13th Annual Symposium on Future Trends in Service-Oriented Computing

Hasso Plattner Institute
Potsdam | April 18 – 20, 2018

Design IT. Create Knowledge.

HPI Hasso Plattner Institut
Digital Engineering · Universität Potsdam

https://hpi.de/future-soc-symposium

Agenda

Wednesday, April 18, 2018

Room: H.E-51

9:30 - 10:15	**Keynote**
	Dr. Maria Keet, Senior Lecturer, University of Cape Town, South Africa More Effective Ontology Authoring with Test-Driven Development
10:15 - 10:30	**Coffee Break**
10:30 - 12:00	**Session DS-1**
	Mina Rezaei, PhD Student, Hasso Plattner Institute, Potsdam Survival-GAN: An Adversarial Framework for Multiple Clinical Tasks
	Vladeta Stojanovic, PhD Student, Hasso Plattner Institute, Potsdam Interactive Visualization for Facility Management 4.0
	Anesu Marufu, PhD Student, University of Cape Town, South Africa Exploring Malware in Resource-Constrained Auction-Augmented Microgrids
	Du Dapeng, PhD student, Nanjing University, China Depth Images Could Tell Us More: Enhancing Depth Discriminability for RGB-D Scene Recognition
	Oren Kalinsky, PhD Student, Technion, Haifa, Israel Efficiency of Join Queries through Algorithms and Hardware
	Lukas Pirl, PhD Student, Hasso Plattner Institute, Potsdam Software Fault Injection for the Internet of Things
12:00 - 13:00	**Lunch Break**
13:00 - 13:30	**Ice Cream & Poster Session**
13:30 - 15:00	**Session DS-2**
	Roman Kaplan, PhD Student, Technion, Haifa, Israel Processing in Memory Architecture for Machine Learning and Bioinformatics
	Leonid Azriel, PhD Student, Technion, Haifa, Israel Using Scan Side Channel to Detect IP Theft
	Hao Rui, PhD Student, Nanjing University, China Recommendation and Clustering System for Crowdsourced Bug Reports
	Selvas Mwanza, PhD Student, University of Cape Town, South Africa Algorithmically Detecting and Tracking Socio-Political Problems in South Africa Using Social Media Data
	Augustine Takyi, PhD Student, University of Cape Town, South Africa Detection of Malicious Nodes using Collaborative Neighbour Monitoring in DSA Networks
	Thomas Brand, PhD Student, Hasso Plattner Institute, Potsdam Focus with Adaptive Monitoring
15:00	**Social Event: Guided Tour through Sanssouci Park and Dinner**

Thursday, April 19, 2018

Room: HS1 (Lecture Building)

12:00 – 12:10 **Opening of the Symposium on Future Trends in Service-Oriented Computing**

Prof. Dr. Christoph Meinel, Prof. Dr. Andreas Polze, HPI, Potsdam

12:10 – 12:30 **Elevator Pitches**

HPI Research School "Service-Oriented Systems Engineering", Potsdam
PhD Students Introduce Their Work

12:30 – 13:00 **Keynote**

Divesh Srivastava, PhD, AT&T Labs-Research, USA
The Confounding Problem of Private Data Release

13:00 – 13:20 **Coffee Break**

13:20 – 14:20 **Session 1: Algorithmics**

Dr. Heiner Ackermann, Fraunhofer Institute for Industrial Mathematics ITWM, Department of Optimization
Coping with Hard Problems in Practice – An Efficient Algorithm for a Packing Problem in Sawmills

Yoshiki Ohshima, PhD, CEO Vision, Inc.
Shadama: A Massively-Parallel Particle Simulation Environment for Everyone

Prof. Assaf Schuster, Technion, Haifa, Israel
Recent Advances in Complex Event Processing

14:20 – 14:30 **Group Photo**

14:30 – 14:50 **Coffee Break**

14:50 – 16:10 **Session 2: Analytics**

Elena Poughia, Managing Director of Dataconomy Media GmbH and Founder & Head Curator of Data Natives
Ex Machina: Artificial Creativity

Qiao Yu, PhD Student, Nanjing University, China
False-name Proof Mechanism for Path Auction in Social Networks

Dr. Jan Schaffner, VP, Head of SAP Innovation Center Network
Scala-based Cloud Application Development Kit for SAP Cloud Platform

Prof. Dr. Martin Monperrus, Ph.D., Professor of Software Technology at KTH Royal Institute of Technology and Chair of the Wallenberg Autonomous Systems and Software Program - WASP
The Future of Automated Program Repair

16:10 – 16:30 **Coffee Break**

16:30 – 17:30 **Session 3: Systems & Threats**

Prof. Dr.-Ing. Jörg Nolte, BTU Cottbus
The Insane Memory Hierarchy – or Why It Makes Sometimes Sense to Do Algorithmically Useless Work To Improve the Performance

Werner Haas, CTO, Cyberus Technology GmbH, Dresden
Meltdown/Spectre: The Dark Side of the Microprocessor

Sven Köhler, PhD Student, Hasso Plattner Institute, Potsdam
Tools for Adaptive Work-Package Choice in Middleware

18:30 **Social Event: Boat Trip to Glienicke Palace and Dinner**

Friday, April 20, 2018

Room: HS1 (Lecture Building)

09:30 - 10:00	**Keynote**
	Jens Mönig, SAP Research Reinventing on Purpose
10:00 - 10:20	**Coffee Break**
10:20 - 11:20	**Session 4: Future in Computing**
	Alexa Gorman, Global VP of SAP.iO Fund & Foundry in Europe at SAP SAP.iO: Driving Corporate Innovation with Startups
	Dr. Luise Pufahl, Research Assistant, Hasso Plattner Institute, Potsdam Modeling and Executing Batch Activities in Business Processes
	Christiane Bauer, SAP SAP Young Thinkers – Inspire and Prepare the Next Generation
11:20 - 11:40	**Coffee Break**
11:40 - 13:00	**Session 5: From CS to Digital Health**
	Prof. Dr.-Ing. Wolfgang Schröder-Preikschat, FAU Erlangen-Nürnberg Adaptive Memory Protection for Many-Core Systems
	Prof. Christopher Longhurst, CIO, UCSD Health, Clinical Professor of Biomedical Informatics and Pediatrics, UCSD School of Medicine Trends in Health Informatics
	Wanjiru Mburu, PhD Student, University of Cape Town, South Africa Needs Assessment Findings for a Neonatal Intensive Care Unit Communication Tool
	Prof. Lawrence Friedman, Professor of Medicine and Associate Dean for Clinical Affairs, UCSD Health and UCSD School of Medicine The Digital Revolution and Transformation of the US Health Care System
13:00	**Closing Remarks**

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.

The Symposium on Future Trends in Service-Oriented Computing

The Symposium on Future Trends in Service-Oriented Computing 2018 is the annual symposium of the HPI Research School and is taking place for the twelfth time. It outlines new trends in the area of Service-Oriented Computing and highlights recent work of select Research School members.

As the HPI Research School is an interdisciplinary undertaking of the HPI research groups, the Symposium on Future Trends in Service-Oriented Computing covers a wide range of topics concerning SOC, which include but are not limited to: cloud computing, software, platform, infrastructure as a service, service description, discovery and composition, service deployment, platform configuration and capacity planning, monitoring, service middleware, service-oriented architectures (SOAs), service management, information as a service, service development and maintenance, novel business models for SOAs, economical implications of web services and SOAs, service science, mobile and peer-to-peer services, data services, quality of service, exception handling, or service reliability and security.

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 School

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:

Human Computer Interaction and Computer Vision as Service (Prof. Dr. Patrick Baudisch), Service-Oriented Geovisualization Systems (Prof. Dr. Jürgen Döllner), Algorithm Engineering for Service-Oriented Systems (Prof. Dr. Tobias Friedrich), Modeling and Verification of Self-Adaptive Service-Oriented Systems (Prof. Dr. Holger Giese), Tools and Methods for Software Engineering in Service-Oriented Systems (Prof. Dr. Robert Hirschfeld), Security Engineering of Service-Based IT-Systems (Prof. Dr. Christoph Meinel), Service-Oriented Knowledge Discovery and Data Mining (Prof. Dr. Emmanuel Müller), Service-Oriented Information Systems (Prof. Dr. Felix Naumann), Evolutionary Transition of Enterprise Applications to Service-Orientation (Prof. Dr. h.c. Hasso Plattner), Operating System Abstractions for Service-Oriented Computing (Prof. Dr. Andreas Polze); and Services Specification, Composition, and Enactment (Prof. Dr. Mathias Weske).

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/research-school

The 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 Insititute 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 Ph.D. 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 Ph.D. students, will follow a more informal pattern: eating ice cream while exploring other people posters and explaining yours!

Social Event on Wednesday: Guided Tour through Sanssouci Park and Dinner

We will have a guided tour (in English) through the New Palace ("Neues Palais") and Sanssouci Park in Potsdam. The park, with its palaces and park structures, may be counted among the most significant park ensembles in Europe and is a UNESCO World Heritage site. The name Sanssouci – carefree – should be understood as the desire and leitmotif of the king, because it was here that he most gladly retreated. After the walk, we will have dinner together.

