Closing**
- 14:30 **Visit to Museum Barberini (Pending Confirmation)**

Hasso Plattner Institute for Digital Engineering

The Hasso Plattner Institute for Digital Engineering in Potsdam is Germany's university excellence center for IT-Systems Engineering. HPI is the only university institution in Germany to offer the Bachelor's and Master's degree in "IT-Systems Engineering" – a practical and engineering-oriented alternative to a conventional computer science study program. Current enrollment is at approximately 500 students. It carries out research noted for its high standard of excellence in its twelve topic areas. Research work is also conducted at the HPI Research School for Ph.D. candidates as well as at its branches in Cape Town, Haifa and Nanjing.

The HPI School of Design Thinking is Europe's first innovation school for university students. It is based on the Stanford d.school model and offers 240 places annually for a supplementary study program. Since 2012 the Hasso Plattner Institute has offered Massive Open Online Courses (MOOCs) at its own interactive platform, openHPI. The courses on IT topics have so far reached more than 170,000 users from over 180 countries. Free via the Internet, learners can access didactically prepared multi-media course materials and learn in close exchange with other course participants through social media.



HPI Research Symposium

The HPI Research Symposium is the annual symposium of the HPI Research Schools. It outlines new trends in the area of Service-oriented Computing and Data Science and Engineering, featuring invited speakers and talks by selected research school members.

The HPI Research Symposium: Connecting Industry and Academia. As the HPI Research Schools are an interdisciplinary effort of the HPI research groups, the HPI Research Symposium covers a wide range of topics related to service-oriented computing, data science and data engineering. Excellent speakers - both from industry and academia - leaders in their respective field of research, are invited to talk about their latest projects and resulting outcomes.



The HPI Research Schools

Research School

Data Science and Engineering

The increasing abundance of data in science and in industry creates many challenges and opportunities. Data science has grown to be a foundational discipline in information technology, allowing new insights from data and creating ever more intelligent applications. Simultaneously, it is becoming increasingly difficult to collect, clean and deliver the vast amounts of data and apply and maintain complex data science processes. Targeting these challenges, the discipline of data engineering has become equally foundational.

The 2019 newly established research school "Data Science and Engineering" unites top PhD students in all areas of data-driven research and technology, including scalable storage, stream processing, data cleaning, machine learning and deep learning, text processing, data visualization and more. We apply our research to many different use cases across the participating interdisciplinary research groups, joining forces whenever possible.

Research School

Service-Oriented Systems Engineering

In October 2005, the HPI started its Research School on "Service-Oriented Systems Engineering", a graduate school based on the model of the DFG (German Research Foundation) "Graduiertenkolleg". The Vision of the Research School Design and implementation of service-oriented architectures impose numerous research questions from the fields of software engineering, system analysis and modeling, adaptability, and application integration. Service-Oriented Systems Engineering represents a symbiosis of best practices in object orientation, component-based development, distributed computing, and business process management. It provides integration of business and IT concerns. Service-Oriented Systems Engineering denotes a current research topic in the field of IT-Systems Engineering with high potential in academic research as well as in industrial application. Supported by an internationally renowned grant, Ph.D. students at our college participate in joint activities such as lectures, seminars, winter schools and workshops.

The Members of the Research School

The professors of the HPI with their research groups are supporting pillars for our Ph.D. school. With its interdisciplinary structure, the research college on Service-Oriented Systems Engineering interconnects the HPI research groups and fosters close and fruitful collaborations. In context of the Research School, the different groups at HPI work on the following topics:

Enterprise Platform and Integration Concepts, Prof. Dr. h.c. Hasso Plattner
Algorithm Engineering, Prof. Dr. Tobias Friedrich
Artificial Intelligence and Intelligent Systems, Prof. Dr. Gerard de Melo
Artificial Intelligence and Sustainability, Prof. Dr. Ralf Herbrich
Business Process Technology, Prof. Dr. Mathias Weske
Computergraphics Systems, Prof. Dr. Jürgen Döllner
Cybersecurity - Enterprise Security, Prof. Dr. Christian Dörr
Cybersecurity - Identity Management, Prof. Dr. Anja Lehmann
Data Analytics and Computational Statistics, Prof. Dr. Bernhard Renard
Data Engineering Systems, Prof. Dr. Tilmann Rabl
Design Thinking and Innovation Research, Prof. Dr. Falk Uebernickel
Digital Health - Connected Healthcare, Prof. Dr. Bert Arnrich
Digital Health - Machine Learning, Prof. Dr. Christoph Lippert
Human Computer Interaction, Prof. Dr. Patrick Baudisch
Information Systems, Prof. Dr. Felix Naumann
Internet-Technologien and Softwarization, Prof. Dr. Holger Karl
Operating Systems and Middleware, Prof. Dr. Andreas Polze
Software Architecture, Prof. Dr. Robert Hirschfeld
System Engineering and Modeling, Prof. Dr. Holger Giese

On the website of the Research School, please find latest information about the Ph.D. students, their research interests, joint projects, and events:

<https://hpi.de/en/research/cooperations-partners/research-schools.html>

The HPI International Research Schools

HPI Research School at Nanjing University (NJU)

Nov 17 2011 marked the inauguration of the "HPI Research School at Nanjing University". HPI director Prof. Christoph Meinel welcomed the first five Chinese doctoral students and their supervising professors. The Ph.D. students that originally hail from the computer science department raised the total head count of the HPI Research School to 50. Together, the highly talented young scientists from all across the globe tackle various issues in the field of service-oriented IT systems. The main research areas of the Chinese Ph.D. students are "Mass Data Analytics and Knowledge Discovery". They will do their work in the former capitol of China, Nanjing, which is situated about 300 kilometers northwest of Shanghai. Since the HPI Research School in Nanjing began its work in 2011, 18 scholarship holders have successfully completed their PhD theses.

Prof. Dr. Song Fangmin is head of the research school at Nanjing University.

HPI Research School at Technion - Israel Institute of Technology

At Israel's technical university in Haifa, 6 students are currently doing research at the "HPI Research School at Technion". Its goal is to allow a multidisciplinary group of researchers in the Computer Science and Electrical Engineering departments to rise up to the new challenges presented by the shift to Chip Multiprocessor Computing (CMP). The students are mentored jointly by professors from the Computer Science and Electrical Engineering departments. Since HPI-Technion began its work in 2010, 30 scholarship holders have successfully completed their PhD theses.

Prof. Dr. Idit Keidar is head of the research school at Technion.

HPI Research School at University of Cape Town (UCT)

The Hasso Plattner Institute Research School at UCT, in Information and Communications Technology for Development (ICT4D), serves as a gateway for educating sub-Saharan African information technologists, and in the grand scheme, to become a leading centre of excellence in designing and understanding computational solutions (both algorithmic and user centric) and systems to address the UN Sustainable Development Goals. To this end, HPI-UCT shall promote and encourage: The education of the brightest PhD students from sub-Saharan Africa. The aim is to offer attractive funding to this pool of candidates for whom few other funding channels are available for study at South African universities. Development of strong research collaborations with others in the Hasso-Plattner Institute (HPI) as a leading world establishment in information technology. There are currently

16 postgraduates from different African countries working at UCT. Since HPI Research School at UCT began its work in 2009, 18 scholarship holders have successfully completed their PhD theses.

Prof. Dr. Melissa Densmore is head of the research school at UCT.

HPI Research School at University of California (UCI)

The Hasso Plattner Institute (HPI) is very pleased to announce the opening of a new Research School branch for doctoral students at the University of California, Irvine (UCI), well known for its academic achievements and high quality education. The new "HPI Research Center in Machine Learning and Data Science at UC Irvine" aims to promote research and educational activities in these two fields between the two leading universities. Through this partnership, 10 UCI scholarship holders are currently jointly supervised by 8 professors from UCI's Donald Bren School of Information and Computer Sciences while closely being integrated into HPI's research activities.

