

HASSO PLATTNER INSTITUT



11th Annual Symposium on Future Trends in Service-Oriented Computing

Hasso Plattner Institute
Potsdam | April 6-8, 2016

Design IT. Create Knowledge.

Agenda

Wednesday, April 6, 2016 (Doctoral Symposium)

Room: [H.E-51](#) (Main Building)

09:15-10:00 **Keynote**

Prof. Dr. Emmanuel MÜLLER, HPI & Deutsches GeoForschungsZentrum GFZ
[Community Detection and Community Outlier Mining in Attributed Graphs](#)

10:00-10:15 **Coffee Break**

10:15-12:15 **Session DS-1**

Daniel RICHTER, Ph.D. student, Hasso Plattner Institute
[Expect Errors: Relaxing System Properties to Increase Availability](#)

Richard MALIWATU, Ph.D. student, University of Cape Town
[Wireless Mesh Routing in White Space](#)

Tim FELGENTREFF, Ph.D. student, Hasso Plattner Institute
[Object-constraint Programming](#)

Xu Ming, Ph.D. student, Nanjing University
[Topic Modeling: Customers' Feedback Information Analysis](#)

Anja JENTZSCH, Ph.D. student, Hasso Plattner Institute
[Exploring Linked Data Graph Structures](#)

Ran BEN BASAT, Ph.D. student, Technion
[Heavy Hitters in Streams and Sliding Windows](#)

12:15-13:15 **Lunch**

13:15-14:00 **Keynote**

Prof. Dr. Avinoam KOLODNY, The Andrew and Erna Viterbi Faculty of Electrical Engineering Technion - Israel Institute of Technology
[The Role of Interconnect in Power-Efficient Computing](#)

14:00-15:00 **Session DS-2**

Alexandra ION, Ph.D. student, Hasso Plattner Institute
[Integrating Mechanisms into Metamaterials](#)

BAO Nianyuan, Ph.D. student, Nanjing University
[EASY: Efficient Adaptive Scheduler on YARN](#)

Ekaterina BAZHENOVA, Ph.D. student, Hasso Plattner Institute
[Decision Process Mining and Refining](#)

15:00-15:30 **Best Paper Session**

Martin KREJCA, Ph.D. student, Hasso Plattner Institute
[Robustness of Ant Colony Optimization to Noise](#)

15:30-16:00 **Poster & Ice Cream Session**

[Ph.D. Students Share Their Ideas & Poster Session](#)

17:00 **Social Event**

[Alternative City Tour](#)

Thursday, April 7, 2016

(Symposium with Industry & Academia, Doctoral Symposium contd.)

Room: HS 3, HS Foyer (Lecture Building)

- 09:15-09:30 **Opening of the Symposium on Future Trends in Service-Oriented Computing**
Prof. Dr. Christoph MEINEL, Hasso Plattner Institute, Director
Prof. Dr. Andreas POLZE, Hasso Plattner Institute, Speaker Research School
- 09:30-10:00 HPI Research School
Ph.D. Students Introduce Their Work—Elevator Pitches
- 10:00-10:45 **Keynote**
Richard P. GABRIEL, Researcher, IBM Research, California
Writing a Writer
- 10:45-11:00 **Coffee Break**
- 11:00-12:20 **Session DS-3**
Wu Xiaotong , Ph.D. student, Nanjing University
The Tradeoff between Statistical Privacy and Utility
Naama KRAUS, Ph.D. student, Technion
Fishing in the Stream: Similarity Search over Endless Data
Anesu M.C. MARUFU, Ph.D. student, University of Cape Town
Cheating Malware in a Resource-Constrained Auction-Based Microgrid
Pedro LOPES, Ph.D. student, Hasso Plattner Institute
Engineering a Spatial Input and Output device using Muscle Stimulation
- 12:20-13:30 **Lunch**
- 13:30-14:15 **Keynote**
Dave PROBERT, Ph.D. , Windows Kernel Architect, Microsoft
Isolated User-Mode, a New Security Architecture for Windows
- 14:15-15:15 **Session 1**
Dr. Sebastian KIEBUSCH, CEO, Saxon IT Services, Dresden
The Future of Governmental IT in Saxony
Dr. Albert MAIER, Chief Architect Information Governance & Data Quality, IBM Analytics Development
Open Source based Information Governance Services for Next Generation Analytics
- 15:15-15:30 **Coffee Break**
- 15:30-16:00 **Invited Talk**
Thomas HARRER, CTO IBM Systems Hardware Sales DACH, IBM Executive Architect, IBM Academy of Technology and Technical Expert Council Central
Evolution in IT with Open Collaboration
- 16:00-16:30 **IBM Power8 Celebration**
Wolfgang WENDT, Vice President Hardware Sales, IBM Systems Europe

- 16:30-17:30 **Session 2**
 Dr.techn. Markus JOBST, Vienna University of Technology
Service-Oriented Mapping: Requirements in Modern Cartography
 Dr. Guido SALVANESCHI, Software Technology Group, Technical University of Darmstadt
Dealing with Change in a Principled Way: The Reactive Programming Approach
- 17:30-17:45 **Coffee Break**
- 17:45-18:45 **Session 3**
 Dr. Matthias KUNZE, Senior Business Process Manager, Zalando SE
Monitoring End-to-End Business Processes in a Microservices Architecture
 Lena FEINBUBE, Ph.D. Student, Hasso Plattner Institute
Experimentally Evaluating the Dependability of Complex Software Systems
- 19:30 **Social Event**
Dinner & Wine Tasting at HPI Main Building Foyer

Friday, April 8, 2016 (Symposium with Industry & Academia)

Room: HS 3, HS Foyer (Lecture Building)

- 09:15-10:00 **Keynote**
 Prof. Dr. Konstantinos PANAGIOTOU, University of Munich
Quantifying Systemic Risk in Inhomogeneous Networks
- 10:00-10:15 **Coffee Break**
- 10:15-11:15 **Session 4**
 Prof. Dr. Idit KEIDAR, Dept. of Electrical Engineering, Technion—Israel Institute of Technology
Transactional Data Structure Libraries
 Prof. Ulrike RIVETT, Department of Information Systems, University of Cape Town
ICT as an Enabler for Development – Is it working?
- 11:15-11:30 **Coffee Break**
- 11:30-12:00 **Invited Talk**
 Dr. Alexander SCHMIDT, Azure Key Vault, Microsoft
Azure in Your Datacenter: Microsoft Azure Stack
- 12:00-12:40 **Session 5**
 Chunrong FANG, Ph.D. student, Nanjing University
Kikbug: A Novel Crowdsourced Testing System
 Thijs ROUMEN, Ph.D. student, Hasso Plattner Institute
Linespace: A Sensemaking Platform for the Blind
- 12:40-13:00 **Closing Remarks**

Hasso Plattner Institute for IT-Systems Engineering

The Hasso Plattner Institute for Software Systems Engineering in Potsdam is Germany's university excellence center for IT Systems Engineering. It is the only university institution in Germany offering bachelor and master programs in "IT Systems Engineering"—a practical and engineering-oriented course of studies in computer science, in which currently 480 students are enrolled. The HPI School of Design Thinking is Europe's first innovation school and is based on the Stanford model of the d.school. It offers 240 places for a supplementary study. There are a total of eleven HPI professors and a further 50 guest professors, lecturers and contract teachers at the Institute. HPI carries out research noted for its standard of excellence in its ten computer science topic areas, as well as at the HPI Research School for Ph.D. candidates, with its further research branches in Cape Town, Haifa and Nanjing. HPI teaching and research focuses on the foundation and application of large-scale, highly complex and interconnected IT systems. The development and exploration of user-driven innovations for all areas of life is an additional area of importance. HPI always earns the highest positions in the CHE university ranking. Since September 2012, the Institute has offered openHPI openhpi.de, an interactive Internet educational network that is open to everyone.



