Progress Report of the Max Planck Institute for Software Systems (MPI-SWS)

November 2013 – July 2015

Max
Planck
Institute
for
Software Systems

Contents

1 State of the Institute

This progress report of the Max Planck Institute for Software Systems (MPI-SWS) covers the period November 2013 – July 2015. We begin with an overview of:

- the mission, goals, and general structure of the institute (Section 1.1), and
- the current state of the institute and our recent accomplishments (Section 1.2).

The subsequent sections of the document provide individual progress reports by the institute's 13 research groups (9 current, 2 adjunct, and 2 former) that were active during this review period. Finally, Section 15 provides summary information and details about the institute and its activities.

1.1 General overview of the institute

This section presents a general overview of the goals, structure, and organization of the institute, not specific to the present review period.

1.1.1 Challenges in software systems

Software systems is the part of computer science that lays the foundation for the practical use of information science and technology. We interpret the term broadly to include all areas of computer science and related disciplines that contribute to the design, analysis, implementation, and evaluation of software-based systems. Thus, we include research in the design and implementation of dependable, distributed, and embedded systems and networks; software engineering and verification; programming languages and programming systems; security and privacy; databases, information retrieval, and data science; social computing systems and human-computer interaction. Across these areas, we emphasize collaboration and combine theory, empirical, and data-driven methodologies to address fundamental challenges in software systems.

During the institute's first decade, the following broad research areas have emerged as focal points; they are each being addressed by multiple groups within the institute and as part of external collaborations:

- **Practical privacy:** The advent of Cloud-scale computing and storage, social media, mobile computing and sensing, combined with advances in algorithms and statistical learning, have enabled the capture,

transmission, aggregation, search, and mining of vast amounts of digital information, and have placed this information at the fingertips of corporations, governments, and individuals. This technology has ushered in the era of Big Data with its fantastic new opportunities for knowledge extraction, optimization, and personalization, but has also created unprecedented new threats to citizens' privacy and freedom. Understanding these threats and devising practical technologies to effectively mitigate them is a key challenge addressed by the institute.

- **Dependable software:** Software systems are fundamental components in many safety- and business-critical applications, such as automotive and avionic control systems, medical devices, energy management systems, and large-scale commercial infrastructure. The software is increasingly complex and heterogeneous, composed from many different components written in different languages at different levels of abstraction, must execute on commodity, multi-core hardware with weak memory models, and must in many cases meet stringent real-time requirements with constrained resources. Developing programming methodologies and verification technologies that enable cost-effective design and verification of software systems remains a key challenge addressed by the institute.

- **Foundations of social computing:** Societal-scale systems like Facebook, Google, Twitter, Amazon, eBay, etc., are rapidly transforming the media landscape, trade, personal and corporate communication, as well as political discourse. In these systems, algorithms trained on users' past behavior increasingly determine what news and information users get to see, who they meet, and what goods and services they are offered at what price. The capability to capture, predict and influence users' behavior, awareness, and opportunities, in the hands of large corporations and governments, raises fundamental questions about freedom, transparency, fairness, bias, and potential discrimination. Understanding these threats and developing appropriate technical means of mitigation while retaining the innovative potential of social computing systems is a key challenge addressed by the institute.

These and other challenges are of fundamental importance to society, and are inadequately addressed by either industrial research (which tends to be focused on new functionality and near-term solutions to emerging challenges) or university research (where it is more difficult to quickly build up significant strengths in emerging areas, especially when cutting across

traditional academic silos).

As a growing research institute in software systems, we emphasize a research environment conducive to long-term, fundamental research on these and other challenges. In particular, we continue to hire faculty who are, individually and as a group, well positioned to address broad challenges in software systems.

1.1.2 Situation

MPI-SWS, one of 83 institutes comprising the Max Planck Society (MPS), was founded in November 2004 and opened its doors in August 2005. The institute has two sites, one located on the campus of Saarland University (UdS), the other on the campus of the Technical University (TU) Kaiserslautern. The sites are 45 minutes apart by car, door-to-door.

Kaiserslautern and Saarbrücken are cities with about 100,000 and 180,000 inhabitants, respectively. The cities offer attractive surroundings and a low cost of living. Access to major metropolitan areas is easy via high-speed rail (two hours to Paris) and low-cost flights from the local airports (Saarbrücken and Luxembourg). Frankfurt airport, the closest international hub, is a 60 minute drive from Kaiserslautern and a 90 minute drive from Saarbrücken.

Several research organizations in related areas are located at the two sites. The computer science department at Saarland University ranks among the top five in Germany. The Max Planck Institute for Informatics (MPI-INF) in Saarbrücken focuses on algorithms, vision and graphics, bioinformatics, databases and information systems. The German Research Center for Artificial Intelligence (DFKI), an applied research lab on artificial intelligence, has locations in both Saarbrücken and Kaiserslautern. The Intel Visual Computing Institute (Intel VCI) in Saarbrücken is a collaborative effort between Intel, MPI-INF, MPI-SWS, DFKI, and UdS. MPI-SWS is part of the Cluster of Excellence on "Multimodal Computing and Interaction" and the newly established Center for IT Security, Privacy and Accountability (CISPA) at Saarland University.