Prof. Dr. Erik Sudderth is head of the research school at UCI.

International Doctoral Symposium

Over the years the HPI Research School has been expanded to a state in which we are excited to – in addition to our members from Hasso Plattner Institute Potsdam, Germany – welcome colleagues from University of Cape Town (UCT), South Africa; the Technion, Israel; and Nanjing University, China. In order to foster collaboration and information exchange, we organized a Doctoral Symposium, so that PhD students from HPI, UCT, Technion, and NJU will get to know each other both on the social and the academic levels. The event includes a mix of a social event, poster session, and demo presentations.

The poster session that dives into the research topics of the attending PhD students, will follow a more informal and virtual event in Gather Town for exploring other people's posters and explaining yours!



Research School Chairs

Speakers of the HPI Research School "Service-Oriented Systems Engineering"

Andreas Polze is the Operating Systems and Middleware Professor at the Hasso Plattner Institute at University Potsdam, Germany. He is also the speaker of the HPI Research School and member of the steering committee of HPI's Future SOC Lab. Andreas received a doctoral degree from Freie University Berlin, Germany, in 1994 and a habilitation degree from Humboldt University Berlin in 2001, both in Computer Science. At HPI, his research focuses on architectures of operating systems, on component-based middleware, as well as on predictable distributed and cloud computing. Andreas Polze was visiting scientist with the Dynamic Systems Unit at Software Engineering Institute, at Carnegie Mellon University, Pittsburgh, USA, where he worked on real-time computing on standard middleware (CORBA) and with the Real-Time Systems Laboratory at University of Illinois, Urbana-Champaign. Current research interests include Predictable Service Computing, Adaptive System Configuration, and End-to-End Service Availability for standard middleware platforms. Together with Charité, GETEMED, and Deutsche Telekom, he has run the Fontane telemedicine project. Joint research with SAP has investigated porting HANA to new processor architectures.

Robert Hirschfeld leads the Software Architecture Group at the Hasso Plattner Institute at the University of Potsdam. His research interests include dynamic programming languages, development tools, and runtime environments to make live, exploratory programming more approachable. Hirschfeld received a PhD in computer science from Technische Universität Ilmenau.



Prof. Dr. Andreas Polze

Head of Operating Systems
and Middleware Group



**Prof. Dr.
Robert Hirschfeld**

Head of Software
Architecture Group

Research School Chairs



**Prof. Dr.
Felix Naumann**

Head of Information Systems
Group

Speakers of the HPI Research School "Data Science and Engineering"

Felix Naumann studied mathematics, economy, and computer sciences at the University of Technology in Berlin. He completed his PhD thesis on "Quality-driven Query Answering" in 2000. In 2001 and 2002 he worked at the IBM Almaden Research Center on topics of data integration. From 2003 - 2006 he was assistant professor for information integration at the Humboldt-University of Berlin. Since then he holds the chair for information systems at the Hasso Plattner Institute at the University of Potsdam in Germany. He is editor-in-chief of the Information Systems journal. His research interests are in data profiling, data cleansing, and text mining.



Prof. Dr. Tilmann Rabl

Head of Data Engineering
Systems Group

Tilmann Rabl holds the chair for Data Engineering Systems at the Hasso Plattner Institute and is Professor at the Digital Engineering Faculty of the University of Potsdam. He is also cofounder and scientific director of the startup bankmark. Tilmann Rabl received his PhD at the University of Passau in 2011. He spent 4 years at the University of Toronto as a postdoc in the Middleware Systems Research Group (MSRG). From 2015 to 2019, he was senior researcher and visiting professor at the Database Systems and Information Management (DIMA) group at Technische Universität Berlin and Vice Director of the Intelligent Analytics for Massive Data (IAM) Group at the German Research Center for Artificial Intelligence (DFKI). His research interests include data processing on modern hardware, stream processing, benchmarking, and end-to-end machine learning systems.