The Symposium on Future Trends in Service-Oriented Computing

The Symposium on Future Trends in Service-Oriented Computing 2016 is the annual symposium of the HPI Research School and is taking place for the eleventh 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-Oriented (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:

<http://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 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 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: Alternative City Tour

Stranded behind the Iron Curtain for nearly 40 years, West Berlin developed its own, unique identity and became a magnet for non-conformists in search of an alternative lifestyle. Sub-cultures, squats, and graffiti art became part of Berlin's distinct style, and with the fall of the Berlin Wall the scene didn't disappear, it simply moved east. Explore the gritty, grungy, underground of the city that Berliners love! Witness the clash between corporate expansion & counter-culture.

Logistics: We depart all together from the main Symposium location at Potsdam Griebnitzsee by S-Bahn train towards Berlin Brandenburg Gate, where we meet our guides. The tour ends in the area of Berlin Warschauer Straße.

Tips: Bring shoes comfortable for long walks.

Straight after the tour, we will have dinner in the area of Berlin Warschauer Straße.

Speakers

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

Christoph Meinel is CEO and Scientific Director of the Hasso Plattner Institute at the University of Potsdam and holds the chair of "Internet Technologies and Systems." He teaches Bachelor and Master courses in "IT-Systems Engineering," oversees numerous Ph.D. projects and is a teacher at the "HPI School of Design Thinking." His research expertise focuses on the areas of Security Engineering, Knowledge Engineering and Web 3.0–Semantic, Social and Service Web. He is also scientifically active in the field of innovation research and the innovative method of Design Thinking. Christoph Meinel is the author or co-author of 14 books and anthologies, as well as editor of various conference proceedings. More than 450 of his papers have been published in scientific journals and at international conferences. He is a member of the National Academy of Science and Engineering (acatech), visiting professor at Beijing University of Technology and serves on many scientific committees and supervisory boards.



Prof. Dr. Christoph MEINEL

Director of the
Hasso Plattner Institute

Prof. Dr. Andreas Polze is the Operating Systems and Middleware Professor at the Hasso Plattner Institute for Software Engineering at University Potsdam, Germany. He is also the speaker of the Ph.D. school on "Service-Oriented Systems Engineering" at HPI 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.



Prof. Dr. Andreas POLZE

Speaker of the
HPI Research School



Prof. Dr.
Emmanuel MÜLLER

Knowledge Discovery
and Data Mining Group,
Hasso Plattner Institute
and Deutsches
GeoForschungsZentrum GFZ

Community Detection and Community Outlier Mining in Attributed Graphs

Attributed graphs are widely used for the representation of social networks, gene and protein interactions, communication networks, or product co-purchase in web stores. For instance, social networks store friendship relations as edges and age, income, and other properties as attributes. These relationships and properties seem to be dependent on each other and exploiting these dependencies is beneficial. Such dependencies are e.g. known as homophily and are widely used. However, looking at multivariate spaces, one can observe that not all given attributes have high dependencies with the graph structure. Consequently, recent graph mining algorithms degenerate for multivariate attribute spaces that lack dependency with the graph structure in some of the attributes.

This talk introduces novel methods that select relevant subspaces, i.e. subsets of the attributes, showing dependencies with the graph. All methods are designed as general pre-processing steps to traditional graph mining and widely applicable. We propose a statistical selection of congruent subspaces, i.e. subsets of attributes showing a dependency with the graph structure. A core challenge in selecting these subspaces lies in the modeling of dependence between graph structure and attribute values. Further, we incorporate the user preference into the selection of relevant subspaces in attributed graphs. For empirical evaluation we consider community detection and community outlier mining.

Prof. Dr. Emmanuel Müller is head of the Knowledge Discovery and Data Mining Group, a joint research group within Hasso-Plattner-Institute at the University of Potsdam and the German Research Centre for Geosciences. Before, he was Independent research group leader at Karlsruhe Institute of Technology and Postdoctoral Fellow at University of Antwerp. Müller received his Diploma and his Ph.D. from RWTH Aachen. He is active member of several program committees including ACM SIGKDD, IEEE BigData, SIAM SDM, organized several workshops and tutorial at major database and data mining conferences. He is leading the open-source initiative OpenSubspace, in which he provides a general contribution to the research community by repeatable and comparable evaluation of recent data mining approaches.

The Role of Interconnect in Power-Efficient Computing

Cloud computing, enabled by the combined capabilities of communication networks and advanced computer technology in data centers, has created exponential growth in demand for data processing. Energy consumption by data centers is growing at a high rate. Similar growth in power consumption had occurred in smaller-scale microprocessor-based computing in recent years, and spurred a wide range of engineering solutions to address the problem. In this talk, the evolution of microprocessor systems is described, highlighting principles and examples of power-saving strategies. The same principles can be applied to large-scale computing centers, in order to address an upcoming power crisis in cloud computing.



Prof. Dr.
Avinoam KOLODNY

The Andrew and Erna Viterbi
Fac. of Electrical Engineering
Technion -
Israel Institute of Technology

Avinoam Kolodny is a professor of electrical engineering at Technion -Israel Institute of Technology. He received his doctorate in microelectronics from Technion in 1980. During twenty years with Intel Corporation, he was engaged in diverse areas including non-volatile memory device physics, electronic design automation and organizational development. He has been a member of the Faculty of Electrical Engineering at the Technion since 2000. His current research is focused primarily on interconnect issues in VLSI systems, covering all levels from physical design of wires to networks on chip, multi-core systems and computing architectures. Prof. Kolodny is a fellow of IEEE.



Daniel RICHTER

Ph.D. Student,
Hasso Plattner Institute

Expect Errors: Relaxing System Properties to Increase Availability


When it comes to make applications more reliable and available, the usual course of action is to remove all faults and prevent errors. Developers spend weeks trying to make the software perfect. Nevertheless, errors and failures are arbitrarily related in time, space, and severity to the cause; in addition developers usually cannot predict in which ways a system will change so all preparations for preventing error were incorrect. Errors should be expected and be inherent in the design. By relaxing system properties like consistency guarantees and result quality, preciseness, or correctness, there are methods given so one does not have to rely on highly-available infrastructure or fault-free components.

With the integration of flexible, imprecise, resilient computation, it is possible to increase the availability of applications without having to rely on special and expensive infrastructure. The goal is to provide a framework to easily integrate the relaxing of system properties into applications, take acceptability properties (relate values in relaxed program to original) into account, and define criteria of improvement (e.g. execute as many optional tasks as possible, get the most precise result, optimize error metrics such as minimize total error, average, or maximum error) that as to be fulfilled.

Daniel is a Ph.D. student at the Operating Systems and Middleware chair of Prof. Andreas Polze at Hasso Plattner Institute in Potsdam and focuses on software implemented fault tolerance. Prior to that, he received his Bachelor's and Master's degree at HPI and worked several years as Software Development Lead for a small ITC company.

Wireless Mesh Routing in White Space


Radio spectrum assignment schemes have traditionally been based on Fixed Spectrum Access (FSA) policies. Under FSA, the regulator assigns frequencies for licensees to operate in and these assignments remain unchanged over the license validity period. White space refers to parts of the licensed radio spectrum that are vacant either spatially or temporally. White spaces are present in any spectrum band but the emphasis is on television broadcast band because they have the most significant amount. These white spaces present an opportunity for wireless communication in urban and rural areas with poor connectivity. Communication using this part of the spectrum offers several advantages such as better transmission range and obstacle penetration in comparison with the current WiFi deployment. Motivated by factors such as channel congestion and limited coverage associated with the 2.4GHz band for the access layer and to some extent the 5GHz band for the back-haul, lately there has been several proposals pushing for a shift from FSA to opportunistic spectrum access i.e. allowing unlicensed users to access the licensed spectrum bands when not being utilised by the licensee. While the concept of opportunistic spectrum access seems fairly simple, the implementation of Dynamic Spectrum Access (DSA) based wireless communication is fraught with several yet to be addressed issues at almost all layers of communication when in ad-hoc mode. The challenges arise from the need for secondary spectrum users to protect the primary users and co-exist with other secondary users. The aim of the proposed research is to build upon the current state of wireless mesh networking technology in DSA and the ideas of community-based infrastructure. The proposed research focuses on the physical, MAC, and network layers to address the problems of clean/optimal channel and optimal route selection in a wireless mesh network.