The computer science department at the TU Kaiserslautern ranks in the top quartile of departments in Germany. Kaiserslautern hosts two applied research institutes, the Fraunhofer Institute for Experimental Software Engineering and the Fraunhofer Institute for Industrial Mathematics. There are also a number of information technology startups and a few mid-sized companies at both sites.

MPI-SWS faculty participate in the Cluster of Excellence for Multimedia Computing and Communication (MMCI) at Saarland University, the

Saarbrücken Graduate School for Computer Science, the Intel Visual Computing Institute (Intel VCI) in Saarbrücken, the new Center for IT Security, Privacy and Accountability (CISPA) at Saarland University, and the Kaiserslautern Science Alliance.

The MPI-SWS has a total budget of just over EUR 10M per year and 17 faculty positions (10 of which are currently filled). The institute buildings at the two sites jointly offer space for over 200 researchers and staff. Additional growth is expected through external funding. In this reporting period, the institute received over €2M in external funding.

1.1.3 Mission and strategic goals

The MPI-SWS mission statement reads as follows: "The Max Planck Institute for Software Systems is chartered to conduct world-class basic research in all areas related to the design, analysis, modeling, implementation and evaluation of complex software systems. Particular areas of interest include programming systems, distributed and networked systems, embedded and autonomous systems, as well as crosscutting aspects like formal modeling and analysis of software systems, security, dependability and software engineering."

As an academic institution dedicated to high-risk, long-term research, the primary goal is to have impact primarily through publications, artifacts, and people. We aim to contribute to a stronger and broader base of software systems research in Germany and Europe. In particular, we seek to attract outstanding talent from all over the world, thus broadening the pool of talent in Germany and Europe. In this reporting period, we placed three students in academic and industry research positions in Germany, and two in the rest of Europe. At the same time, we expect our graduates to be competitive for academic and research positions at top universities and laboratories worldwide.

1.1.4 Tenure-track: attracting top talent to the institute

The core principle for achieving our goals is to hire the most talented researchers (including faculty, postdocs, and students) available within our areas of interest. When hiring faculty, this principle trumps, for instance, trying to develop one specific area or another. Hiring the most talented people necessarily means recruiting from all over the world, and in particular means being able to compete with the top CS departments in the USA. This is primarily because US universities house many of the best Software

Systems groups in the world.

For this reason, we have eschewed the traditional German academic model in favor of a faculty model, similar to that of many internationally renowned computer science departments. With the exception of one temporary project-funded (soft-money) researcher, all researchers above the postdoc level are independent faculty: either *tenured* or *tenure-track*. Our tenure evaluation model follows those of US schools: tenure-track faculty are evaluated for tenure during their sixth year. (An internal mid-term review is conducted during the third or fourth year.) Tenure cases are reviewed by a committee appointed by the Chemistry, Physics and Technology (CPT) Section of the MPS and chaired by the vice-president of the CPT.

The tenure-track model enables the institute to compete for top talents internationally, and offers junior faculty a merit-based career path at the institute and thus a stake in its leadership and direction. Tenure-track faculty participate, along with tenured faculty and directors, in most institute-level decisions, including faculty hires, budget allocation, institute policy and other aspects of the institute's academic governance. The only decision in which non-tenured faculty do not participate are decisions to grant tenure.

While common in the USA, the type of flat institute organization engendered by the tenure-track model is fairly unconventional relative to other MPIs and most German universities. However, the Max Planck Society (MPS) grants its individual institutes substantial flexibility in their organizational and research strategy, which has allowed us to adopt this structure.

The model has worked very well for the institute. We have established ourselves as leaders in areas well outside the expertise of the directors. The institute has demonstrated that it can adapt quickly to take on new research challenges in software systems. We have done this both through existing faculty changing their own research agendas (Gummadi and Francis, from networking to social computing and privacy respectively), and through hiring faculty in different but complementary areas (machine learning and, at least briefly, natural language, both complements to social computing). At the same time, we continue to have a healthy rate of churn, with both Clement and Danescu-Niculescu-Mizil leaving in this reporting period (Google Research and Cornell, respectively). Both Gummadi and Dreyer, our "younger" tenured faculty, increasingly take leadership roles in the institute, including mentoring junior faculty and playing a larger role in establishing policy.

A key factor in attracting top talent has been our ability to offer competitive salaries and generous base funding. Compared to top universities, we can offer our faculty permanent base funding for postdocs, doctoral stu-

dents, travel and equipment. This funding has proved quite attractive to junior faculty, as it gets them started early, and reduces the risk of failing to get funding for their chosen research topics. Flexibility in when and what they teach, as well as a reduced teaching load relative to our peer institutes, is also attractive to faculty members.