Speakers

Building Database Management Systems for Modern Enterprise Applications Past, Present, and Future Challenges

With hundreds of CPU cores and DRAM capacities in the multiple terabytes, recent developments in modern hardware led to a new class of database management systems. These operational analytics systems can process demanding transactional (OLTP) and analytical (OLAP) workloads on a single copy of data. They provide tremendous benefits and simplification potential for applications and end-users, as various data management tasks can be consolidated in the context of a single system. For researchers and industry practitioners alike, the combination of OLTP and OLAP workloads raises additional questions how to best engineer such database systems. In this talk, we'll discuss various challenges with respect to workload management, scheduling, and non-uniform memory architectures (NUMA) that have a significant impact on the performance and scalability of operational analytics systems. We also dig into cost-related aspects such as memory footprint reduction by smart data placement on different layers of the storage hierarchy (DRAM, SSD), or more efficient query plan operator scheduling. We conclude the talk with an outlook on current engineering and research challenges in the context of modern, cloud-based enterprise database systems, i.e. the popular pattern of separating compute and storage layers, or challenges and opportunities of using machine learning to further improve these systems.



Dr. Alexander Böhm
Principal Engineer, Google

Dr. Alexander Böhm is a principal engineer working on Google's cloud database portfolio. His current focus is on AlloyDB for PostgreSQL, an open-source compatible, cloud-native database management system for demanding enterprise workloads. Prior to joining Google in 2022, he was a development executive at SAP, and one of the chief architects working on SAP's in-memory database management system, HANA. There, his focus was on performance optimization and holistic improvements of enterprise architectures, in particular application server/DBMS co-design. Alexander received his PhD from the University of Mannheim, Germany in 2010, where he worked on the development of efficient and scalable applications using declarative message processing.



**Prof. Dr.-Ing.
Stefan Schulte**

Head of the Institute for Data
Engineering, Hamburg
University of Technology

A Short History of Service-Oriented Computing

These days, Everything-as-a-Service (XaaS) is a major design principle in both distributed computing and software engineering, allowing to provide pretty much "every" functionality as a remote service. The roots of service-oriented computing (SOC) can be traced back at least to the origins of Remote Procedure Calls (RPCs) in the 1970s, and different evolutionary steps have led to the current state-of-the-art. Within this talk, we will give a brief history of the most important technical developments that led to XaaS, including RPCs, Web Services, REST, and the Cloud. In addition, we will have a brief look at ongoing SOC topics currently investigated at the Institute for Data Engineering at Hamburg University of Technology.

Dr.-Ing. Stefan Schulte is Full Professor at Hamburg University of Technology, head of the Institute for Data Engineering, and leads the Christian Doppler Laboratory Blockchain Technologies for the Internet of Things (CDL-BOT). He received a diploma degree in economics and a Bachelor in Computer Science from the University of Oldenburg, Germany, and a Master of Information Technology (with Merit) from the University of Newcastle, New South Wales, in 2005 and 2006, respectively. He is a member of the Editorial Board of Journal of Parallel and Distributed Computing, IEEE Transactions on Network and Service Management, and Computing. In 2023, he is the PC Co-Chair of the IEEE International Conference on Fog and Edge Computing (ICFEC) and General Chair of the IEEE International Conference on Cloud Engineering (IC2E). He is also an elected member of the Steering Board of the "Communication and Distributed Systems" Special Interest Group of the German Informatics Society.

Unleashing Affordable Cloud Resources with SpotOS

Community wisdom tells us that while the cloud early vision promised economy of scale, in practice, renting cloud resources is an expensive business. There are many reasons for this unfortunate state of affairs, some of which can be attributed to non-utilized opportunities. SpotOS is an effort which is aimed to use state of the art tools to help cloud users make better rental choices. In particular, SpotOS enables wider, easier, and better use of spot instances, which are “empty” resources sold in high discount with the pitfall that they may have to be evacuated in short notice. SpotOS can be viewed as a new kind of a distributed “operating system” which extends the OS traditional tasks by optimizing task allocation to resources, optimizing batch resource rental from the provider, and supporting failover mechanisms. To achieve these goals SpotOS uses advanced modern tools such as optimization algorithms, predictive machine learning, efficient external storage, and stateful task management in distributed environments. The talk will survey the challenges and opportunities in SpotOS.