Logistics: We depart all together from the main Symposium location at Potsdam Griebnitzsee by train to the visitor center "Neues Palais", where we meet our guides. Please bring shoes that your are comfortable walking in.

The HPI Future SoC Lab

The HPI Future SOC (Service-Oriented Computing) Lab is a cooperation of the Hasso Plattner Institute and the industrial partners EMC, Fujitsu, SAP and Hewlett Packard Enterprise. Its mission is to enable and promote exchange and interaction between the research community and the industrial partners.

The Lab provides researchers with free of charge access to a complete infrastructure of state of the art hard- and software. This infrastructure includes components, which might be too expensive for an ordinary research environment. The offerings address researchers particularly from but not limited to the areas of computer science and business information systems. Main areas of research include cloud computing, parallelization, and In-Memory technologies.

Future SOC Lab Day – Spring 2018

On Tuesday, April 17, 2018, the projects of the previous Future SOC Lab period got a chance to present the results of their research activities. Additionally, selected requesters of new projects can expand their ideas. You are welcome to hand in a project proposal and to apply for the usage of the lab's resources. More information: https://hpi.de/future-soc-lab

Speakers

Opening of the Symposium on Future Trends in Service-Oriented Computing

Prof. Dr.
Christoph Meinel

Dean of the
Digital Engineering Faculty
at the University of Potsdam

Christoph Meinel (Univ. Prof., Dr. sc. nat., Dr. rer. nat., *1954) is CEO and Scientific Director of the Hasso Plattner Institute for Digital Engineering gGmbH (HPI) and Dean of the Digital Engineering Faculty at the University of Potsdam. Christoph Meinel is full professor (C4) for computer science at HPI and the University of Potsdam, and he holds the chair of Internet Technologies and Systems. He teaches courses on IT Systems Engineering in the HPI Bachelor and Master Degree programs and in the MOOC platform developed by his team: openHPI. He supervises numerous Ph.D. projects and is a teacher at the HPI School of Design Thinking. His research currently focuses on security engineering, knowledge engineering, and Web 3.0-Semantic, Social, Service Web. He is also scientifically active in research on the innovation method Design Thinking. Earlier scientific work concentrated on efficient algorithms and complexity theory. Christoph Meinel is author or co-author of more than 25 books, anthologies, as well as numerous conference proceedings. He has had more than 550 (peer-reviewed) papers published in scientific journals and at international conferences and holds a number of international patents. He is a member of the National Academy of Science and Engineering (acatech), director of the HPI-Stanford Design Thinking Research Program, honorary professor at the TU Beijing, visiting professor at Shanghai University, concurrent professor at the University of Nanjing, and member of numerous scientific committees and supervisory boards.

Prof. Dr.
Andreas Polze

Vice Dean and Speaker of the
HPI Research School

Prof. Dr. 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.

Speakers

Dr. **Maria Keet**

Senior Lecturer, University of Cape Town, South Africa

More Effective Ontology Authoring with Test-Driven Development

Ontology authoring is a complex process, where commonly the automated reasoner is invoked for verification of newly introduced changes, therewith amounting to a time-consuming test-last approach. Test-Driven Development (TDD) for ontology authoring is a recent test-first approach that aims to reduce authoring time and increase authoring efficiency. Current TDD testing falls short on coverage of OWL features and possible test outcomes, the rigorous foundation thereof, and evaluations to ascertain its effectiveness. We aim to address these issues. We propose a succinct logic-based model of TDD testing and novel TDD algorithms that cover also any OWL 2 class expression for the TBox and for the principal ABox assertions. The algorithms use methods from the OWL API directly such that reclassification is not necessary for test execution, therewith reducing ontology authoring time. The algorithms were implemented in TDDOnto2, a Protégé plugin. TDDOnto2 was evaluated on editing efficiency and by users. The editing efficiency of TDDonto2 is faster than a typical ontology authoring interface, especially for medium size and large ontologies. Modellers make significantly less errors with TDDOnto2 compared to the standard Protégé interface and complete their tasks better using less time. Thus, the results indicate that Test-Driven Development is a promising approach as an ontology development methodology.

Dr. C. Maria Keet is a Senior Lecturer with the Department of Computer Science, University of Cape Town. She focuses on knowledge engineering with ontologies, conceptual data modeling, and natural language generation. Before her employment at UCT, Maria worked at the School of Computer Science, UKZN, South Africa, and at the KRDB Research Centre, Free University of Bozen-Bolzano, Italy. She obtained a PhD in Computer Science at the KRDB Research Centre in 2008, following a BSc(hons) 1st class in IT & Computing (OU UK), and 3.5 in the IT industry. She also holds an MSc in Food Science (WUR, the Netherlands) and an MA 1st class in Peace & Development Studies (UL, Ireland).

Survival-GAN: An Adversarial Framework for Multiple Clinical Tasks

Mina Rezaei

PhD Student, Hasso Plattner Institute, Potsdam

Inspired by the recent success of generative adversarial networks (GANs), we propose a novel adversarial network called Survival-GAN for multiple clinical tasks. We introduce an automatic conditional generative adversarial network (cGAN) for medical image semantic segmentation and prediction of the patient overall survival (OS). The Survival-GAN comprises two components a generator and a discriminator. The generator is trained on sequential brain magnetic resonance images (MRI) and clinical data to learn the segmentation label's map and overall survival, while the discriminator is trained to discriminate a generator output, coming from the ground truth or from the generator network. The generator network and the discriminator network simultaneously train via adversarial loss, with the well-known strategy of two-player mini-max games. The proposed architecture shows notable results on the BraTS-2017 benchmark where the achieved accuracy in term of Dice for segmentation whole tumor region is 0.89, and 0.82 for overall survival days prediction.

Mina is a Ph.D. candidate at the Internet Technologies and Systems chair of Prof. Meinel at Hasso Plattner Institute in Potsdam and focuses on deep learning for medical image analysis. Prior to that, she received her Master's degree at artificial intelligence at Shiraz university in Iran and Bachelor's degree in software engineering at Arak university in Iran. She worked six years as software developer at Statistics Center of Iran and was an intern at Technical University Munich.

Vladeta Stojanovic

PhD Student, Hasso Plattner Institute, Potsdam

Interactive Visualization for Facility Management 4.0

The ability to capture and visualize Building Information Modelling (BIM) data is becoming increasingly important in the field of Facility Management (FM). Facility managers can make use of as-is BIM data in order to create and enhance planning and operational documentation for a building throughout its lifetime. Generation of as-is BIM data from point cloud data poses a particular challenge. The visualization-based analytical output of combined as-designed and as-is BIM, point cloud and sensor data based on a Generalized Building Information Model (GBIM), allows for representation of a Digital Twin for lifecycle management topics for stakeholder engagement within the emerging Real Estate 4.0 realm. The use of service-based interactive visualization extends the outputs of the visualization scenario to thin clients. This could be potentially beneficial for increasing stakeholder engagement by allowing complex visualization results to be transmitted to clients in real-time on mobile and portable computer devices and accessed through a simple web-based portal. This research investigates the development of new methods and techniques for development of semantically rich 3D point cloud representations for use as basis data for Digital Twins along with combined as-is BIM and sensor analytics visualization.

Vladeta Stojanovic is a second year PhD student working with the Computer Graphics Systems group under the supervision of Prof. Dr. Jürgen Döllner, at the Hasso Plattner Institute in Potsdam, Germany. He originally comes from a computer games and real-time 3D graphics programming background. The topic of his PhD research is the application of interactive 3D visualization for Facility Management applications. He also has a strong interest in real-time interactive 3D visualisation technology in order to help solve sustainability assessment, land use planning, and energy modelling and construction problems in the built environment. His other professional and hobbyist interests include 3D graphics programming, machine learning and web-based 3D visualization.

Exploring Malware in Resource-Constrained Auction-Augmented Microgrids

Anesu Marufu

PhD Student, University of Cape Town, South Africa

The principle of Continuous Double Auctioning (CDA) is known to provide an efficient way of matching supply and demand among distributed selsh participants with limited information. However, the literature indicates that the classic CDA algorithms developed for grid-like applications are centralised and insensitive to the processing resources capacity, which poses a hindrance for their application on resource constrained smart micro-grids (RCSMG). In this research, we design and evaluate a CDA algorithm for power allocation in a RCSMG. First, we extend the original CDA scheme to enable decentralised auctioning. Second, to handle CDA performance issues due to malfunctioning devices on an unreliable network (such as a lossy network), we extend our proposed CDA scheme to tolerate failure. In addition, we propose a decentralised consumption scheduling scheme that complements the auctioning scheme in guaranteeing successful power allocation within the RCSMG. Thirdly, we propose threat models centred on cheating attacks aimed at foiling the extended CDA scheme. Fourth, in order to mitigate these cheating attacks, we propose an exception handling (EH) scheme.