Our faculty recruiting timeline is aligned with that of US schools. We take applications for faculty candidates in December and January. An internal committee screens the applications and we interview faculty candidates individually in the February through April timeframe. The faculty then nominates their selected candidates to an MPG appointment committee led by the vice-president CPT, which issues a recommendation to the MPS president, who extends the official offers by early May.

1.1.5 The two locations: Kaiserslautern and Saarbrücken

A key challenge of the location split is maintaining a single, unified institute, with as much cross-site collaboration and shared institute culture as possible. We have had in place three primary mechanisms for achieving this, and they have been working well:

- A policy of selecting the site for new faculty that is independent of their research area.
- Frequent cross-site visits and video-conferencing.
- A uniform MPI-SWS graduate program with easy transfer between the two associated universities (UdS and TU-KL).

Our strategy for placing new faculty is to hire for the site with fewer faculty, or if both sites have the same number, to alternate sites. As a general rule, we do not give new faculty a choice of site unless there is a two-body situation that can only be solved at one site, or the new faculty has strong personal reasons. In particular, we neither try to co-locate nor try to separate researchers who work in similar research areas. By "letting the chips fall where they may", we have effectively prevented the institute from splitting into sub-institutes with very different emphases. Instead, we have a healthy mix of researchers working in different areas at both locations, and this seems to be working out well. For example, we have many cross-site collaborations, both within and across disciplines. There are of course pros and cons to any such hiring approach, but by making the choice of site "algorithmic", we have managed to avoid complex and ultimately unresolvable discussions about where best to place new faculty.

Second, to maintain the collaborative environment that is engendered by MPI-SWS's flat structure, it is critical that we effectively bridge the distance gap between the two sites. We use a combination of face-to-face visits and video conferencing.

Each week, all members from one site visit the other site via a chartered full-sized bus, alternating between the two sites. The visits result in weekly face-to-face interactions between members of each site. Indeed, we have found that they tend to foster interaction, both because we feel motivated to exploit the opportunity, and because the visitor is less distracted by the demands of his or her office. There is a well-attended afternoon "tea time" on these visit days where people can mix and socialize. Faculty conduct lunch-time meetings during visit days, where institute matters can be discussed face-to-face. In addition, during the hour-long bus ride between the two locations, people tend to have spontaneous as well as planned meetings, so the time is spent productively. We feel that it may even improve communication among the researchers of one site.

Finally, although visit days partially make up for the lack of contact between institute members at the two sites, they often don't suffice. Connecting the two sites in real time requires state-of-the-art electronic communications equipment. Towards that end, we have set up videoconferencing facilities at both sites, which we use to transmit talks and classroom lectures. The setup supports separate channels for the speaker view, his or her presentation slides, and views of both audiences. The quality of these videocasts is quite good and we are finding this to be an adequate solution.

For one-on-one or small group discussions, we have installed multiple mobile videoconferencing units in both locations, as well as a high-quality Cisco T3 Telepresence system that supports face-to-face meetings with up to 6 participants at each site.

1.2 The state of the institute, and recent accomplishments

In this section, we briefly summarize the state of the institute, as well as some key statistics concerning our research output, notable accolades, etc.

Personnel. During the reporting period, the institute has hired two new tenure-track faculty: Manuel Gomez Rodriguez (applied machine learning and social computing) and Eva Darulova (software reliability and verification). The institute has also recruited 8 new postdoctoral researchers and 16 new doctoral students.

Two junior faculty have also left the institute: Allen Clement (Google Research) and Cristian Danescu-Niculescu-Mizil (Cornell University).

Publications and talks. MPI-SWS has produced 125 peer-reviewed publications during the reporting period. The majority of these are in top-tier conferences (see individual group sections for details). Of these publications, 19 are co-authored by members of two or more groups. In this period we have published in many of the major conferences in our areas, including:

- Programming Languages and Verification: POPL, PLDI, ICFP, OOPSLA, CAV, ECOOP, ESOP, CONCUR
- Distributed, Dependable and Mobile Systems: OSDI, SIGCOMM, NSDI, MobiSys, ASPLOS, CoNext, EuroSys, DISC, IMC
- Real-Time and Embedded Systems: EMSOFT, RTSS, RTAS, ECRTS
- Security and Privacy: CCS, S&P, Usenix Security, CSF, SOUPS
- Social Computing and Social Media Systems: WWW, ICWSM, COSN, CHI, CSCW
- ML, Data Mining, NLP, IR: NIPS, ICML, AAAI, KDD, ACL, EMNLP, CIKM, RecSys

MPI-SWS faculty gave invited talks at 8 conferences and 16 workshops. We collectively served as program chairs or co-chairs for 9 conferences and 7 workshops. MPI-SWS faculty and postdocs have served on the PCs of 75 conferences and workshops. The details are given in the individual group sections.