**Prof. Dr.
Assaf Schuster**

Professor of the Computer
Science Department Technion,
Israel Institute of Technology

Prof. Dr. Assaf Schuster of the Computer Science Department Technion is a world leading expert of distributed and scalable data Mining, Big Data technologies analytics & prediction, Cyber security and system vulnerabilities, privacy preserving, cloud resource management and more. Prof. Schuster published more than 200 papers in highly selective conferences and journals, some of which won prestigious awards. He consulted leading hi-tech companies, such as IBM, HP, Microsoft, and Verint. He participated in the bumpy journey of quite a few startups, some of which were successful. His algorithms on data-race detection were implemented in Intel's Thread Checker. His papers triggered a rewrite of the Java Memory Model. His patents were sold by the Technion. He has built scalable production systems to handle petabytes of storage with off-the-shelf hardware. His research group is well known in its contributions to the field of big data and scalable, real-time knowledge discovery in distributed data streams. The annual funding of his group from national, international, and industrial sources, reaches hundreds of thousands of Euro. The group consists of several post-doctoral researchers and quite a few graduate students, as well as promising undergraduate students. Prof. Schuster established CSL - the Computer Systems Laboratory, and TCE - the Technion Center for Computer Engineering, which he headed until recently.



Dr. Belen Bermejo

Assistant Lecturer &
Researcher at the University
of the Balearic Islands

Performance and Energy Trade-off in Server Consolidation (STANDARD Project)

With the large and continually increasing data capacity demands placed on data centers worldwide, efficient use of data center energy is extremely important. Moreover, datacenters represent from 2% to 6% of worldwide emissions of greenhouse gases. Servers have poor power proportionality and the reduction in power consumption in datacenters can be achieved by lowering the servers performance, however, if the users services or applications consume more time to be performed, the servers will consume more energy. On the one hand, the ISO/IEC 30134 standard specifies datacenter energy effectiveness KPIs to help datacenter operators measure and improve specific aspects of datacenter energy effectiveness. However, it does not provide a method for comparing individual server energy effectiveness across datacenters. On the other hand, the ISO/IEC 21836 standard provides a Server Energy Effectiveness Metric (SEEM) to measure and report the energy effectiveness of specific server configurations. The main goal of the STANDARD project is the scientific development of KPIs to support future standardization which quantifies the trade-off between performance degradation and energy consumption of server consolidation for datacenters. These indicators would allow data centers operators to know the efficiency of a certain configuration of virtual machines and/or containers per physical server. In addition, the indicators will allow making comparisons between different physical and consolidation techniques to determine the most suitable consolidation option in a scenario.

Dr. Belen Bermejo received the PhD degree in Computer Science from the University of the Balearic Islands in 2020. She is currently assistant Lecturer at the mentioned university and also researcher. Her research is focused on the energy efficiency and performance of virtualization and cloud system's. She has authored several indexed journals and international conferences.

Fast and Efficient Java Applications with GraalVM Native Image

In this talk, we present GraalVM Native Image and show how it can turn Java applications into native executables. Such executables start up and run fast, have a small memory footprint, and provide other advantages over running on a Java virtual machine. We explain how Native Image works, why it is more than just AOT compilation, what limitations it imposes, and how it is used today and could be used in the future.



Dr. Fabio Niephaus

Principal Researcher, Oracle
Labs

Fabio Niephaus is a researcher on the GraalVM project at Oracle Labs. His research interests include developer tools, programming languages, and virtual machines. Before joining Oracle Labs and as part of his Ph.D. research at HPI, he created TruffleSqueak, a Squeak/Smalltalk VM and Polyglot Programming Environment for the GraalVM.