Anesu Marufu is a PhD fellow at the University of Cape Town within the Computer Science department. He is currently part of the Information Security Group working in the cyber-physical infrastructure security space. This research is part of the SANCOOP conjoint project between NRF (South Africa) and Norwegian Research board. His Research interests are in Distributed Systems, Double Auctions, Auction Security, Multi-agent Systems, Attack Modelling.

Du Dapeng

PhD student, Nanjing University, China

Depth Images Could Tell Us More: Enhancing Depth Discriminability for RGB-D Scene Recognition

Recently depth-modal information has been witnessed effectively in computer vision community, especially for scene analysis related tasks. However, it still suffers severely from depth data scarcity as well as improperly transferring pre-trained RGB models to fit depth-modal data. In this study, we propose a novel two-step training strategy to address these problems and focus on enhancing the recognition power for depth-modal images in RGB-D scene recognition task. Specifically, we build an effective "Res-U" architecture on a GAN (generative adversarial networks) based RGB-to-depth modality translation model, which is endowed with both short and long skips for residual learning. On one hand, this could first well pre-train a depth-modal-specific discriminator network from scratch in an unsupervised manner, which is effectively transformed for the subsequent recognition task instead of directly fitting pre-trained RGB model to depth-specific one. On the other hand, new depth images with helpful perturbations, generated from the modality translation model, help argument the original training set and regularize the learning process in some sense. This two-step training strategy makes it more effective for training a modal-specific network to discriminate depth scenes. Besides, we extensively explore the modality translation network to investigate the effects in recognizing depth-modal scenes, which encourages a reasonable way to take full advantage of multi-modalities. The proposed method achieves state-of-the-art accuracy on NYU Depth v2 and SUN RGB-D benchmark datasets, especially on depth data only evaluation.

DU Dapeng joined the Department of Computer Science and Technology of the Nanjing University as a PhD student in 2015 and became a member of the MCG Group, led by Professor Gang-Shan Wu. His PhD research topic is about visual multimedia processing & analysis and understanding as well as computer vision, particularly in indoor/outdoor place recognition, location based estimation (localization) in the scene. Before that, he had been working as a programmer at FNST Inc. participating in developing several Java-based enterprise-level projects for customers at home and abroad.

Efficiency of Join Queries through Algorithms and Hardware

Oren Kalinsky

PhD Student, Technion,
Haifa, Israel

A core problem in data analysis is that of evaluating complex join queries over large datasets. Such queries are highly expensive in terms of time, memory, and energy. We attack this challenge by exploring the connection between join algorithms and hardware. I will present our join algorithm based on the recent concept of a "worst-case optimal" join. By combining theoretical and system perspectives, the algorithm can utilize any available memory without compromising the theoretical guarantees. We achieve a speedup of orders of magnitude over state-of-the-art algorithms and systems on various benchmark tasks. I will show an extension of our approach in a system for exploratory queries over datasets in the Semantic Web, a growing field that employs complex joins (e.g., SPARQL). Finally, I will present our efforts in designing a hardware accelerator for database and graph queries. Our accelerator is based on the above algorithm, specifically its hardware friendliness, and can achieve further speedups and energy efficiency.

Oren Kalinsky is a Ph.D. candidate of Computer Engineering at the Technion, Haifa, Israel. He is supervised by Prof. Yoav Etsion, specializing in Computer Systems, and Prof. Benny Kimelfeld, specializing in Database Theory. Oren received his B.Sc. and M.Sc. from the Technion. He worked at Intel for three years, primarily in chip design. Later, he worked as a Navy Lieutenant at the Israeli Navy for six years, managing system architecture and software infrastructure teams.

Software Fault Injection for the Internet of Things

Lukas Pirl

PhD Student, Hasso Plattner Institute, Potsdam

In the context of the Internet of Things (IoT), unprecedented circumstances narrow the applicability of established best practices for fault tolerance. Foremost, IoT devices face notable resource constraints regarding computation, storage and communication. Further, they largely operate in uncontrolled environments. On a higher level, the IoT is a widely distributed system, where ensuring dependability is notoriously difficult. Additionally, the economic pressure for short time-to-market cycles hampers the maturing of the systems. Overall, this suggests that best practices need to be re-evaluated and alternatives might need to be reconsidered.

The concept of Software Fault Injection (SFI) offers promising prospects to experimentally assess the dependability of complex systems. Thus, SFI can aid in understanding systems' dependability properties and improving their dependability. Since SFI is not unfolded to its full potential yet, we demonstrate its versatility and promote its incorporation into development practices.

Differing from other SFI approaches (e.g., *Fuzzing*), we promote structured SFI assessments. Therefore, detailed fault models (i.e., *What can go wrong?*) and dependability models (i.e., *What is done about it?*) need to be developed. The SFI investigations are then based on the foregoing models. A case study is conducted in the publicly funded project *RailX*, where Vehicle-to-Vehicle applications for railway services are developed.

Lukas is a PhD candidate at the Hasso Plattner Institute's Research School for Service-Oriented Systems Engineering since 2017. In the same year, he earned his master's degree with his thesis on a structured approach to software fault injection, which was exemplified by a case study on an infrastructure as a service platform. Since then, his research focuses on the assessment of distributed systems' dependability through software fault injection. Specifically, Lukas is currently interested in dependability aspects in the Internet of Things, where uncontrolled environments and manifold resource constraints pose unprecedented challenges.

Processing in Memory Architecture for Machine Learning and Bioinformatics

Roman Kaplan

PhD Student, Technion, Haifa, Israel

User-generated and bioinformatics database volumes has been increasing exponentially for more than a decade. With the slowdown and approaching end of Moore's law, traditional technologies cannot satisfy the increasing demands for processing power. In this talk, I will present PRINS, a highly-parallel in-storage processing architecture. PRINS combines non-volatile memory with processing capabilities on every bitcell. An emerging technology, memristors, form the basis for the design and will also be discussed. In addition, I will present our implementation of algorithms from bioinformatics (DNA local sequence alignment) and machine learning (K-means, K-nearest neighbors), and compare PRINS performance and energy efficiency to that of CPU, GPU, FPGA and the Automata-processor.

Roman received his MSc from the faculty of Electrical Engineering at Technion in 2015, after spending 5 years as a software engineer. He is now a PhD candidate in the same faculty under the supervision of Prof. Ran Ginosar. Roman's research interests are parallel computer architectures, in-data processing, accelerators for machine learning, bioinformatics and computational biology.

Leonid Azriel

PhD Student, Technion,
Haifa, Israel

Using Scan Side Channel to Detect IP Theft

In the growing heterogeneous IoT market, which embraces a plurality of vendors and service providers, IP protection plays a central role. Our work proposes a process for detection of IP theft in VLSI devices that exploits the internal test scan chains, designed for production test automation. The scan chains supply direct access to the internal registers in the device, enabling combinational analysis of the device logic. By using Boolean function learning methods, the learner creates a partial dependency graph of the internal flip-flops. The graph is further partitioned using the Shared Nearest Neighbors graph clustering method, and individual blocks of combinational logic are isolated. These blocks can be matched with known building blocks that compose the original function. This enables reconstruction of the function implementation to the level of pipeline structure. The IP owner can compare the resulting structure with his own implementation to confirm whether an IP violation has occurred. We demonstrate the power of the presented approach with a test case of an open source Bitcoin SHA-256 accelerator, containing more than 80,000 registers. With the presented method, we discover the microarchitecture of the module, locate all the main components of the SHA-256 algorithm, and learn the module's flow control.

Leonid Azriel received his BSc and MSc from the Technion-Israel Institute of Technology. Currently, he is pursuing a PhD also in the Technion-Israel Institute of Technology. All the degrees are in Electrical Engineering. Prior to beginning his PhD studies, he served in different technical and managerial positions at National Semiconductor, Winbond Electronics and Nuvoton Technologies companies, where he was involved in the development of the Trusted Platform Module (TPM).

Recommendation and Clustering System for Crowdsourced Bug Reports

Hao Rui

PhD Student, Nanjing University, China

With the popularity of mobile phones, the application market has been expanding in recent years and the app publish cycle is becoming shorter and shorter to survive in the fierce competitive environment. It has become a new challenge to ensure the quality of applications under frequent fast iterations. Thus, crowdsourced testing gains recent attention in the software-engineering community against this background. Generally, workers on crowdsourced platforms are requested to submit bug reports which consist of screenshots and textual description relevant with the bugs they have encountered. Considering the large number of crowdsourced workers, the bug reports accumulate vastly and bug reports inspection and triaging turn into quite challenging works. Thus, we propose several methods to help discerning and prioritizing useful information from bug reports. First, semi-supervised clustering and topic model are utilized to assist developers in quickly understanding the overall state of submitted reports, then a bug report prioritization technique picks out the high-quality reports for developers. Moreover, we also apply the recommendation to exploring the breadth or depth of the dataset. These approaches have been adopted by real industry leaders and the feedback is very effective.