Richard MALIWATU

Ph.D. Student,
University of Cape Town

Richard did his BSc Computer Science at the Copperbelt University in Kitwe, Zambia in 2009, and BSc (honours) specialising in Information Technology at the University of Cape Town (UCT) in 2012. He completed his Masters in Computer Science in 2014 and is currently working towards his Ph.D. at UCT's ICT4D lab. His general research interest is in computer networks and the current focus is on wireless mesh networking with dynamic spectrum access.



Tim FELGENTREFF

Ph.D. Student,
Hasso Plattner Institute

Object-Constraint Programming

Constraints provide a useful technique for ensuring that desired properties hold in an application. They have been used in a wide range of applications, including graphical layout, simulation, and scheduling. We describe the design and implementation of Babelsberg, an Object Constraint Programming language that cleanly integrates constraints with an object-oriented language in a way that respects encapsulation and OO programming techniques. We have used this language to minimize the coding effort for self-healing applications, load balancers, as well as algorithms such as available-to-promise. We also present performance results that show that the language is competitive, albeit slower than comparable purely object-oriented solutions.

I am a Ph.D. student at the Hasso Plattner Institute in Potsdam. My interests are in language and virtual machine design. I focus on developing language features meant to improve expressibility as well as VM technologies to make sure those language features offer good performance. I am involved in the PyPy/RPython project and various dynamic language VMs built on top of it, and I am the creator of the Babelsberg project, which has spawned a family of languages that integrate constraint solving with object-oriented programming. I prototype much of my research in the self-supporting development environments Squeak/Smalltalk and the Lively Kernel.

Topic Modeling: Customers' Feedback Information Analysis



Xu Ming

Ph.D. Student,
Nanjing University

Nowadays, the competitions among mobile operators are increasingly heating up. To prevent customers' lose, the operators have to improve the quality of services to gain customers' degree of loyalty. But, how can they do more to achieve the goal. A good way is to solve customers' compliments. So we collect customers' feedback information to analyse what they are complaining and help operators make decision.

Topic model is known to help find the latent topics in documents. And LDA(Latent Dirichlet Allocation) is the usual method. We try to use it in many fields, such as the latent risk of banks or secret information analysis. However, LDA is designed for long texts rather than short ones. So it performs not well in such situations. According to this, we came up with an idea to extend the short texts to long ones then we can analyse them to help find out what customers are complaining. And the mobile operators can take measures to solve them based on the information and gain the degree of loyalty.

A second-year Ph.D. candidate in Nanjing University. Interested in recommendation system recently.



Anja JENTZSCH

Ph.D. Student,
Hasso Plattner Institute

Exploring Linked Data Graph Structures

As Linked Datasets grow on the web, their entities and links among them form intricate graphs, and intrinsic patterns emerge. While graph patterns and pattern mining are known concepts with many methods, we propose an explorative and visual approach to engage in and understand the semantics of new datasets. To this end, we formally define a set of frequent patterns based on our initial observations. We then propose algorithms to efficiently extract such patterns from very large Linked Data sets and simultaneously allow the definition of new patterns to extract. Interestingly, we observe many re-occurring motifs across various heterogeneous datasets, suggesting an underlying regularity of how data is accretes.

Anja Jentzsch is a Ph.D. student in the Information Systems Group at Hasso Plattner Institute Potsdam. She is a Linked Data enthusiast, being involved in several Linked Data projects since 2007. Currently she is working on extracting frequent and common graph structures from Linked Datasets. In the past she has been working on Wikidata (structured data storage for Wikipedia), DBpedia (Wikipedia as Linked Data), Silk (interlinking Linked Data sets) and LODD (Linking Open Drug Data).

Heavy Hitters in Streams and Sliding Windows

Identifying heavy hitter flows is a fundamental problem in various network domains. The well established method of using sketches to approximate flow statistics suffers from space inefficiencies. In addition, flow arrival rates are dynamic, thus keeping track of the most recent heavy hitters poses a challenge. Sliding window approximations address this problem, reducing space at the cost of increasing point query time. We present two novel algorithms for identifying heavy hitters in streams and sliding windows. Both algorithms use statically allocated memory and support constant time point queries.



Ran BEN BASAT

Ph.D. Student,
Technion

Ran Ben Basat is a Computer Science Ph.D. candidate at the Technion Institute, Israel. He received his bachelor's degree (summa cum laude) in 2012 and master's (cum laude) in 2015, both from the Technion. Currently, he is working on a whole range of fields including streaming algorithms, information retrieval, and game theory.



Alexandra Ion

Ph.D. Student,
Hasso Plattner Institute

Integrating Mechanisms into Metamaterials

Recent advances in fabrication technology, such as high-resolution 3D printers, allow fabricating objects with internal microstructure, also known as mechanical metamaterials. In this work, we create a new class of metamaterials: mechanisms, i.e., devices that transform forces and movement. These mechanisms from metamaterials require no assembly since they are made from one single material and in one piece. We arrange cells of different topologies to create the desired output force or motion, e.g. a door latch mechanism where the latch is pulled in when the user pushes the handle down.

Designing the internal microstructure—potentially consisting of billions of cells—is a complex task. As a first step, we approach this by creating an editor that helps experts to design such mechanisms from metamaterial. We argue that the key to more complex metamaterials is to allow users to manipulate individual cells and sub-cell elements interactively. We propose a system that converts metamaterials to (stacks of) bitmaps and back, enabling the use of efficient image processing tools. We read color information from pixels and construct the cells of the 3D model's microstructure. Depending on objects' size and (future) 3D printer resolution the number of cells can be arbitrarily large, leading to challenges in terms of space and time complexity. Therefore, we investigate efficient data structures and caching algorithms to reuse already processed blocks of geometry. In the future, we will computationally generate the cells and their arrangement in order to synthesize a mechanism.

Alexandra is a Ph.D. student at the human computer interaction lab at the Hasso Plattner Institute (Potsdam, Germany) of Patrick Baudisch. She is interested in computational aspects of mechanical material properties in fabrication and in haptic feedback systems.

EASY: Efficient Adaptive Scheduler on YARN



Bao Nianyuan

Ph.D. Student,
Nanjing University

The initial design of Apache Hadoop was tightly focused on running massive MapReduce jobs to process a web crawl. However, in recent years, MapReduce has become the defacto scheme for large scale semi-structured or un-structured data processing. Hadoop YARN, the second generation in Hadoop ecosystem, provides a fine-grained resource management. YARN has many resource schedulers, such as FIFO, Fair and Capacity scheduler. Their primary performance concerns are focused on how to minimize the total completion length, i.e., makespan of a set of MapReduce jobs. However, those existing resource schedulers do not consider the dependency between tasks, which is also very crucial for the MapReduce jobs. We find that fully utilizing the dependency between tasks can greatly enhance the efficiency of the resource utilization. Therefore, we propose a new YARN scheduler, named EASY (Efficient Adaptive Scheduler on YARN), using Artificial Fish Swarm Algorithm, which can effectively improve the resource utilization and decrease the makespan of the total MapReduce jobs by taking into account the job profile information and dependency between tasks. We use two partial implementation of EASY as a pluggable scheduler in the recent version of Hadoop YARN and use some classic MapReduce benchmarks to evaluate it. Experimental results shows that, compared with schedulers provided by Hadoop YARN, EASY achieves better efficiency with less execution time.