Awards and honors. Paul Francis's cloaked database system received a TÜViT Trusted Process certification for generalized anonymized analytics. This is to our knowledge the first such certification given for a general-purpose anonymizing analytics system (i.e. independent of use case).

MPI-SWS and Aircloak jointly won the 2014 Cisco IoT Security Grand Challenge competition.

Isabel Valera and Rijurekha Sen were both awarded two-year Humboldt postdoctoral fellowships in March 2015, while Saptarshi Ghosh was awarded a one-year Humboldt postdoctoral fellowship in Ocotber 2014.

Deepak Garg was awarded a Google Faculty Research Award, jointly with Peter Druschel in 2014.

Peter Druschel, Rupak Majumdar, and Michael Backes, jointly with Gerhard Weikum (MPI-INF), received an ERC Synergy Grant in 2014—one of Europe's most distinguished research awards and the only one in computer science so far.

Michael Backes was selected as a member of the German Academy of Science and Engineering (AcaTech) in 2015. In 2014, he received an IEEE Outstanding Community Service Award, the CS Teaching Award of Saarland University for the best CS lecture in 2014, and the Teaching Award of the State of Saarland. Moreover, he was named one of Germany's digital minds by Germany's federal minister of Science and Education Johanna Wanka in 2014.

Rupak Majumdar won the 'Most Influential Paper Award' in January 2014 at POPL 2014.

Institute members won the following 'Best Paper' awards:

- Pramod Bhatotia won the best student paper award with his Middleware 2014 paper "Slider: Incremental sliding window analytics". [42].
- Alexey Reznichenko received a Best Student Paper Award at CCS 2014 for his paper "Private-by-Design Advertising Meets the Real World".
- Mainack Mondal, Bimal Viswanath, and Krishna Gummadi received the Distinguished Paper award at the SOUPS 2014 conference for their work on "Understanding and Specifying Social Access Control Lists".
- Manuel Gomez Rodriguez won an Outstanding Paper Award at the Neural Information Processing (NIPS) Conference in December 2013.

Institute members also won the following Best Thesis awards:

- Aaron Turon received the 2014 ACM SIGPLAN John C. Reynolds Doctoral Dissertation Award for his PhD thesis, *Understanding and Expressing Scalable Concurrency* [173].
- Rijurekha Sen won the ACM-India Best Doctoral Dissertation Award 2014 for her thesis "Different Sensing Modalities for Traffic Monitoring in Developing Regions".
- Sadegh Soudjani won the DIC Best PhD Thesis Award for his thesis "Formal Abstraction for Automated Verification and Sythnesis of Stochastic Systems".

External grant funding. Although the institute provides its faculty members with base funding to run their research groups, we actively encourage all faculty to seek external funding. Securing such funding is important not only in terms of bringing additional resources to the institute, but also in providing junior faculty with grant-writing experience that will be essential for their future careers.

Deepak Garg and UdS faculty Christian Hammer jointly secured external funding from DFG, program RS3 for a period of 4 years. Deepak also won a Google Faculty Research Award.

Krishna Gummadi received an unrestricted grant from AT&T, as well as two Humboldt Postdoctoral Fellowships. IMPECS also awarded him two fellowships and a grant on 'Analysis and Design of Online Social Networks'.

Rupak Majumdar, Peter Druschel, and Michael Backes, together with MPI-INF director Gerhard Weikum, won a prestigious ERC Synergy Award for 'ImPact: Privacy, Accountability, Compliance, and Trust in Tomorrow's Internet'.

Rupak has also been the recipient of a Toyota Research Contract since 2013.

Viktor Vafeiadis will continue until March 2016 to use his European Commission ADVENT Fund.

Derek Dreyer obtained a Microsoft Research PhD Scholarship for David Swasey.

Paul Francis received EXIST funding throughout the reporting period for the spinoff Aircloak.

Michael Backes and Peter Druschel received grants from BMBF for the CISPA center.

Teaching. Teaching is not a formal requirement for institute faculty, but we strongly encourage faculty to regularly teach courses regardless. Teaching is important both in terms of training our doctoral students and ensuring that our faculty are well-prepared for any future positions they may hold at other academic institutions. During this review period, institute faculty taught 9 courses, 2 of them core courses. (For further details, see Section 15.)

Recruiting world-class postdocs. Institute members' base funding includes support for postdoctoral positions. We have strived to make these positions as attractive as possible, by emphasizing to potential candidates the ability to both work on existing institute projects and also develop their own independent research agendas. As a result, we have been able to attract

strong postdocs from around the world.