Rui Hao is a Ph.D. student in software engineering at Nanjing University, supervised by Prof. Zhenyu Chen. Her research interests lie in software maintenance, data mining, and software testing (especially on mobile testing and crowdsourcing testing).

Selvas Mwanza

PhD Student, University of Cape Town, South Africa

Algorithmically Detecting and Tracking Socio-Political Problems in South Africa Using Social Media Data

Social surveys have been used by researchers and policy makers as an essential tool for understanding social and political activities in society. Social media has introduced a new way of capturing data from large numbers of people. Unlike surveys, social media deliver data more rapidly and cheaply. Research has shown that social media has become pivotal in shaping the political discourse around the world. In this talk, I will discuss how we are using social media data to detect social and political issues in South Africa.

Selvas Mwanza obtained his Bachelor of Computer Science at the University of Zambia in 2010, a Masters of Engineering Software Engineering at Beijing University of Posts and Telecommunication in China in 2013 and is currently doing a PhD in Computer Science at University of Cape Town in South Africa. His research interests include Text Mining and Natural Language Processing.

Detection of Malicious Nodes using Collaborative Neighbour Monitoring in DSA Networks

Augustine Takyi

PhD Student, University of Cape Town, South Africa

Dynamic spectrum access (DSA) based networks are forecast to provide internet access to rural and remote areas. That includes areas with little or no access to network infrastructure. However, the deployment of such networks is susceptible to several kinds of attack. Securing the networks while exploiting the flexibilities offered by DSA-based networks remains a daunting challenge. This work addresses this challenge by introducing a system model of a DSA-based network with a fusion centre, modelling the possible threat of malicious nodes, and presenting a novel physical layer malicious node detection strategy. Our detection strategy uses collaborative neighbour monitoring by the secondary nodes within the deployment area to detect malicious nodes. Each node monitors its neighbours by sensing received power values and forwarding this information to a fusion centre for analysis. The simulation results obtained show that collaborative neighbour monitoring in hybrid (centralized and distributed) networks works well in detecting position falsification attacks in the dynamic spectrum access networks, provided that the distance between the actual malicious node position and the falsified position is at least 0.035km.

Augustine Takyi is a PhD student at the Department of Computer Science in University of Cape Town, South Africa and a Lecturer from University of Energy and Natural Resources in Ghana. He has varying teaching and research experience in areas ranging from Networks and Security, programming, Database Development, etc. He holds a BSc Computer Science & Statistics (combined major) from University of Ghana, Ghana and MEng Communication and information System from Huazhong University of Science and Technology, China. His research interest is in securing computer networks for effective communication.

Thomas Brand

PhD Student, Hasso Plattner Institute, Potsdam

Focus with Adaptive Monitoring

Advances in information technology allow capturing and processing vast and further increasing amounts of data. While technical constraints vanish other aspects gain weight in the decision making about data. After ethical aspects, economical and ecological aspects become eventually significant, when deciding how much and which data to capture and process.

The economical and ecological aspects are addressed by focusing on valuable data. But the value of data can change over time. Large scales, high speeds and frequencies of changes suggest automating the finer grained adaptation of data capturing and processing mechanisms.

In this talk our current work on generic adaptive monitoring is presented. It deals with capturing and processing data about software systems. The monitoring configuration is automatically adjusted during the system runtime. Adaptations can be triggered due to changing data demands but also by changes in the behavior or structure of the monitored system.

Existing work on adaptive monitoring is specific to particular monitoring purposes, such as service level agreement monitoring or anomaly detection. In order to reduce the effort of applying adaptive monitoring for other purposes the presented approach aims to be more generic.

Thomas is a Ph.D. student at the System Analysis and Modeling chair of Prof. Holger Giese at Hasso Plattner Institute in Potsdam. He is interested in how software changes over time to adapt to new requirements. Thus he is investigating adaptive monitoring and how it can be used to comprehend the usage of enterprise software.

Speakers

PhD Students Introduce Their Work: Elevator Pitches

HPI Research School "Service-Oriented Systems Engineering", Potsdam

Ahmed Shams:	Improving MOOC in Poor Network Performance
Andreas Fricke:	Geo-Servicification
Ankit Chauhan:	Analysis of Algorithms Using Structural Properties of Real-World Networks
Anton Titsulin:	Graph Representation Learning
Christian Adriano:	Predicting tasks in human in the loop mechanisms
Christoph Matthies:	Data-Driven Software Development Processes
Christiane Hagedorn:	Fostering Collaboration with Gameful Learning in Scalable E-Learning Environments
Erik Scharwächter:	Event Impacts on Time Series
Fabio Niephaus:	Polyglot Programming
Francesco Quinzan:	Evolutionary Algorithms in Combinatorial Optimization
Julian Risch:	News Comment Analysis - Deep Learning for News Rooms
Lan Jiang:	Data Preparation Exploration
Martin Krejca:	Understanding Swarm Algorithms
Ralf Rothenberger:	Power-Law Distributions in Random Satisfiability
Robert Kovacs:	TrussFormer: 3D printing large kinematic structures
Sebastian Marwecki:	Scenograph
Sona Ghahremani:	Utility-Driven Architecture-based Self-adaptive systems
Stefan Ramson:	Active Expressions - A Basic Building Block for Reactive Programming
Thijs Roumen:	Kyub: a 3D Editor for Modeling Sturdy Laser-Cut Objects
Tobias Bleifuß:	Change Exploration - A New Dimension of Data Analytics
Toni Mattis:	Learning Semantic Concepts from Source Code

The HPI Research School "Service-Oriented Systems Engineering" focuses on the design and implementation of service-based architectures. These architectures target various research problems from software engineering, system modeling and analysis, as well as the adaptability and integration of applications. It represents a synthesis of best practices in object orientation, component-based development, distributed computing, and business process management. The topics of the HPI Research School thus have a unique potential not only in terms of academic research but also in industrial application.

Divesh Srivastava, PhD

AT&T Labs-Research, USA

The Confounding Problem of Private Data Release

In our Big Data era, as data-driven decision making sweeps through all aspects of society, the demands to make useful data available are growing ever louder. For example, the ubiquity of GPS-enabled devices has resulted in a wealth of data about the movements of individuals and populations, which can be analyzed for useful information to aid in city and traffic planning, disaster preparedness, and so on. But the problem of releasing such data without disclosing confidential information, such as the places people visit, is a subtle and difficult one. Is "private data release" an oxymoron? This talk will delve into the motivations of private data release, explore the challenges, and outline some of the historical and recent approaches developed in response to this confounding problem.

Divesh Srivastava is the head of Database Research at AT&T Labs-Research. He is a Fellow of the Association for Computing Machinery (ACM), the Vice President of the VLDB Endowment, and the managing editor of the Proceedings of the VLDB Endowment (PVLDB). He has served as an associate editor of the ACM Transactions on Database Systems (TODS), and as an associate Editor-in-Chief of the IEEE Transactions on Knowledge and Data Engineering (TKDE). He has presented keynote talks at several international conferences, and his research interests and publications span a variety of topics in data management. He received his Ph.D. from the University of Wisconsin, Madison, USA, and his Bachelor of Technology from the Indian Institute of Technology, Bombay, India.

Coping with Hard Problems in Practice – An Efficient Algorithm for a Packing Problem in Sawmills

Dr. **Heiner Ackermann**

Fraunhofer Institute for Industrial Mathematics ITWM, Department of Optimization

From a theoretical perspective, many real-world planning problems are computationally intractable because of NP-hard subproblems. Nonetheless, we need to solve them quickly. The key question then is: How can this be achieved? There is certainly no universal algorithm, but there is a powerful pattern: Divide the problem into tractable subproblems and combine their solutions. In this talk, I will apply this pattern to an example that arises in sawmills in the production of glued laminated timber. The problem is to set up pressing operations with minimum waste (packing problem) and minimum setup times (sequencing problem). In practice, our approach is fast and significantly outperforms planning experts.

Dr. Heiner Ackermann studied Computer Science at TU Dortmund University and ETH Zurich. Thereafter, he was a doctoral candidate in the Algorithms and Complexity Group at RWTH Aachen headed by Berthold Vöcking. Since 2009, he is a research associate in the Department of Optimization at the Fraunhofer Institute for Industrial Mathematics ITWM in Kaiserslautern. At Fraunhofer ITWM, he analyzes and optimizes production processes in various kinds of industries, designs packing and scheduling algorithms, and implements decision support systems.

Yoshiki Ohshima, PhD

CEO Vision, Inc.

Shadama: A Massively-Parallel Particle Simulation Environment for Everyone

We present a prototype of a programming system called Shadama. Shadama is designed for writing programs that create, control and visualize large numbers of objects. The basic execution model follows the tradition of StarLogo and its "turtles and patches" abstraction. This abstraction has been proven to be effective and easy to use. The primary goal of the language is to facilitate the writing of scientific simulations by students at the high school level.