BAO Nianyuan is now a doctor student in Nanjing University since 2015. He has been in the Multimedia Computing Group in Nanjing University since 2012, and has been participating in a National Science and Technology Major Project of the Ministry of Science and Technology of China, which is under the collaboration of Multimedia Computing Group of NJU and Sinopec. So far, he has accomplished a National Invention Patent in China, and has accomplished a paper accepted by Algorithms and Architectures for Parallel Processing, Springer International Publishing, 2015.



Ekaterina BAZHENOVA

Ph.D. Student,
Hasso Plattner Institute

Decision Process Mining and Refining

Enterprise business process management is directly affected by how effectively it designs and coordinates decision making. In recent years, a number of decision management frameworks have appeared in addition to existing business process management systems. Whether companies develop packaged decision management systems, adopt business rules, or apply advanced analytics to their business, a thorough decision understanding, modeling and execution is critical. To ensure optimal process executions, decision management should incorporate decision logic documentation and implementation. To achieve the separation of concerns principle, the OMG group proposes to use Decision Model and Notation (DMN) complementary to Business Process Model and Notation (BPMN). However, often in practice, decision logic is either explicitly encoded in process models through control flow structures, or it is implicitly contained in process execution logs. Our work proposes an approach of semi-automatic derivation of DMN decision models from process event logs with the help of decision tree classification. As well, we propose a method of improving business process executions by prioritization of input data in decision tasks with the help of dynamic restructuring of underlying decision structures. We demonstrate the approach on a use case from banking industry.

I received the Bachelor and Master degrees in Applied Mathematics and Informatics at the St. Petersburg State Polytechnic University in 2009 and 2011. Back then, I worked as a lecturer and a researcher in the field of integration of enterprise business processes at the SAP Academic Centre of Competence (St. Petersburg). In 2012 I worked on coordinating the supplying of ERP-systems for institutions as a specialist of SAP University Alliances (Moscow). Followed this, I worked as an invited researcher at SAP Research (Pretoria, South Africa) exploring business process management at SMEs on emerging markets. Since March 2013, I am a Ph.D. student at the HPI Business Process Technology Group in Potsdam.

Robustness of Ant Colony Optimization to Noise



Martin KREJCA

Ph.D. Student,
Hasso Plattner Institute

Recently, Ant Colony Optimization (ACO) algorithms have been proven to be efficient in uncertain environments, such as noisy or dynamically changing fitness functions. Most of these analyses focus on combinatorial problems such as path finding.

We rigorously analyze an ACO algorithm optimizing linear pseudo-Boolean functions under additive posterior noise. We study noise distributions whose tails decay exponentially fast, including the classical case of additive Gaussian noise.

Without noise, the classical $(\mu+1)$ -EA outperforms any ACO algorithm, with smaller μ being better; however, in the case of large noise, the $(\mu+1)$ -EA fails, even for high values of μ (which are known to help against small noise). In this talk, we show that ACO is able to deal with arbitrarily large noise in a graceful manner, that is, as long as the evaporation factor ρ is small enough dependent on the variance σ^2 of the noise and the dimension n of the search space, optimization will be successful. We also briefly consider the case of prior noise and prove that ACO can also efficiently optimize linear functions under this noise model.

Martin is a Ph.D. student at the chair for algorithm engineering of Tobias Friedrich at the Hasso Plattner Institute. His research interests lie in the field of evolutionary computation, especially theoretical analyses of estimation of distribution algorithms.

Speakers

Ph.D. Students Introduce Their Work: Elevator Pitches

HPI Research School “Service-Oriented Systems Engineering”, Potsdam

Ahmad Samiei:	Incremental Duplicate Detection
Andreas Grapentin:	Rack Scale Architectures: Resource Management for Future Server Systems
Ankit Chauhan:	Faster Algorithms using Structural Properties of Social Networks
Aragats Amirkhanyan:	Public Safety Intelligence: Social Geotagged Data from the Perspective of Public Safety Awareness
Arian Treffer:	Omniscient Debugging in Database Applications
Cheng Wang:	Neural Visual Translator: Image Captioning with Deep Network
Christoph Matthies:	ScrumLint—Data-Driven Development Process Feedback
Davide Mottin:	Interactive Graph Exploration
Dietmar Funck:	Processing and Analyzing of Multi-Temporal 3D Point Clouds
Dominique Blouin:	Model Synchronization for Complex Industrial Systems
Mina Rezaei:	Brain Image Analysis by Deep Learning
Ralf Rothenberger:	Models for Industrial SAT
Robert Kovács:	TrussMaker—Scaling Personal Fabrication by Embedding Ready-Made Objects
Sankalita Mandal:	Using Events to Implement Processes
Sona Ghahremani:	Utility-Driven Modularized MAPE-K Loop Architectures for Self-Adaptive Systems
Sören Discher:	Scalable Visualization of Massive 3D Point Clouds
Stefan Lehmann:	Abstraction-Aware Debugging Tools
Thomas Brand:	Runtime Data-Driven Software Evolution in Enterprise Software Ecosystems
Toni Mattis:	Code Mining in Live Programming Environments

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.

Writing a Writer

I have written a program that writes haiku and revises text to mimic other writers. In Summer 2015, eighteen haiku written by it were critiqued by three professional poets in a writers' workshop, and those poems were judged to be good poems written by a person—a form of the Turing Test. How does it work and what does it mean?



Richard P. GABRIEL

Researcher,
IBM Research, California

Richard Gabriel is an eclectic researcher, working in Artificial Intelligence, programming languages, programming environments, parallel programming, software patterns, open source theory and practice, and now natural language processing (again). He has a Ph.D. in Computer Science from Stanford, and an MFA in Creative Writing from Warren Wilson College. He has written a poem a day since March 2000.



Wu Xiaotong

Ph.D. Student,
Nanjing University

The Tradeoff between Statistical Privacy and Utility

Statistical privacy is a multi-disciplinary field that studies how to mine the valuable information from the raw data while preventing the leakage of the sensitive information. Naturally, one of the important challenges for data dealer is to focus on the tradeoff between utility and privacy, which is to maximize the utility measure constrained by privacy considerations. Unfortunately, current privacy definitions and utility metrics still change, so that basic properties of optimization solutions can differ. In view of this challenge, we attempt to survey the existing privacy definitions and utility metrics in the previous literatures and present the corresponding findings about the tradeoff in the experiments. We also list some recent research about desirable properties of privacy definitions and utility measures by generic mathematical classes and interesting results. Finally, we discuss the future work and give some possible ways of the tradeoff between statistical privacy and utility.

WU Xiaotong is currently working towards the Ph.D. degree at the Department of Computer Science and Technology, Nanjing University, China. He has received his Bachelor's degree and Master's degree in Software Engineering from Central South University and Department of Computer Science and Technology from Nanjing University of China, respectively. His research interests include cloud computing, resource allocation, privacy protection and network security.

Fishing in the Stream: Similarity Search over Endless Data



Naama KRAUS

Ph.D. Student,
Technion

Similarity search is the task of retrieving data items that are similar to a given query. In this talk, we introduce time-sensitive similarity search over endless data-stream (SSDS), such as news and social media feeds.

SSDS is challenging as it processes unbounded data, whereas computation resources are bounded. We propose Stream-LSH, an effective and efficient randomized SSDS algorithm using bounded space capacity. Our algorithm is based on Locality Sensitive Hashing (LSH), a widely used similarity search algorithm for massive high dimensional datasets. LSH builds a hash-based index with some redundancy in order to increase the search recall. Stream-LSH exploits capacity resources efficiently by 1) adjusting the redundancy of indexed items based on their quality, and 2) using a retention policy which bounds space capacity by eliminating old items from the index. We suggest a novel randomized retention policy, which gradually eliminates index entries as a function of their age.