In this reporting period, the following postdocs have joined our institute:

Oana Goga (LiP6 Paris), Anne-Kathrin Schmuck (TU Berlin), Vinayak Prabhu (Berkeley), Ori Lahav (Tel Aviv University), Sadegh Soudjani (TU Delft), Przemyslaw Grabowicz (University of Palma De Mallorca), Isabel Valera (Carlos II University, Madrid as a Humboldt Fellow), Rijurekha Sen (IIT Bombay as a Humboldt Fellow), Saptarshi Ghosh (IIT Kharagpur as a Humboldt Fellow), and Denzil Correa (IIT Delhi as a IMPECS Fellow),

Placing our postdocs and students. Five students graduated from MPI-SWS during this reporting period. Nuno Santos became an Assistant Professor at IST Lisbon. Pramod Bhatotia joined the TU Dresden as an independent research group leader. Alexey Reznichenko joined the Microsoft Advanced Development Lab in Munich. Bimal Viswanath and Ekin Istemi Akkus have accepted positions at Bell Labs in Stuttgart. Two PhD students took postdoc positions: Pedro Fonseca at the University of Washington, and Beta Ziliani at the National University of Cordoba in Argentina.

In addition, a former student, Matt Hammer, obtained a faculty position at the University of Colorado Boulder after doing a postdoc at the University of Maryland with Michael Hicks.

Among our postdocs, Dario Fiore became an assistant professor at IMDEA Software Institute, Madrid. Aaron Turon joined Mozilla Research.

Part I
Current Research Groups

2 The Real-Time Systems Group

2.1 Overview

The report covers the period from November 2013 – July 2015. The Real-Time Systems Group's efforts are centered on the theoretical foundations and practical challenges of multiprocessor real-time systems. A particular focus is the design and implementation of analytically sound real-time operating systems (RTOSs) that enable and simplify the development of provably predictable applications for safety-critical domains. To this end, the group both investigates the relevant algorithmic foundations (e.g., the design of predictable locking algorithms and static timing analysis methods) and engages in prototyping and system building efforts (mainly in the context of LITMUS$^{\mathrm{RT}}$, the group's Linux-based RTOS).

Personnel. The group is led by Björn Brandenburg and currently consists of four graduate students (Alexander Wieder, Arpan Gujarati, Manohar Vanga, and Felipe Cerqueira), one (non-PhD) graduate research assistant (Mahircan Gül, enrolled at TUKL as a Master's student), and one undergraduate research assistant (Felix Stutz, enrolled at TUKL).

Pedro Fonseca, who was co-advised by Björn Brandenburg, graduated in June 2015. A visiting graduate student, Maolin Yang (UESTC, China), stayed with the group from October 2014 until the end of May 2015 to collaborate on a paper on novel analysis methods for global real-time locking protocols (currently under submission).

Roy Spliet (now at U. Cambridge) interned with the group from September 2013 until May 2014; his project resulted in a paper that was published at RTSS'14 [168]. From May 2014 until August 2014, Yuvraj Patel (formerly of Netapp, moving to U.Wisconsin) and Akshay Aggarwal (IIT Kanpur) joined the group as summer interns.

Darshit Shah (BITS Pilani) visited the group from February 2014 to May 2014 to carry out research towards completion of his undergraduate thesis on the analysis of the read-copy-update (RCU) synchronization primitive in real-time systems. Tobias Blaß (UdS) completed his Bachelor's thesis on the implementation of a new real-time locking protocol in the Linux kernel under the supervision of Björn Brandenburg in March 2015.

Collaborations. The group has active collaborations with Andrea Bastoni and Henrik Theiling of SYSGO AG.[1]

The group is in the process of setting up a collaboration with Sophie Quinton (INRIA Grenoble), Jean-François Monin (Verimag), Jean-Bernard Stefani (INRIA Grenoble), and Rolf Ernst (TU Braunschweig) concerning the mechanized verification of real-time scheduling theory. A joint grant proposal has been submitted and is under review.

There are further early-stage discussions with Rob Davis (U. York), Marko Bertogna (U. Modena), and several of their collaborators concerning an opportunity for joint work on the analysis and evaluation of different scheduling approaches for limited preemptive global scheduling.

Internally, the group has collaborated with the Dependable Systems group. In the Winter semester 2014/2015, Rupak Majumdar and Björn Brandenburg jointly offered a course on the foundations of cyber-physical systems at TUKL.

Publications. In the reporting period (November 2013 – July 2015), group members have published papers in the three major real-time systems conferences RTSS [41, 56, 59, 168, 185, 186], RTAS [60], and ECRTS [55], and in the journals *Real-Time Systems* [113] and *Leibniz Transactions on Embedded Systems* [54].

In collaboration with the Dependable Systems group, the group further published at OSDI'14 [94] and ASPLOS'15 [43].