The Shadama environment supports liveness. Once changes to the program are saved, the effect on the running simulation is immediate; there is no need to restart the simulation. The language supports both 2D and 3D environments.

Shadama programs are run on the GPU by means of code translation to the OpenGL Shading Language. The OpenGL Shading Language Program data resides entirely on the GPU, which enables high performance.

We are experimenting with the ideas for other kinds of end-users, including company executives who need to deal with massive amount of data and make sense of them. We are hoping to provide some solution for such a problem domain as well.

Yoshiki Ohshima has research interests in interactive and educational computer systems, software architectures and programming languages. Yoshiki graduated from the Tokyo Institute of Technology in 1994. He was awarded his PhD for the creation of "Kedama", a massively parallel particle programming system, from the Tokyo Institute of Technology in 2006. Yoshiki worked on theme park related research project at Walt Disney Imagineering R&D from 2000 through 2002. In 2002, he joined Twin Sun, Inc. From 2007 he worked at the Viewpoints Research Institute. At VPRI, Yoshiki worked on projects include bringing the etoys environment to children via the "XO" and the One Laptop Per Child (olpc) initiative, multilingualization of software environments, and the STEPS project. From 2013 He was a principal investigator at SAP CDG Labs and Y Combinator Research until 2017.

Recent Advances in Complex Event Processing

Prof. **Assaf Schuster**

Technion, Haifa, Israel

Following some real-world trends, the computer sciences are also following with a slow process of focus changes. Instead of the one-time offline algorithms operating on stationary data, we now see a surge of interest in online processing of streaming data. Unfortunately, the algorithms and technologies for online stream processing are lagging behind and the gap can be estimated in years of research. In this talk I will focus on Complex Event Processing (CEP), an algorithmic field whose goal is to find interesting patterns in rapid streams of events (or data). CEP is already an industry of billions of dollars and is growing very fast. I will show that state-of-the-art algorithmic methods for CEP are fairly primitive as compared with other algorithmic fields in computer science. I will then provide an overview of recent methods we developed in an effort to bridge the gap.

Prof. Assaf Schuster of the Computer Science Department at the Technion is an ACM Fellow and an IEEE Fellow and 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. He 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. In 2016, he co-founded CY-OT, a startup company to cyber-protect organizations in the era of the Internet of Things.

Elena Poughia

Managing Director of Dataconomy Media GmbH and Founder & Head Curator of Data Natives

Ex Machina: Artificial Creativity

The talk will provide an intro to computational creativity and its goal to model, simulate or replicate creativity using a computer, to achieve one of several ends: -to construct a program or computer capable of human-level creativity -to better understand human creativity and to formulate an algorithmic perspective on creative behavior in humans. We will examine the connection between artificial intelligence and art from a sociological, philosophical and psychological perspective through real life examples in music, arts, movies, beauty and writing.

Elena Poughia is Managing Director of Dataconomy Media GmbH and the Founder & Head Curator of Data Natives, Europe's largest data science conference. Started her career in arts at established galleries such as Gagosian Gallery and pioneering institutions such as the Venice Biennale and Athens Biennale. She has co-founded an Art Events Company (2012) and an Independent Art Publication (2014) before moving to work for Dataconomy, the #1 Media Platform for the data-driven generation.

False-name Proof Mechanism for Path Auction in Social Networks

Qiao Yu

PhD Student, Nanjing University, China

We study path auction mechanisms for buying a path between two given nodes in a social network. In social network environments, the mechanism is vulnerable to falsename manipulations where agents can profit from placing multiple bids under fictitious names. In this demo, we present core-selecting path mechanisms that are robust against false-name bids and address the overpayment problem. Specifically, we provide a new formulation for the core, which greatly reduces the number of core constraints. Based on the new formulation, we present a Vickery-nearest pricing rule, which finds the core payment profile that minimizes the distance to the VCG payment profile.

Currently Yu Qiao is a second year Ph.D. student of the Department of Computer Science and Technology at Nanjing University and a member of the IIP Group led by Prof. Chongjun Wang. He received his M.Sc. degree in computer science in June 2016 from Xi'an Jiaotong University, Xi'an China. In September 2016, he was admitted to study for a Ph.D. degree at Nanjing University.

Dr. **Jan Schaffner**

VP, Head of SAP Innovation Center Network

Scala-based Cloud Application Development Kit for SAP Cloud Platform

SAP has an impressive track record in efficiently developing applications based on ABAP for the on-premise world. With the new challenges induced by the cloud development, the question about efficiency and maintainability gets a new spin.

The goal of the cloud application development kit for Scala is to take the first steps towards a comparable efficiency for developing cloud services and applications for the SAP Cloud Platform. The kit is based on the modern programming language Scala, leveraging its strong type system and powerful metaprogramming capabilities to offer easy-to-use components for the application developers. The core components of the kit are a simplified access to different (SAP-)data sources via unified Domain Specific Languages (DSLs) and a Scala wrapper for SAP's default user interface components (SAP UI5). That reduces a lot of the boilerplate code for the application development teams. Paired with the conciseness and improved readability of functional programming code, this leads to a situation where an application developer can rapidly write maintainable code.

Jan Schaffner heads the SAP Innovation Center Network (ICN) – a unit that explores and develops meaningful new technologies to fuel transformative growth at SAP. The ICN drives collaborations with leading research institutes as well as market exploration with key customers and co-innovation partners. Jan reports to SAP's Chief Innovation Officer Juergen Mueller. Prior to his current role, Jan led the Office of the Chief Technology Officer (CTO). He and his team worked on strategic projects, evaluated how cutting-edge technologies can be implemented at SAP, and advised on architectural questions. Jan has vast experience in building and leading development teams and deep expertise in database systems and distributed systems. After joining SAP in 2006, Jan worked at the Hasso Plattner Institute for Software Systems Engineering in Potsdam at Prof. Dr. Hasso Plattner's research chair. He was one of the founding members of the prominent joint research project with SAP, which built the first prototype of SAP's award winning in-memory platform SAP HANA. Since 2013, Jan acted in various roles at SAP, such as the technology advisor of the company's CTO, or as the CTO for SAP's product offerings in finance. Jan graduated from the University of Potsdam and holds a PhD degree in Computer Science from the Hasso Plattner Institute. He did parts of this work at the University of California in Berkeley. Jan lives in Berlin is an enthusiastic cyclist and has a passion for coffee.

The Future of Automated Program Repair

Prof. Dr.
Martin Monperrus,
Ph.D.

Professor of Software Technology at KTH Royal Institute of Technology and Chair of the Wallenberg Autonomous Systems and Software Program - WASP

Automatic program repair is the process of fixing software bugs automatically and it is a very active research area. Is it a crazy academic idea or something that can work in practice? What are the main achievements and the main challenges? This talks presents the state-of-the-art and the main research problems in automated program repair.

Martin Monperrus is Professor of Software Technology at KTH Royal Institute of Technology, Sweden. In 2011-2017, he was associate professor at the University of Lille, France and adjunct researcher at Inria. He received a Ph.D. from the University of Rennes, and a Master's degree from the Compiègne University of Technology (UTC). His research lies in the field of software engineering with a current focus on automatic program repair, self-healing software and chaos engineering.

Prof. Dr.-Ing. **Jörg Nolte**

BTU Cottbus

The Insane Memory Hierarchy – or Why It Makes Sometimes Sense to Do Algorithmically Useless Work To Improve the Performance

Modern CPUs hide memory access latencies by means of sophisticated memory hierarchies in combination with other techniques. Three layers of internal CPU-caches are common these days and good cache optimisation is mandatory to achieve acceptable performance. Thus, data access must be organized in a way, that the cache hit rate is maximised. In this talk we will discuss several aspects of the memory hierarchy and show its impact on a Barnes-Hut particle simulation. We will show that algorithmically useless data rearrangements will pay of significantly.

Jörg Nolte is professor for computer science at the Brandenburg University of Technology (BTU) in Cottbus (Germany) where he holds the chair for distributed systems and operating systems. Prior to that position he was a senior researcher at the Fraunhofer Gesellschaft, Institute for Computer Architecture and Software Technology (FIRST), Berlin. He received his M.S. (Dipl.Inform.) in computer science in 1988 and his Ph.D. (Dr.-Ing.) in 1994, both from the Technical University of Berlin. He was a principal member and finally the deputy head of the PEACE group that developed the operating system for Germany's first massively parallel supercomputer. In the 90s he was a post doc fellow and senior researcher in the Tsukuba Research Center (TRC) of the Real World Computing Partnership (RWCP) in Tsukuba Science City, Japan. Since that time his research concentrated on scalable, low-latency middleware and operating system platforms for clusters and other parallel architectures, including rather strange ones such as wireless sensor networks. He is the member of the board of the special interest group for operating systems of the German GI and is currently the dean of the Faculty 1 (Mathematics, Computer Science, Physics, Electrical Engineering and Information Technology) of the BTU. His major research interests are operating systems, middleware and programming languages for parallel, distributed and embedded systems.