To the best of our knowledge, Stream-LSH is the first SSDS algorithm that combines similarity, quality, and temporal metrics in its indexing process. We theoretically analyze its effectiveness, and show that it increases the probability to retrieve similar, timely, and high quality items compared to state-of-the-art alternatives using the same space capacity. We further conduct an empirical study using the Reuters real world news stream dataset, which shows that Stream-LSH increases recall in 10% compared to the competition.

Naama Kraus is a Ph.D. student in the Electrical Engineering Department, Technion IIT, under the supervision of Prof. Idit Keidar and Dr. David Carmel. She holds MSc. in Computer Science from the Computer Science department, Technion IIT, and Bsc. in Computer Science and Mathematics (Cum Laude) from the Bar Ilan University in Israel. Naama's research lies on the border of information retrieval and distributed computing. She explores search in large scale environments, such as peer-to-peer networks and data streams. Prior to her Ph.D. studies, Naama was a research staff member in IBM Haifa research lab, where she was involved in a variety of projects in the areas of information retrieval, parallel computing, and file systems.



Anesu M.C. MARUFU

Ph.D. Student,
University of Cape Town

Cheating Malware in a Resource-Constrained Auction-Based Microgrid

A noticeable trend evident within power systems is the transformation from a centralised, producer-controlled smart-grid to a more distributed and consumer-interactive microgrid. Introduction of the microgrid concept includes the creation of intelligent electricity supply networks to allow efficient use of energy resources and reduction of carbon emissions. The usual practice in the microgrid is matching of supply to demand, as electricity must be used as it is being generated. This is made possible by use of auction mechanisms to shed spikes in load consumption. Among the plenitude of auction formats available, the Continuous Double Auction (CDA) is the most popular. Actors in the auction market system are self-interested and adversarial attacks may be realised. Security of CDA (especially within microgrids) has been given limited attention, which raises a major concern considering the rise of malware and limited research work within: the smart grid, resource constrained platforms and auction systems. Our work formally proposes a market framework to allocate energy in a resource constrained environment; highlights plausible malware augmented cheating attack models; demonstrates some attack scenarios and seeks to provide some mitigation measures to the cheating attacks.

Anesu Marufu is a Ph.D. student 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 malware threat modeling, Double Auction security, Multi-agent system applications and security.

Engineering a Spatial Input and Output Device Using Muscle Stimulation

We explore how to create interactive systems based on electrical muscle stimulation of high expressiveness. We present a system that assists mobile users in cognitively demanding activities, such as solving math or the aerodynamics of objects inside a wind tunnel, by providing them with access to a computer system. Using pen-on-paper interaction, users write formulas and the system responds by drawing graphs. Unlike earlier work, however, our system uses the pen not only for input, but also for output: it uses electrical muscle stimulation to make the user's hand plot.



Pedro LOPES

Ph.D. Student,
Hasso Plattner Institute

Pedro is a Ph.D. student of Prof. Patrick Baudisch's Human Computer Interaction lab in Hasso Plattner Institut. Pedro creates wearable interfaces that read & write directly to the user's body through our muscles [proprioceptive interaction]. Pedro augments humans & their realities by using electrical muscle stimulation to actuate human muscles as interfaces to new virtual worlds. His works have been published at ACM CHI and UIST. Enjoys writing about tech and is the digital content editor of the ACM XRDS.



Dave PROBERT, Ph.D.

Windows Kernel Architect,
Microsoft

Isolated User-Mode, a New Security Architecture for Windows

Windows 10 introduces a significant addition to the kernel architecture of Windows: Isolated User-Mode (IUM). IUM is built on a Windows 10 hypervisor feature - Virtual Secure Mode (VSM)—which uses the second-level address translation tables to create normal/secure modes which are orthogonal to the usual user/kernel modes. IUM runs the user-mode part of specially signed processes (called Trustlets) in secure user-mode, and relies on the OS running in normal kernel-mode for OS services. These services are accessed using a proxy, the SecureKernel running in secure kernel-mode. The SecureKernel hardens IUM against attacks from a compromised OS. In Windows 10 IUM was used to build a handful of facilities - collectively called Virtualization-Based Security (VBS)—to demonstrate the effectiveness of this architecture. This talk will cover the VSM and IUM architectural changes to Windows, the nature of Trustlets, and the VBS facilities that were added to Windows using this technology.

Dave Probert has been a kernel architect in Windows for almost 20 years, as both a manager of kernel development and an individual contributor. He has developed a number of Windows technologies, including Isolated User-Mode, the background execution model for modern apps, the Windows Notification Framework, User-Mode Scheduling, and Hotpatching. As part of the Windows Academic Program, Dave published the Windows kernel source (as the WRK) and has spoken at universities around the world. Dave has a Ph.D. from University of California at Santa Barbara, where he developed the SPACE OS architecture, which was later used at Bell Labs to build Pebble.

The Future of Governmental IT in Saxony

The state-owned company "Saxon IT services" (SID) was founded in 2008. It has 450 employees at 5 sites and realizes a total revenue of 64 Mio €. The consolidation of further state-owned IT sites, employees and IT services is an ongoing process. This process imposes a major challenge for the SID since the requirements of the particular Saxon ministries and departments regarding data integrity and confidentiality have to be balanced. SID meets this challenge both by offering customized IT services based on well-proven technologies and by attending a dedicated research program for cloud computing.



Dr. Sebastian KIEBUSCH

CEO,
Saxon IT Services, Dresden

Dr. Sebastian Kiebusch is since 2012 the chief executive officer of the state-owned company Saxon IT Services (SID). Previously he served as a deputy managing director for the Brandenburg IT services (ZIT-BB) and worked as an IT consultant for the management of a well-known private bank. Moreover he was a researcher at the University of Leipzig and worked for the innovation management of a large retail company.



Dr. ALBERT MAIER

Chief Architect Information
Governance & Data Quality,
IBM Analytics Development

Open Source based Information Governance Services for Next Generation Analytics

Information Governance is a key enabler for any “Next Generation Analytics Platform” (NGAP). This session briefly explains what Information Governance is about and - using a Data Lake solution architecture as example - discusses IBM’s approach to NGAP, stressing the role and importance of metadata and governance services. IBM’s new platform is based on open source. Using the “Shared Discovery Platform” (SDP) and services like profiling services, term classification services, join recommendation services as example, this approach is motivated and detailed. SDP is based on Apache Atlas and Kafka. IBM is cooperating with the HPI in this area.

Dr. Albert Maier is a Senior Technical Staff Member in IBM’s Analytics development organization. He is the chief architect for IBM’s Information Governance, Metadata and Data Quality portfolios. Prior to this job, Albert has held a variety of technical lead positions in various IBM organizations including Research and Global Services. Albert’s interests have ranged broadly across information integration, middleware integration, and databases. Albert is an IBM Master Inventor and has received many patents and awards for his contributions to Information Management products. He is also a member of the IBM Academy of Technology.

Evolution in IT with Open Collaboration

A new set of IT application, which leverage cognitive systems and machine learning, is evolving quickly—these applications will innovate many areas of economy and life. They will contribute to a continued exponential growth in data volumes (50% per year). On the other hand, the advancement of information technology has slowed down in the last years and traditional computing will hit more and more physical limits in the foreseeable future. Will the slow down of technology hit the innovation capabilities? Probably not, since there are still some breakthroughs expected in basic technologies—and there is a large opportunity to accelerate and optimize the whole stack of computing. The presentation will provide insights into industry drivers and open collaboration based initiatives to address the optimization of the compute and storage stack of the future.