Systems and tools. LITMUS$^{\mathrm{RT}}$, a real-time extension of the Linux kernel, is the group's primary open-source project and has been continuously maintained since 2006. (As of July 2015, 15 major releases have been made, spanning 29 Linux kernel versions.) LITMUS$^{\mathrm{RT}}$ has been used by scholars in North America, Africa, Asia, and Europe, and has served as the basis of more than 45 publications and eight PhD and MS theses.[2]

The group's second major open-source software project is SchedCAT, a toolkit for schedulability and blocking analysis, which makes the group's analytical contributions available in runnable form.

In the reporting period, the group heavily used and developed both LITMUS$^{\mathrm{RT}}$ and SchedCAT, in the context of the research activities related to publications [56, 60, 168] and [55, 59, 60, 113, 185], respectively.

[1] SYSGO AG is headquartered in Mainz (about one hour northeast of Kaiserlautern) and is a major real-time operating system vendor for safety critical systems and avionics.

[2] The list of publications is available on the LITMUS$^{\mathrm{RT}}$ homepage: `https://wiki.litmus-rt.org/litmus/Publications`.

Teaching. Björn Brandenburg offered a seminar on "Operating System Design and Implementation" (Winter 2013/2014, with Allen Clement) and co-taught a lecture course on "Foundations of Cyber-Physical Systems" (Winter 2014/2015, with Rupak Majumdar). In December 2013, Björn Brandenburg taught an intensive, one-week lecture course on multiprocessor real-time synchronization at Scuola Superiore Sant'Anna, Pisa, Italy.

Service. Internally to MPI-SWS, Björn Brandenburg served on the graduate admissions committee in 2015, and chaired said committee in 2014. Since Fall 2014, he is responsible for screening incoming post-doc applications. In April 2014, he organized and carried out the institute's participation in the Kaiserslautern science night ("Nacht, die Wissen schafft").

Within the Max Planck Society, Björn Brandenburg has served throughout 2014 as the Germany-wide elected representative of the *Max Planck Research Group Leaders* (MPRGLs) in the *Chemical-Physical-Technical Section* (CPTS) of the Max Planck Society, which involved attending bimonthly meetings of the section's planning committee ("Perspektivenkommission") and two section meetings. In October 2014, he co-organized and directed the annual MPRGL meeting in Berlin.

Externally, Björn Brandenburg has served as program chair of the RTAS 2015 *Applications, RTOSs and Run-Time Software and Tools* track, as program co-chair of the OSPERT'14 and OSPERT'15 workshops, and as program chair of the RTAS 2014 Work-in-Progress track. In March 2015, he served as local organizer for the ECRTS'15 program committee meeting, which was hosted at MPI-SWS.

He served on the program committees of RTAS'14, ECRTS'14, ECRTS'15, EMSOFT'15, RTNS'14, SAC'15, and ETFA'14, and as a reviewer for the journals *Real-Time Systems*, *Journal of Systems Architecture*, *ACM Computing Surveys*, *Software Practice & Experience*, and *IEEE Transactions on Industrial Informatics*, and for the German science foundation (DFG).

2.2 Research agenda

The group's research activities in the reporting period were centered around three focus areas: real-time scheduling with affinity constraints, RTOS implementation issues, and lock-based real-time synchronization.[3] The first two areas are discussed in some detail; the latter area is summarized briefly.

[3]The overview assumes a some familiarity with real-time systems. A high-level introduction to the group's general area of research was provided in the 2013 progress report.

Affinity constraints. The literature on multiprocessor real-time scheduling to date is focused primarily on well-structured approaches such as partitioned, global, and clustered scheduling. However, commercial RTOSs such as QNX and VxWorks — and also mainstream general-purpose OSs such as Linux, FreeBSD, and Windows — do not actually implement these specific approaches. Instead, they implement a more flexible mechanism that permits each task to have an *arbitrary processor affinity* (APA), which is simply a user-specified set of processors on which a task may be dispatched.

Prior work had considered processor affinities to be a mere implementation detail of little consequence from an analytical point of view. In 2013 (prior to the reporting period), the Real-Time Systems group at MPI-SWS was the first to realize [112] that APA scheduling actually strictly dominates global, clustered, and partitioned scheduling combined, with respect to the class of *job-level fixed-priority* (JLFP) scheduling policies, and proposed an initial analysis.

In the reporting period, the group has continued to investigate APA scheduling from an analytical perspective. In particular, it has extended the generality result to *job-level dynamic-priority* (JLDP) policies, where global and APA scheduling are equally powerful [113], and developed a novel *schedulability analysis* (i.e., static analysis to determine if all deadlines will be met at runtime) for the type of fixed-priority APA schedulers found in Linux, QNX, VxWorks, and most other current multiprocessor RTOSs [113].