Meltdown/Spectre: The Dark Side of the Microprocessor

Werner Haas

CTO, Cyberus Technology GmbH, Dresden

Meltdown and Spectre, these are the names given to microprocessor vulnerabilities disclosed at the beginning of the year that have dominated security discussions ever since. And the names speak for themselves: presumed rock-solid, hardware-anchored protection mechanisms simply melt away and Spectre spreads shock and fear among security experts as a key principle of operation opens the door for malicious goals. Ever increasing processor performance has almost become taken for granted. But have we built the corresponding complex, digital systems on a foundation of sand? This talk will look behind the scenes of performance numbers and illustrate the architectural features that make such attack vectors possible. How can it be that microprocessors regularly perform speculative execution when their results are usually considered the golden reference for unbiased, highly precise operation? Besides the plain technical aspects, the talk will also provide some broader context around the disclosure of the vulnerabilities: how they were discovered, what happened behind the scenes, and which lessons should be drawn from the event.

Werner Haas is co-founder of Cyberus Technology and is responsible for its longer term technology strategy. Previously, he worked for the FireEye cyber-security company and spent over 10 years at Intel, mainly in its research division. His main interests centered around the memory-related aspects of processor architecture where he contributed, for example, to the recent Protection Key extension. Before joining Intel in 2004, he was a scientific assistant at the Institute for Computer Aided Circuit Design of the University of Erlangen-Nuremberg, where he received his Dipl.-Ing. degree in electrical engineering in 1997. Together with Thomas Prescher, he disclosed the Meltdown vulnerability to Intel.

Sven Köhler

PhD Student, Hasso Plattner Institute, Potsdam

Tools for Adaptive Work-Package Choice in Middleware

Hardware performance counters are a popular means to evaluate software behavior in regards of, for instance, consumed time, issued instructions, and failed branch prediction.

I investigate how measurement jitters—usually avoided in experiments—can provide detailed information on the overall system load and properties of currently running processes on the same hardware as my probe program. At the example of memory access times on an IBM POWER8 processor, I present a set of probing tools. These tools enable not only a better understanding of how black-box processes' resource usage changes over time, but allow further the construction of middlewares that can dynamically start or defer annotated work packages based on the current cache utilization.

As a third application, I show how cache timing delays can be employed for the construction of cross-process covert channels. They allow for breaking the isolation of kernel process groups and potentially logical partitions (LPARs).

Sven is a Ph.D. student at the Operating Systems and Middleware group of Prof. Dr. Andreas Polze at the Hasso Plattner Institute (Potsdam, Germany). His work focusses on accelerators for parallel data processing and cryptography on heterogenous platforms. In his spare time he mentors a CoderDojo and introduces school kids to the marvels of computing.

Reinventing on Purpose

Jens Mönig
SAP Research

A concept currently discussed in many novice programming environments is that of custom microworlds. I present work-in-progress around this idea in the Snap! programming language.

Snap! lets educators remove blocks from the palette. Often the reason for hiding blocks is to let students focus on a specific task without getting distracted by an overwhelming multitude of features not relevant in this situation. A variation of this is to entice students into inventing their own versions of previously removed blocks, which can open up discussions around programming language design and user interface choices. I will hide some key control structures from Snap! and try to reinvent them "live". I will also demonstrate how we apply the idea of reinventing on purpose to designing a new way to interact with hardware such as the BBC micro:bit and the Calliope circuit board.

Snap! Is a Scratch-like programming language that treats code-blocks as first class citizens instead of confining them to an editing modality. Embracing nested data structures and higher order functions Snap! lets learners create arbitrary control structures and even custom programming languages with just blocks. Snap! has been developed for UC Berkeley's new introductory computer science course named "The Beauty and Joy of Computing".

Jens Mönig is a researcher at SAP and makes interactive programming environments. He is fanatical about visual coding blocks. Jens is the architect and lead programmer, together with Brian Harvey, of UC Berkeley's "Snap! Build Your Own Blocks" programming language, used in the introductory "Beauty and Joy of Computing" curriculum. Previously Jens has worked under Alan Kay on the GP programming language together with John Maloney and Yoshiki Ohshima, helped develop Scratch for the MIT Media Lab and written enterprise software at MioSoft. Jens is a fully qualified lawyer in Germany and has been an attorney, corporate counsel and lecturer for many years before rediscovering his love for programming through Scratch and Squeak. For leisure Jens likes guitar picking and strumming his mandolin.

Alexa Gorman

Global VP of SAP.iO Fund & Foundry in Europe at SAP

SAP.iO: Driving Corporate Innovation with Startups

In this presentation, you will hear how SAP is driving open innovation with startups through the SAP.iO Fund and Foundries. Learn more about how the SAP.iO Fund was set up to invest in early stage startups that have complementary solutions to SAP's portfolio and could benefit from API integrations with SAP solutions as well as access to SAP's customers. An overview of the investment strategy and criteria will be given and how these investments benefit both SAP's customers and SAP itself. In addition, the presentation will dive into the different acceleration programs SAP offers for startups, what the benefit are and how to approach SAP.iO if an investment or acceleration program are of interest.

Alexa is Global VP of the SAP.iO Fund & Foundry in Europe and responsible for early stage startup investments as well as acceleration and incubation. In this role, Alexa leads the Berlin and Paris SAP.iO Foundries. She has 18+ years of experience in business development and strategy. The SAP.iO Fund invests in early stage startups that can leverage SAP APIs, data sets, business content and/or platform technologies. The SAP.iO Foundry is a global network of accelerators providing deep, programmatic mentorships to startups as well as access to SAP customers.

Modeling and Executing Batch Activities in Business Processes

Dr. **Luise Pufahl**

Research Assistant, Hasso Plattner Institute, Potsdam

Business process automation improves the efficiency of organizations to perform their work. For this, the business processes are first documented as a process model which then serve as blueprint for a number of process instances representing the execution of business cases. In existing BPM languages and systems, process instances are assumed to run independently from each other. However, batch processing – the collectively execution of several instances at specific process activities – is a common phenomenon in operational processes to reduce cost or time. Currently, batch processing is organized manually or hard-coded in software. In this talk, a new process modeling element – the batch activity - is presented supporting the specification of batch work in process models and their automatic execution. Its benefits will be demonstrated in this talk and its integration in an existing BPM system.

Luise Pufahl is research assistant at the Business Process Technology research group at Hasso Plattner Institute since 2012. Her research interest is batch processing in business processes, data and business processes, business process simulation and the flexible execution of business processes in IoT environments. Her research appeared in major journals and conferences (e.g. IS, BPM, CAiSE, EDOC, ICSOC). In March 2018, she finished her doctor's thesis about Modeling and Executing Batch Activities in Business Processes. She leads the development of Scylla, an extensible BPMN process simulator. Luise is involved in teaching with regards to BPM, Modeling in Information Systems, etc. Additionally, she has been involved in the organization of two Massive Open Online Courses on BPM and Decision Modeling.

Christiane Bauer

SAP

SAP Young Thinkers – Inspire and Prepare the Next Generation

SAP partners with high – and vocational schools worldwide to support holistic thinking and computational literacy to foster a new generation of out-of-the-box-thinkers. The program aims to engage novices and advanced programmers in creative computing with Snap! on one hand side and on the other hand side invites to think early on of a potential business value of the applications by drafting first business models. To experiment, spread and apply knowledge and more importantly excitement, we offer age-diverse learning events like Go Digital Nights bringing people from different backgrounds, age and expertise level together to enjoy learning.

Christiane Bauer is a manager at SAP, leading teams and programs in the context of Education since many years. Today she leads the global SAP Young Thinkers Community, a program focusing on knowledge and attitude sharing with the next, next generation decision makers. Christiane is passionate about non-standard ways of learning and working together. She likes to try out new things and believes in a 'can do' mentality and 'failure – tolerant' environment as basis for innovative organizations. As a Business & Design Thinking Coach she likes to work with people and is inspired by diversity in all flavors and the power of collaboration combined with creativity and freedom of mind.

Adaptive Memory Protection for Many-Core Systems

Prof. Dr.-Ing.
Wolfgang Schröder-Preikschat

FAU Erlangen-Nürnberg

Hardware-based memory protection is widely applied in all areas of computing, it is a fundamental building block for safety and security. However, improving protection measures by means of dedicated hardware such as a memory management unit (MMU) is not carved in stone, but rather depends on the application domain and the facts of the programming system and hardware, respectively.

State of the art operating systems statically determine whether or not software entities are subject to memory protection. If enabled, costly operations will follow due to multi-level page-table handling, TLB (translation look-aside buffer) invalidations, and the release of inter-processor interrupts. This all reduces performance, increases operating-system noise, and makes system behavior unpredictable. If application programs are type-safe, then enabled hardware protection becomes a pretty much superfluous features. Contrariwise, if the programs are type-unsafe, disabled hardware protection opens door and gate for malware.

This talk is on adaptive memory-protection that is capable of dynamically changing the per-program protection state. The feature is triggered at load/unload time of application programs and applies in the background of running processes. Measurement results are presented and discussed in terms of timing predictability of the system.