Thomas HARRER

CTO IBM Systems
Hardware Sales DACH,
IBM Executive Architect,
IBM Academy of Technology
and Technical Expert Council

Thomas Harrer is CTO for IBM Systems HW Sales for Germany, Switzerland and Austria (DACH). In this role, he is responsible for a holistic technical strategy for IBM infrastructure solutions for the local markets. As an experienced Client Technical Architect, he designed critical IT infrastructures for large clients. As an IBM Executive Architect, he works on interdisciplinary challenges and bridges between business innovation and information technology. Thomas Harrer is member of the IBM Technical Expert Council Central (TEC), a group of leading technical experts in IBM DACH. Between 2011 and 2013, Thomas Harrer was chairman of the TEC Central. He is a member of the IBM Academy of Technology, the worldwide IBM community of the top technical leaders.



Wolfgang WENDT

Vice President
Hardware Sales,
IBM Systems Europe

IBM Power8 Celebration

As an extension of the long term cooperation between HPI and IBM, IBM Systems will provide Power hardware for research purposes focusing on topics such as in-memory technologies, profiling, and monitoring on IBM Power architecture systems; integration of hardware accelerators; and evaluating the effects of non-volatile memory.

Besides the new high-end Power system that will enable new dimensions of research opportunities, IBM will promote several scholarships at HPI, too.

Wolfgang Wendt is Vice President, IBM Systems Hardware Sales, with a focus on Europe. He is responsible for the infrastructure portfolio of IBM, covering physical and software defined server and storage infrastructures, and is particularly focused on helping clients achieve business agility through the deployment of hybrid cloud and analytics to enable their digital transformation. Prior to this position, Wolfgang was Vice President, IBM Systems Hardware Sales for a region spanning Austria, Germany and Switzerland. In recent years, Wolfgang has held a number of roles, including Vice President Mainframe Sales for IBM's Growth Markets in Asia, Australia, Eastern Europe, Africa and Latin America, as well as Vice President for Mainframe Hardware, Software and Services in Germany. Prior to that, Wolfgang managed the global relationship with Daimler-Benz, with responsibility for sales to Mercedes-Benz production plants worldwide.

Wolfgang joined IBM in 1991, having previously completed University Studies in Computer Science and Economics, and undertaken a number of Business consulting projects. Wolfgang is married with 2 children, and lives in Stuttgart Germany.

Service-Oriented Mapping: Requirements in Modern Cartography


Modern cartography undergoes massive changes of paradigm how maps are produced and used. Whereas well established production lanes follow a procedure of data collection, modelling, production and dissemination, the quite new paradigm of Service-Oriented Mapping has to build upon decentralized sources.

The disadvantages are still missing semantics for data and services, misunderstandings in registries, inhomogeneous metadata and inconsistent spatial data. Therefore the interoperability framework is not exactly brilliant and results in failures of the pragmatic dimension: the usability of maps.

The advantages of Service-Oriented Mapping are related with availability, accessibility and its value adding behaviour. Geospatial data combinations sourced in various distributed (sensor) networks become available for map applications in almost real-time. Actual application fields can be found in meteorology, environment or industry. The accessibility to resulting maps is transparently managed and its existence is progressively published. The growing number of geoportals and geospatial catalogue services broaden the access to available sources and illustrate their importance. A value adding behaviour can be observed by serendipity effects within Service-Oriented Mapping. New data and service combinations lead to unpredictable/unforeseen outcomes that add value for specific use cases. The movements of open (government) data announce the number of created applications as indicator of the business value.

In terms of a sustainable Service-Oriented Mapping we have to ask for the main requirements in modern cartography. Do service-oriented architectures help to fulfil the requirements or is this new paradigm a blind alley for modern cartography?

Markus Jobst is deputy head of department information management at the Austrian Federal Office of Metrology and Surveying (BEV) and researcher at the Vienna University of Technology. He is coordinating the implementation of the European spatial data infrastructure INSPIRE at BEV and acts as consultant for Austrian ministries and technical liaison to the European Commission. As subgroup chair in the UNGGIM:Europe regional committee he works together with national bodies on the future requirements of geoinformation management. The cross-domain network for Service-Oriented Mapping is brought together in the ICA Commission on Map Production and Geoinformation Management, in which he is one of the chairpersons. Markus Jobst received his Ph.D. with the focus on "a semiotic model for the cartographic communication with 3D" at the Vienna University of Technology in 2008.



Dr. techn. Markus Jobst

Vienna University
of Technology



Dr. Guido SALVANESCHI

Software Technology Group,
Technical University
of Darmstadt

Dealing with Change in a Principled Way: The Reactive Programming Approach

Reactive applications are ubiquitous and still surprisingly hard to develop and maintain. The Observer design pattern, the preferred way to implement reactivity in the object-oriented setting, decouples observers from the observables, but has a number of limitations—such as lack of composability—that has been recognized in literature for long.

Starting from the first investigations with strictly functional languages, reactive programming has gained popularity as -the- programming paradigm for this class of software. Reactive programming improves software design by making change propagation in reactive applications part of the language runtime, thus providing the developer relief from managing updates explicitly. The resulting programming style, based on event streams and signals, enhances composability and is easier to comprehend. This talk will present recent research in the area aiming to provide a language that integrates reactive abstractions and Object-oriented programming. We will also outline current work on extending reactive programming to the distributed setting and its evaluation with case studies and controlled experiments.

Guido Salvaneschi is a postdoctoral researcher at TU Darmstadt. He completed his Ph.D. at Politecnico di Milano, under the supervision of Prof. Carlo Ghezzi, visiting Mira Mezini's group at TU Darmstadt (Spring 2011) and Barbara Liskov's group at CSAIL MIT (Fall 2011). His current research interests focus on programming language design of reactive applications, such as event-based languages, data flow languages and functional reactive programming. His work includes the integration of different paradigms, incrementality and distribution. Some of Guido's recent publications appear in ICSE, PLDI, OOPSLA, FSE, Modularity, DEBS and IEEE Software. He served in the PC of Modularity'16 and '14 and co-organized the REBLS workshop series at SPLASH'13-'15.

Monitoring End-to-End Business Processes in a Microservices Architecture



Dr. Matthias Kunze

Senior Business
Process Manager,
Zalando SE

Zalando is moving swiftly from an online retailer to a multi-service platform that connects people with fashion. New business channels involving a myriad of different partners demand for an architecture for innovation based on a highly flexible and scalable infrastructure - both technically and organizationally.

In 2015, Zalando introduced Radical Agility - a tech culture that reinforces values to ensure our products and technologies grow and thrive. The new approach gives teams full autonomy in the development and operation of services and allows for fast and disruptive innovation. Microservices as technological paradigm enable us to manage scalability and complexity.

At the same time, customer satisfaction remains among our top objectives and requires us to make absolutely sure that our business processes are not stalled or delayed. Yet, in a tremendously distributed and heterogenous environment, where every single service in our control works as expected, things can go amiss. Business process monitoring safeguards our operations by tracking our operations end-to-end, agnostic of the technical details of the underlying infrastructure.

We present some of the technical challenges of process monitoring in a radically agile architecture and our approach to tackling them, enabling autonomous and continuous monitoring, and reporting of business processes.

Matthias Kunze is a business process manager and business process consultant in the Business Excellence team at Zalando SE. The team is responsible for ensuring the reliable performance of Zalando's operations and business continuity, and to enable stakeholders in Zalando to effectively manage and control their business processes. Before joining Zalando in 2015, he was a student, research associate, and postdoctoral researcher at the Hasso Plattner Institute. Besides his contributions to many practical aspects of BPM, his research work focussed on the behavioural analysis of business process models. Matthias Kunze received his Ph.D. in 2013 for his dissertation on "Searching Business Process Models by Example".