The key challenge, compared to the analysis of classic approaches such as global or clustered scheduling, is that higher-priority tasks can potentially interfere with (i.e., delay) lower-priority tasks on some, but not all processors, which makes it difficult to devise effective closed-form bounds on total higher-priority interference during a job's scheduling window (i.e., the interval from a job's arrival to its deadline). To sidestep the problem, Gujarati et al. [113] devised a new response-time analysis method based on linear programming that does not rely on closed-form aggregate interference bounds. Rather, the problem is expressed as a linear optimization problem, where the objective is to maximize the analyzed task's response time subject to constraints that restrict higher-priority workload to specific cores.

In the linear program (LP) formulation, however, concrete workload bounds are required, which are themselves dependent on the analyzed task's response time. A circular dependency thus exists, which is resolved, as in classic response-time analysis, with a fixed-point search (where an LP is generated and solved in each iteration). The new method [113] is directly applicable to current RTOSs and was shown both to analytically dominate the prior analysis [112] and to be fast enough to be practical, despite the

many invocations of an LP solver during the fixed-point search.

Complementing the work on APA schedulability analysis, a collaboration with Sanjoy Baruah (UNC Chapel Hill) yielded a *feasibility test* for APA scheduling [41], which is a kind of decision procedure that is useful for assessing the precision of schedulability tests. Most multiprocessor schedulability analyses, including the one developed by Gujarati et al. [113], are only *sufficient* but not *necessary*, in the sense that they may wrongly identify some actually schedulable task sets as being at risk of missing a deadline (i.e., such tests are safe, but not precise). In contrast, a necessary feasibility test identifies all workloads that can be successfully scheduled by *some* scheduling policy, but does not shed light on whether a given workload can be supported by a particular policy (i.e., such tests are precise, but not useful safety arguments). A necessary feasibility test, such as the one developed by Baruah and Brandenburg [41], thus establishes an objective ground truth and provides a measure of the limits of current scheduling and analysis techniques, i.e., it can be used to quantify the gap between currently unsupported and truly infeasible workloads.

In work aiming to narrow this gap, Cerqueira et al. [59] investigated migration rules for APA schedulers and found that current implementations sometimes make pathological task placement choices, which can result in needless deadline misses (and thus contributes to the feasibility-vs.-schedulability gap). The current APA schedulers in Linux, QNX, etc. never reassign (or *shift*) high-priority tasks from one core to another core to make room for a lower-priority (but still urgent) task with a more-constraining affinity. In other words, once a high-priority task has commenced execution, it cannot be "dislodged" by a lower-priority task with a more-constrained affinity, even if the higher-priority task could be scheduled just as well elsewhere. Cerqueira et al. [59] observed that it can be beneficial to shift higher-priority tasks in some cases, but — if implemented as in a typical RTOS scheduler — it is far from obvious how one should determine which task to shift, where to shift it to, and when to shift it.

Modeling the problem instead as a *maximum vertex value matching problem (MVM)*, Cerqueira et al. [59] found an isomorphic problem that had been studied previously in the operations research community in the context of assigning nurses to open positions in a hospital under consideration of both per-nurse skill sets and preferences ($\approx$ affinity constraints), and *weak* or *strong* employee seniority constraints ($\approx$ task priorities).

Analogously, Cerqueira et al. [59] found that the schedulers implemented in current RTOSs are only *weak* APA schedulers, whereas *strong* APA schedulers that rely on shifting can schedule a substantially broader range of real-

time workloads (both empirically and analytically). Because existing MVM algorithms are unsuitable for a kernel environment, Cerqueira et al. proposed a simpler strong APA scheduler and matching schedulability analysis [59].

RTOS implementation. Whereas the work on APA scheduling has focused to date primarily on analytical issues, the group also continued to investigate practical systems issues in the context of the LITMUS$^{\mathrm{RT}}$ project.[4] In particular, Cerqueira et al. [60] studied the efficiency of global schedulers.

A global multiprocessor scheduler must at all times maintain a global invariant, such as "given m processors, always schedule the up to m ready tasks with the earliest deadlines." To satisfy the invariant, tasks may be migrated freely as needed.

Global approaches are a convenient choice in dynamic environments because they do not require manual load-balancing operations when the workload changes (i.e., if tasks join or leave the system, or if the execution requirement or activation frequency of a task changes) because a work-conserving global scheduler implicitly load-balances at each scheduling decision. This is especially useful if no processor affinity has been specified by the user; Linux and many RTOSs such as QNX and VxWorks thus default to global scheduling unless explicitly overridden with APAs. However, global schedulers unfortunately also suffer from high runtime overheads due to the need to coordinate all processors and to maintain a consistent global state.

Fundamentally, when implementing a global policy, there is a tension between correctness (always making the right decision) and efficiency (deciding quickly). The baseline global scheduler in LITMUS$^{\mathrm{RT}}$ uses a shared ready queue protected by a single spin lock that serializes all scheduler invocations, which greatly helps to ensure correctness of the scheduler, but comes with obvious scalability limitations. In contrast, the Linux scheduler uses per-processor ready queues and explicit load-balancing steps, called "push" and "pull" operations, to maintain the global scheduling invariant, which ostensibly lowers overheads. However, the Linux scheduler fails to actually maintain the global scheduling invariant in all cases due to certain pathological (but not so improbable) race conditions [52], which renders it analytically unsound and undermines all schedulability analysis.