Dr. Wolfgang Schröder-Preikschat studied computer science at the Technical University of Berlin, Germany, where he also took his doctor's degree and venia legendi. After a long-term period of extramural research at the GMD, he became full professor for computer science at the Universities of Potsdam, Magdeburg, and Erlangen-Nürnberg. He is reviewer for many international scientific conferences and journals and served as general chair as well as programme committee co-chair and member in numerous conferences such as ATC (USENIX), EMSOFT (ACM), EuroSys (ACM), Middleware (ACM), OSDI (USENIX), and RTSS (IEEE). Dr. Schröder-Preikschat is member of ACM, EuroSys, GI, IEEE, and USENIX. His main interest is on resource-aware (parallel) operating systems, notably process coordination, especially as to time/energy-dependable application and problem domains.

Prof.
Christopher Longhurst

CIO, UCSD Health, Clinical Professor of Biomedical Informatics and Pediatrics, UCSD School of Medicine

Trends in Health Informatics

As Chief Information Officer, Dr. Longhurst is responsible for all operations and strategic planning for information and communications technology across the multiple hospitals, clinics, and professional schools which encompass UC San Diego Health. Dr. Longhurst is also a Clinical Professor of Biomedical Informatics and Pediatrics at UC San Diego School of Medicine, and continues to see patients. As a result of his efforts to leverage technology to improve patient experience in the UCSD Jacobs Medical Center, he was voted the 2017 Top Tech Exec in the education category for San Diego. He previously served as Chief Medical Information Officer for Stanford Children's Health and Clinical Professor at the Stanford University School of Medicine, where he helped lead the organization through the implementation of a comprehensive electronic medical record (EMR) for over a decade. This work culminated in HIMSS stage 7 awards for both Lucile Packard Children's Hospital and 167 network practices in Stanford Children's Health.

Dr. Longhurst has published dozens of scientific articles in peer-reviewed journals on how technology and data can improve patient care and outcomes and was elected as a fellow in the prestigious American College of Medical Informatics, among other distinctions. He is a board-certified pediatrician and clinical informaticist, and founded Stanford's fellowship in clinical informatics, first in the nation to receive accreditation. Described as a pragmatic academician, Dr. Longhurst, serves as an advisor to several companies and speaks internationally on a wide gamut of healthcare IT topics.

Needs Assessment Findings for a Neonatal Intensive Care Unit Communication Tool

Wanjiru Mburu

PhD Student, University of Cape Town, South Africa

Premature delivery and hospitalization in a neonatal intensive care unit (NICU) are stressful events for most mothers. Communication between mothers and NICU staff is an essential part of the support offered to the mothers and can reduce their emotional stress. However, in low-income settings where most NICU are short staffed, interventions to improve infants' overall well-being often take precedence over providing psychosocial support and care to mothers. In this study, we explore how we can involve mothers and NICU staff in the design process of technological intervention that can improve mother-staff communication in the NICU. In the first phase of this study, we used ethnographic observation and semi-structured interviews to assess the needs of mothers. We identify four categories of information that mothers of premature infants need from the NICU staff namely: 1. Breastfeeding information 2. Infant's health information 3. Infant's hospital discharge information and 4. Explanation of commonly used medical terms. Moreover, the interviewed mothers and NICU staff provided several suggestions for suitable interventions that may be used to relay this information.

Wanjiru is an ICT4D researcher who is passionate about using Information and Communication Technologies (ICTs) to bridge healthcare digital divide in developing countries. She holds a bachelor's degree and master's degree both in Computer Science. She is currently a Ph.D. student at the ICT4D center, University of Cape Town. Her research focuses on understanding how Information and Communication Technologies can be used to enhance communication between mothers of hospitalized preterm infants and Neonatal Intensive Care Unit (NICU) staff. She is using co-design approach to involve participants (mothers and NICU staff) throughout the design process. Her research interests are mainly in human-computer interaction for development (HCI4D) and mobile health field.

The Digital Revolution and Transformation of the US Health Care System

Three major forces related to digitalization have emerged that are pushing the US toward integrated health care information systems. 1) "consumerism" and the role of health information transparency for patients, payers and providers, 2) Big data and the practical applications being brought to the front line health care provider and 3) Value-Based health care and the profound shift in provider and payer reimbursement and incentives motivating provider and consumer behaviors.

Prof. **Lawrence Friedman**

Professor of Medicine and Associate Dean for Clinical Affairs, UCSD Health and UCSD School of Medicine

Lawrence Friedman, MD, is Associate Dean for Clinical Affairs; in this role he is responsible for the management oversight of managed care, accountable care, occupational medicine and telemedicine. He also serves as the CEO of International Clinical Programs and is the physician champion for Telemedicine for the University of California with the goal of developing a system-wide telemedicine collaborative. Dr. Friedman is a critical member of the leadership team and plays an integral role in the growth of the Health System as it expands its footprint regionally, nationally and internationally. He is working on strategies to develop the primary and specialty care clinical workforces, as well as supporting the clinical integration between Health System and community practices. Since he began with the Health System in 1994, Dr. Friedman founded the UC San Diego Telemedicine Learning Center and Telemedicine Clinical programs which provide specialty consultation globally. He also co-chaired the University Health System Consortium (UHC) Ambulatory Care Steering Committee; served as a member of the Pay for Performance Implementation Advisory Task Force at the National Quality Forum; and is currently a member of the Agency for Health Care Research Study Section on Health Care Technology and Decision Science. Before assuming full-time administrative responsibilities at UC San Diego Health System, Dr. Friedman was the division chief for General Pediatrics and Adolescent Medicine. Prior to UC San Diego Health System, he was the acting division chief of the Division of General Internal Medicine at the Deaconess Hospital in Boston and Assistant Professor of Medicine at Harvard Medical School. Dr. Friedman completed his residency and internship training at the Deaconess Hospital (now Beth Israel Deaconess) in Boston, followed by a general medicine fellowship at Harvard Medical School. He earned his medical degree at Georgetown University in Washington D.C. Dr. Friedman has authored more than fifty publications including original research, book chapters and review articles, and is currently a reviewer for several peer-reviewed medical journals.

International Branches

International cooperation enables the HPI to extend its research scope and to provide its international partners with an opportunity for close exchange and with access to the HPI's international research contacts. Junior researchers work at their home universities, but receive their scholarships from the HPI in Germany. While being mentored jointly by their home university professors, they will be integrated into the Potsdam HPI Research School's research activities and participate in joint conferences and symposia.

Nanjing University, Nanjing, China

In November 2011 the "HPI Research School at Nanjing University" was opened.

Technion, Israel Institute of Technology, Haifa, Israel

Since April 2010 there is a HPI Research School at Haifa in cooperation with Technion – Israel Institute of Technology too. The Ph.D. students are working at the "HPI Research School at Technion, Haifa" within the HPI Research School team.

University of Cape Town, Cape Town, South Africa

In April 2009, the HPI Research School launched a "branch" in South Africa. The "HPI Research School at University of Cape Town" is mainly researching information and communication technologies that are relevant for developing and emerging countries.

Expansion of the Research School

Annually, the Institute's Research School seeks talented junior researchers and accordingly offers

8 Ph.D. Scholarships and 2 Postdoc Scholarships

The HPI Research School focuses on the foundation and application of large-scale, highly complex and interconnected IT systems. With its interdisciplinary and international structure, the Research School interconnects the HPI research groups as well as its international branches at Cape Town University, Technion – Israel Institute of Technology and Nanjing University. The HPI Future SOC Lab, a state-of-the-art computer center, enriches the academic work at the HPI Research School.

The HPI professors and their research groups ensure high quality research and will supervise Ph.D. students in the following topic areas: Human Computer Interaction, Prof. Dr. Patrick Baudisch; Computer Graphics Systems, Prof. Dr. Jürgen Döllner; Algorithm Engineering, Prof. Dr. Tobias Friedrich; System Engineering and Modeling, Prof. Dr. Holger Giese; Software Architecture, Prof. Dr. Robert Hirschfeld; Internet Technologies and Systems, Prof. Dr. Christoph Meinel; Knowledge Discovery and Data Mining, Prof. Dr. Emmanuel Müller; Information Systems, Prof. Dr. Felix Naumann; Enterprise Platform and Integration Concepts, Prof. Dr. h.c. Hasso Plattner; Operating Systems and Middleware, Prof. Dr. Andreas Polze; and Business Process Technology, Prof. Dr. Mathias Weske.

If you have prior experience in any of these areas, you are invited to submit a full application with the following documents: curriculum vitae and copies of certificates / transcripts, brief research proposal, work samples / copies of relevant scientific work (e.g. master's thesis), and a letter of recommendation.