Lena FEINBUBE

Ph.D. Student,
Hasso Plattner Institute

Experimentally Evaluating the Dependability of Complex Software Systems

Fault injection is a useful tool for practical experimentation. It can be used to test the fault tolerance features by artificially inserting faulty behaviour and error states into a running system. Thus, fault injection represents the notion that the faultload needs to be incorporated into the software testing process.

A dependable software system must be fault tolerant at all layers of the software stack.

This talk discusses fault injection at the operating system and distributed middleware layers.

At the operating system level, we show how different classes of software faults in libraries can be simulated using link-time API hooking. The fault classes are chosen as representatives from the community-maintained CWE database of real world software problems.

At the distributed level, modern cloud services have to fulfil needs such as scalability, availability and security simultaneously. Therefore, experimental dependability assessment using distributed fault injection is steadily gaining relevance. On-going research on fault injection in OpenStack is discussed.

Lena Feinbube is a Ph.D. candidate in the HPI research school. As a member of the Operating Systems and Middleware Group supervised by Prof. Polze, her research focusses on software dependability and fault injection in real world applications.

Quantifying Systemic Risk in Inhomogeneous Networks

Networks of several types—biological, technological, social—have been in the focus of science for a long time, as they enable us to describe in an abstract way relationships and interactions between elements of complex and heterogeneous systems. In this context, understanding and quantifying the notion of “systemic risk”, which determines to what extent local adverse shocks can propagate to large parts of the system through a cascade, plays an increasingly central role. In this talk the focus will be on two aspects: first, the formulation of a general mechanism (“bootstrap percolation”) that can be used to describe how a certain activity disperses on a given finite graph. Second, the introduction of a broad family of models that allows for a fine description of not necessarily symmetric relationships between the vertices. We conduct a thorough mathematical analysis of the model; among other results, we quantify the notion of systemic risk, and we discover novel features that make graphs prone or resilient to initially small activity.



Prof. Dr.
Konstantinos PANAGIOTOU

Professor for Discrete and
Algorithmic Mathematics,
Department for Mathematics,
University of Munich

Konstantinos Panagiotou has been a professor for Discrete and Algorithmic Mathematics at the University of Munich since 2012. Before going to Munich, he was a senior researcher at the Max Planck Institute for Informatics and a visiting professor at the Université Pierre-et-Marie-Curie, Paris. He received his doctoral degree from ETH Zurich in 2008 and his Diploma in 2004, both in Computer Science. His research is concerned with various algorithmic and combinatorial questions regarding the complexity of optimization problems, with a special focus on random discrete structures and randomized algorithms.



Transactional Data Structure Libraries

Prof. Dr. Idit KEIDAR

Dep. of Electrical Engineering,
Technion—
Israel Institute of Technology

We introduce transactions into libraries of concurrent data structures; such transactions can be used to ensure atomicity of sequences of data structure operations. By focusing on transactional access to a well-defined set of data structure operations, we strike a balance between the ease-of-programming of transactions and the efficiency of custom-tailored data structures. We exemplify this concept by designing and implementing a library supporting transactions on any number of maps, sets (implemented as skiplists), and queues. Our library offers efficient and scalable transactions, which are an order of magnitude faster than state-of-the-art transactional memory toolkits. Moreover, our approach treats stand-alone data structure operations (like put and enqueue) as first class citizens, and allows them to execute with virtually no overhead, at the speed of the original data structure library.

Joint work with Alexander Spiegelman and Guy Golan-Gueta, to appear in PLDI 2016.

Idit Keidar is a Professor and the Associate Dean for Graduate Studies at the Viterbi Department of Electrical Engineering at the Technion. She also heads the Networked Software Systems Laboratory (NSSL). She received her BSc (summa cum laude), MSc (summa cum laude) and Ph.D. from the Hebrew University of Jerusalem in 1992, 1994, and 1998 resp. Prof. Keidar is a world renown expert in distributed computing and concurrency, having developed numerous distributed and networked systems, including for storage, multicast, group communication, group membership, distributed transactions and atomic commit, cloud services, trusted access to data in cloud storage, and concurrent programming. Her awards include the Yanai Award for Excellence in Academic Education, the Muriel and David Jackow Award for Excellence in Teaching, the David Dudi Ben-Aharon Research Award, the Allon Fellowship, the Rothschild Yad-Hanadiv fellowship for postdoctoral studies, and a Wolf Foundation Prize for Ph.D. students.

ICT as an Enabler for Development - Is it working?



Prof. **Ulrike RIVETT**

Department of
Information Systems,
University of Cape Town

ICT4D was a phrase that expressed hope, trust and an ambition for technology becoming the change we wanted to see in the world. The term development meant many things to many people. For most of the members of the ICT4D community, ICT was the distinguishing element. The use of mobile technology and the development of an application that would address a social ill or a need, would justify the use of the term ICT4D. Development was closely linked to “improving the lives of people” and “people” referred to someone in an under-resourced or developing nation. One of the reasons for trusting that ICT could have this impact, was linked to the notion of the “bottom of the pyramid”. The poorest of the poor were prepared to pay for airtime, for phones, for communication. By having a common tool that the majority of the world population either had or wanted to have, the hope lay with creating new avenues of engagement. A decade later, we have more failed ICT4D projects than we could have ever imagined. The enthusiasm and positive spirit of yet another application has worn thin and the doubt is growing if there is any change in sight. Will the poor remain poor and only use their mobile phone to send money to their rural family members rather than holding a government responsible? This presentation reflects on research done by the iCOMMS team at the University of Cape Town looking at the success and failure of ICT4D.

Ulrike Rivett is a Professor in the Department of Information Systems at the University of Cape Town. Her research contribution over the last decade has been to “connect the dots” between the academic knowledge of ICT4D, the creation of ICT applications that offer an innovative approach to existing problems and the engagement with users. She heads up the iCOMMS research team which focuses on understanding the use ICT systems for the benefit of society. Her work and publications have been in the field of mHealth, ICT4D and ICT in the Water and Sanitation sector.



Dr. Alexander SCHMIDT

Azure Key Vault,
Microsoft

Azure in Your Datacenter: Microsoft Azure Stack

After giving a brief introduction to Microsoft Azure and the hyper-scale of the platform, the talk will introduce Microsoft Azure Stack, the “downloadable version of Azure.” Microsoft Azure Stack is a new hybrid cloud platform product that enables organizations to deliver Azure services from their own datacenter. Organizations get the best of two worlds: the power of cloud services, while maintaining control of their datacenter for true hybrid cloud agility. Organizations decide where to keep their data and applications—in their own datacenter or with a hosting service provider.

Alexander Schmidt is an alum of the HPI Research School and the Operating Systems and Middleware group from which he received his doctorate degree in 2011. After joining Microsoft, Alexander worked with the Windows Server group on Windows Server Update Services, Microsoft’s on-premise technology for managing software updates, and on an effort which is now known as the Microsoft Azure Stack. Alexander currently works with the Azure Key Vault team, which provides a cryptographic key management Cloud service in Azure.

Kikbug: A Novel Crowdsourced Testing System



Chunrong FANG

Ph.D. student,
Nanjing University

Crowdsourced testing is perfect for mobile testing, especially on the Android system where apps need to be tested on diverse realistic platforms by different real users. However, the current crowdsourced testing systems are not able to effectively support construction of reports with comprehensive information that directly lead to a bug resolution. To address the need for an improved crowdsourced testing system, we introduce a novel solution, called Kikbug, that benefits both requesters and crowd workers in crowdsourced testing practice.

Kikbug provides a light-weight android app to facilitate testing jobs. During the testing process, crowd workers focus on the jobs, with Kikbug Android automatically collecting logs, recording user's actions, and providing an easy-to-use button for taking screenshots. With the collected information, Kikbug provides the requesters a series of tools on testing tasks, such as bug reproduction (REACT), bug summarization (S-CAT) and job quality evaluation (E-CAT). Kikbug has served several large corporations, such as Alibaba and Netease, for 5 popular apps. And the current results show that (1) Kikbug can effectively reduce the workers' efforts on reporting bugs, (2) a majority of bug reports can be replayed by REACT, and (3) the summarized reports can facilitate report inspection.