Speakers



Jonas Cremerius

PhD Student, Hasso Plattner
Institute, Potsdam

Change Patterns in Dynamic Event Attributes

Process mining uses data recorded during process execution to obtain an understanding of the actual execution of a process in different application domains. The healthcare domain is especially interesting, as it comes with unique characteristics and challenges. Processes in healthcare are complex and produce a high amount of data, such as laboratory measurements. Many of these values change during a typical process quite frequently. Hence, we refer to those as dynamic event attributes. My research deals with the analysis of dynamic event attributes and focusses on the detection of change patterns in processes. This allows to relate changes in the actual values to conducted activities, such that a change in laboratory measurements can be observed for a given treatment. In this talk, I give an overview of my research project in which I analyze change patterns in dynamic event attributes.

Jonas Cremerius is a PhD student at the Business Process Technology research group at Hasso Plattner Institute, University of Potsdam. His research focuses on process mining methods and techniques in healthcare. While conventional methods focus on activities and their sequence, he incorporates domain-specific data, such as laboratory measurements of patients. This makes it possible to analyze how the patient's well-being develops during the treatment process.

Discretization-Based Algorithms for Repulsive Gibbs Point Processes

Gibbs point processes are a popular way to model particle distributions of fluids and gasses in Euclidean space. Sampling their Gibbs distribution and computing their normalizing constant, the partition function, is highly relevant for studying such point processes. However, until recently, very few rigorous computational results existed. In this talk, I give a brief introduction to Gibbs point processes and present recent algorithmic results. In particular, I focus on discretization-based algorithms, which reduce the algorithmic tasks at hand to related problems for discrete spin systems, making use of the rich literature in that area. My main focus will be on a recent approach that employs hard-core models on carefully constructed families of geometric random graphs to obtain sampling and approximation algorithms for Gibbs point processes. This results in efficient algorithms for arbitrary repulsive pair potentials ϕ up to a fugacity of $\lambda < e/C_\phi$, where C_ϕ is the temperedness constant of ϕ .



Marcus Pappik

PhD Student, Hasso Plattner
Institute, Potsdam

Marcus Pappik is a PhD student at the Algorithm Engineering research group at Hasso Plattner Institute, University of Potsdam. His research interests include a variety of topics such as data science, computational statistics (especially causal inference), game theory and stochastic processes. He is currently focusing on the application of Markov chains to discrete and continuous systems from statistical physics and the resulting algorithmic applications.



Aikaterini Niklanovits

PhD Student, Hasso Plattner
Institute, Potsdam

Algorithmic Results on Connected Graph Partitioning Problems

Partitioning a graph into connected components is a fundamental task in graph algorithms, and those connected partitions occur as desirable structures in areas such as image processing, network decomposition and robotics. Moreover partitioning a graph into components that are not only connected but also maximize the overall density of the partition is one of the most efficient ways to model community structures. The first problem can be found under many names depending on the size constraints also applied to each connected component of the partition, or on the number of components, while the second one is called dense partition. In this talk I am going to introduce these problems and present some algorithmic and hardness results on computing such partitions.

Aikaterini Niklanovits is a PhD student at the Algorithm Engineering research group at Hasso Plattner Institute, University of Potsdam. Her research focuses on developing efficient algorithms for various partitioning problems restricted to specific graph classes. In particular, Aikaterini focuses on studying such problems on classes defined through the absence of certain induced subgraphs.

Optimizing the Mechanical Fit Between the Parts of Laser Cut Models

Since laser-cutting is inherently limited to producing flat parts, users produce 3D models with the help of so-called "box joints" cut into those plates. If properly designed, such box joints hold the model together without glue, but just by the tightness of the joints. Producing just the right tightness is challenging though, as even tiny variations in the dimensions of the parts can cause joints to fall apart or make them impossible to assemble.