Applications must be submitted by August 15th of the respective year. Positions are usually available at the beginning of October. Please send your applications to:

research-school-application@hpi.de

For more information on HPI and its HPI Research School see:

https://www.hpi.de/research_school

Social Event on Thursday: Boat Trip to Glienicke Palace & Dinner

We will leave with a boat at 6:30 pm from Griebnitzsee, via Jungfernsee to the Glienicke Bridge. This bridge connects Berlin with Potsdam and formerly was one of the checkpoints between West Berlin and the GDR. During the Cold War, the bridge was used several times for the exchange of captured spies and thus became known as the *Bridge of Spies*.

After the boat trip, we will have dinner at the Glienicke Palace, more particularly in the Remise ("Depot"), which was built in 1928. The Glienicke Palace follows Neoclassical designs of architect Karl Friedrich Schinkel and was developed "in the style of antiquity", after 21-year-old Prince Carl of Prussia returned to Berlin from his first trip to Italy in 1823. The Glienicke property, with Glienicke Palace, the Casino, and the "Great" and "Small Curiosities", characterize the Mediterranean nature of this summer residence to this day.

After the dinner, there will be a shuttle transfer at around 10:45pm back to HPI with a stop in between at Potsdam main station.

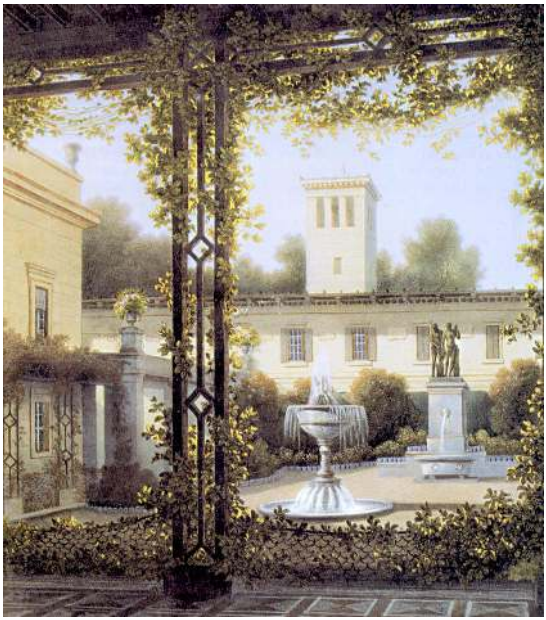

Garden-courtyard in Glienicke, Picture by August C. Haun, 1837, © Public Domain

Maps

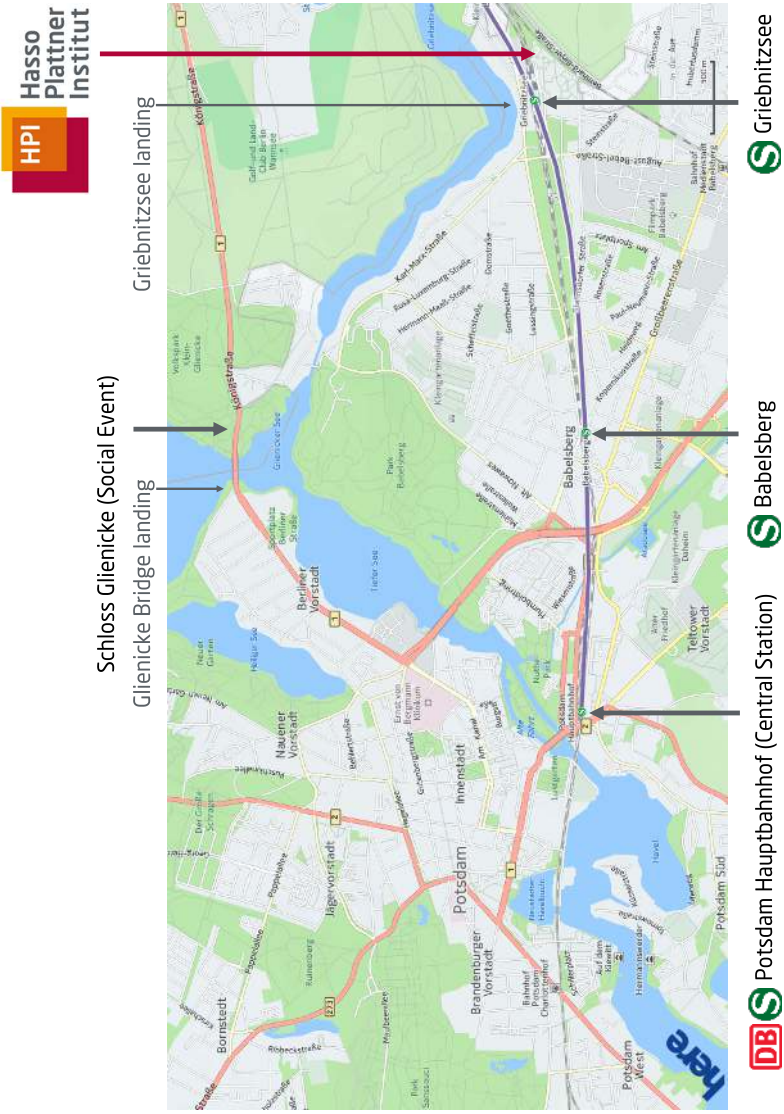

https://www.here.com/?map=52.39625,13.08678,14

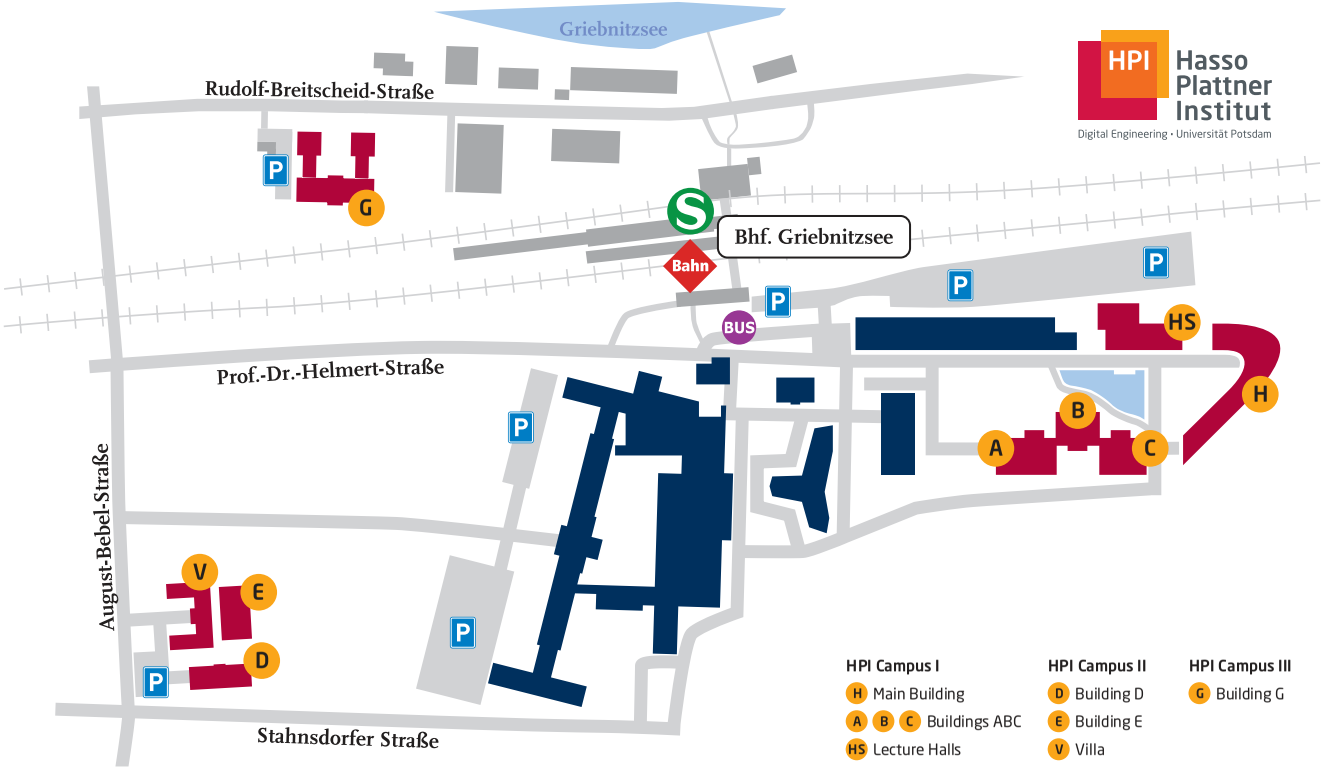
Griebnitzsee
Rudolf-Breitscheid-Straße
Bhf. Griebnitzsee
Bahn
BUS
Prof.-Dr.-Helmert-Straße
August-Bebel-Straße
Stahnsdorfer Straße
HPI Hasso Plattner Institut
Digital Engineering · Universität Potsdam
HPI Campus I
H Main Building
A B C Buildings ABC
HS Lecture Halls
HPI Campus II
D Building D
E Building E
V Villa
HPI Campus III
G Building G

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

Photo: HPI / Kay Herschelmann, Charles Yunck