Chunrong Fang is a Ph.D. student in software engineering at Nanjing University. His research interests lie in software testing, especially on test case prioritization, mobile testing and crowdsourcing testing. His work has been published in TR, SQJ, SCIS, FSE, ISSRE etc.



Thijs ROUMEN

Ph.D. Student,
Hasso Plattner Institute

Linespace: A Sensemaking Platform for the Blind

For visually impaired users, making sense of spatial information is difficult as they have to scan and memorize content before being able to analyze it. Even worse, any update to the displayed content invalidates their spatial memory, which can force them to manually rescan the entire display. Making display contents persist, we argue, is thus the highest priority in designing a sensemaking system for the visually impaired. We built a tactile display system designed with this goal in mind. The foundation of our system is a large tactile display, which we achieve by using a modified 3D printer to print raised lines of filament. The system's software then uses the large space to minimise screen updates. Instead of panning and zooming, for example, our system creates additional views, leaving display contents intact and thus preserving user's spatial memory. In this talk I will introduce the concept of Linespace and show various new applications that have been developed using the linespace framework.

Thijs Roumen has been a Ph.D. student at the Hasso Plattner Institute for the last 1,5 years under supervision of Prof. Dr. P. Baudisch (Human Computer Interaction). He has a strong interest in digital fabrication and ways to use this technology. He has been involved in various projects and currently is driving the linespace project forwards which he developed together with Saiganesh Swaminathan.

International Branches

International cooperations enable the HPI to extend its research scope and to provide 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.

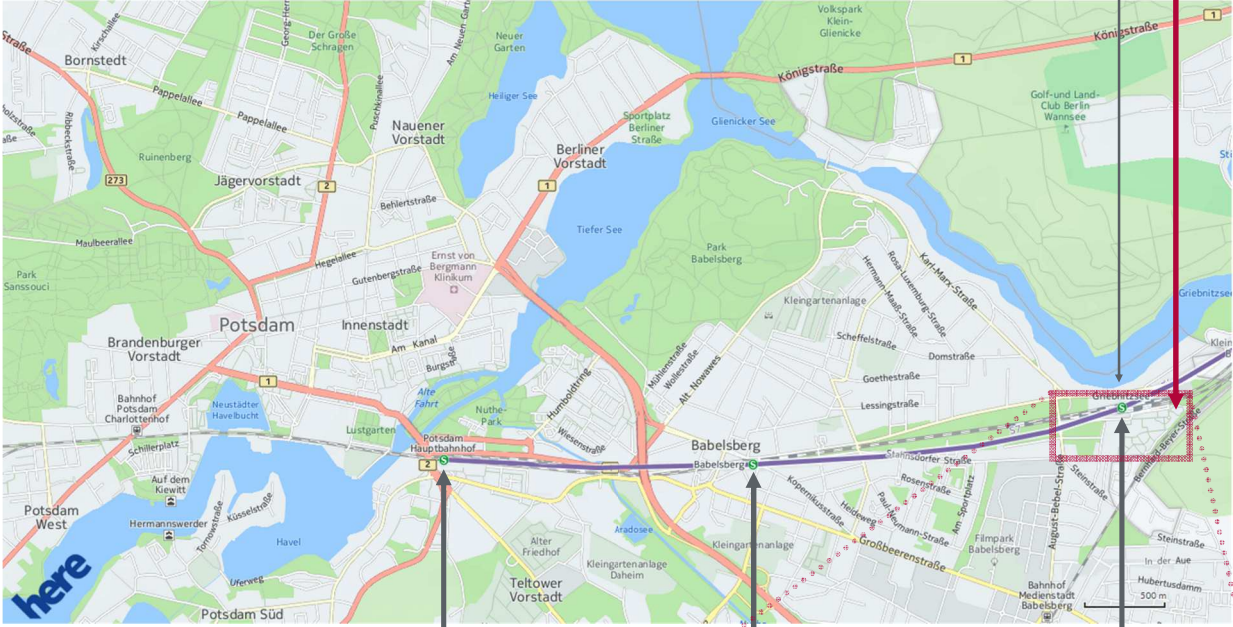
Applications must be submitted by August 15th of the respective year. Positions are usually available at the beginning of October. Please send your applications to:


research-school-application@hpi.de


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


http://www.hpi.de/research_school



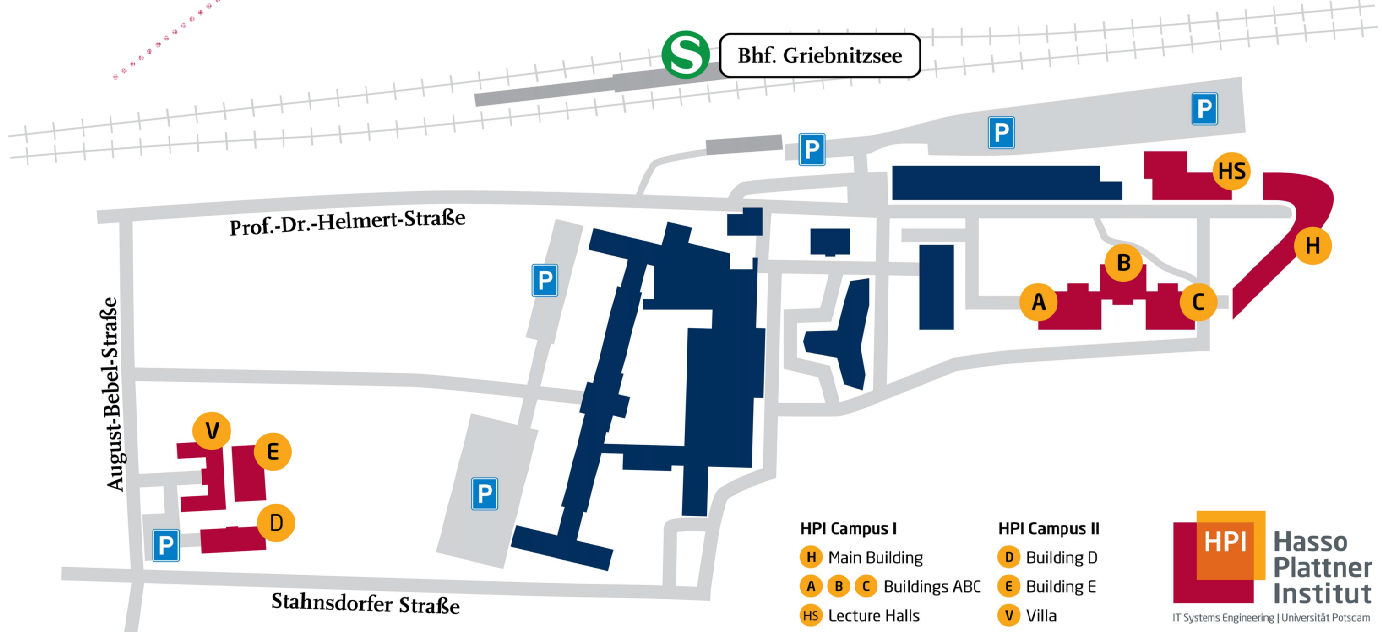


 Potsdam Hauptbahnhof (Central Station)

 Babelsberg

 Griebnitzsee

Campus Griebnitzsee



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



Contact

Hasso-Plattner-Institut für Softwaresystemtechnik

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

Phone +49 331 5509-0
E-Mail hpi-info@hpi.de
Web www.hpi.de

Photo: HPI/Kay Herschelmann