In this talk, I present a series of projects aiming to eliminate sources of such tiny variations, including (1) making the laser cutter itself more precise (by improving "repeatability") (2) Helping users calibrate the laser cutter, and specifically the amount of material it burns of during cutting (also known as "kerf"), (3) adjusting the tightness according to the specifics of the respective joint, and (4) according to assembly order.

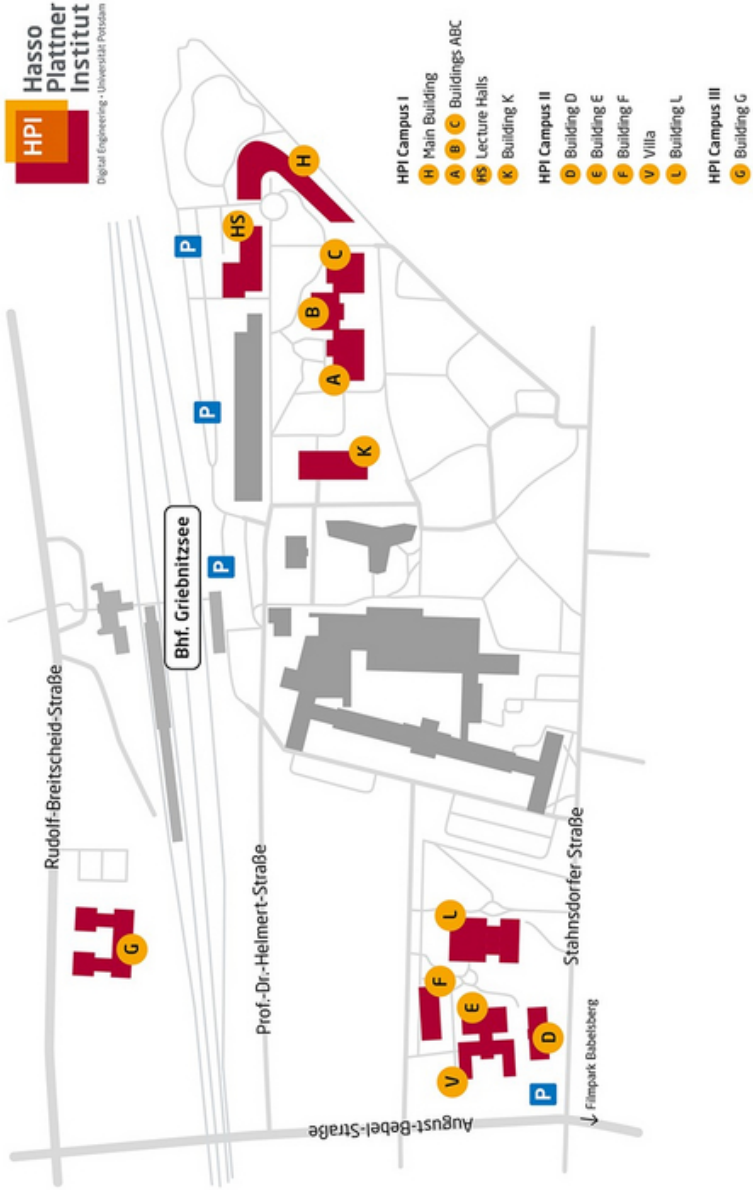


Shohei Katakura


PhD Student, Hasso Plattner
Institute, Potsdam


Shohei Katakura is a PhD student at the Human Computer Interaction research group at Hasso Plattner Institute, University of Potsdam. His research interest is in human-computer interaction (HCI), especially personal fabrication. He designs a relationship and interaction with manufacturing machines by developing software and hardware systems. His research goal is to enable non-experts to create interactive and moving physical objects using 3D printing technologies.

Maps



Notes





Contact

Hasso-Plattner-Institut für Digital Engineering gGmbH
Campus Griebnitzsee | University of Potsdam

Prof.-Dr.-Helmert-Straße 2-3
D-14482 Potsdam, Germany

Phone +49 331 5509-0
E-Mail hpi-info@hpi.de
Web <https://www.hpi.de>