Attempting to reconcile the two extremes of "correct but unscalable" (LITMUS$^{\mathrm{RT}}$) and "scalable but wrong" (Linux), the group investigated various approaches for implementing global schedulers more efficiently without sacrificing correctness. After exploring several dead ends of successively in-

[4] `http://www.litmus-rt.org`

creasing complexity — based on evermore fine-grained locking, reader-writer locking, and ultimately lock- and wait-free solutions — a much simpler solution based on message passing was found [60].

In a real-time context, the problem with Linux-like per-processor state and other "commonsense scalability techniques" (like adopting lock-free designs) is that, while such techniques are indeed effective at reducing *average*-case overheads, they fail to reduce *peak contention*—in the *worst* case, all cores are still contending for the same datum, be it explicitly for a single lock (as in the LITMUS$^{\mathrm{RT}}$ baseline), explicitly for a per-processor lock (as in Linux), or implicitly for cache lines (as with CAS-based solutions). In fact, all tested designs based on fine-grained synchronization, and in particular the Linux scheduler, exhibited (far) worse maximum overheads than the "naïve" baseline design in LITMUS$^{\mathrm{RT}}$, thus debunking the (at least in the real-time community) widely-held belief that fine-grained synchronization performs better: it is better only in the average case (i.e., in terms of throughput), but not with regard to the worst case, which favors simplicity over intricate synchronization strategies.

Embracing this observation, Cerqueira et al. [60] designed a simple solution based on message passing that significantly lowered *both* average- and worst-case overheads. To avoid cache-related overheads, they centralized all scheduler state onto a single master core (thus avoiding cache-line bouncing) and used simple, cache-friendly mailboxes and inter-processor interrupts to notify client cores of scheduling decisions. The new implementation was shown to scale well up to 64 cores, achieving much lower worst-case overheads than both the baseline design (up to 23x reduction) and Linux's scheduler (up to 36x reduction), and achieving average-case overheads comparable to Linux (a >100x improvement over the baseline design) [60].

In another project aimed at reconciling average- and worst-case performance, Spliet et al. [168] investigated how to extend Linux's futex (i.e, "fast userspace mutex") mechanism to various predictable real-time locking protocols. A futex exploits the observation that, if a lock is uncontested, then it can be acquired and released without invoking (costly) system calls since no tasks need to be blocked or resumed. As most lock acquisitions tend to be uncontested (in a well-designed system), futexes can offer great savings in synchronization-intensive workloads. However, prior work on futex mechanisms had considered only *reactive* real-time locking protocols, which take effect only once a problematic resource conflict exists. While this greatly simplifies the futex implementation (there is nothing to do in the absence of contention), many more advanced real-time locking protocols use *anticipatory* techniques, in the sense that they prevent problematic contention in the

first place (e.g, by adjusting the priority of lock-holding tasks), which cannot be simply omitted in all uncontended cases (as problematic contention may still arise after the lock has already been acquired). To work around this limitation, Spliet et al. [168] devised simple techniques for realizing anticipatory futexes and showed them to be effective with an evaluation in LITMUS$^{\mathrm{RT}}$.

Real-time synchronization. The group has continued its well-established line of synchronization work with several contributions concerning the foundations of real-time synchronization, practical protocol design, and static blocking time analysis methods.

Regarding the former, Wieder and Brandenburg [186] showed that the worst-case blocking analysis (WCBA) problem on multiprocessors in the presence of nested critical sections is NP-hard for both priority- and FIFO-ordered wait queues, and regardless of whether waiting is implemented by busy-waiting (i.e., spin locks) or suspending (i.e., semaphores), which came as a surprise since polynomial-time solutions exist for the WCBA problem both on uniprocessors and, if nesting is disallowed, also on multiprocessors. Brandenburg further extended the theory of blocking optimality to distributed real-time locking protocols [54] and devised the FMLP^{+}, the first real-time locking protocol for clustered JLFP scheduling that is asymptotically optimal under suspension-aware analysis [55]. (The FMLP^{+} is also of practical interest and implemented in LITMUS$^{\mathrm{RT}}$.)

Motivated by the inclusion of spin locks in the AUTOSAR OS standard (pervasive in the automotive industry), Wieder and Brandenburg [185] developed novel WCBA techniques for various types of spin locks, including spin locks with weak or no progress guarantees (as permitted by AUTOSAR) that were previously considered "too unpredictable" to analyze.

Most recently, Brandenburg [56] designed and implemented in LITMUS$^{\mathrm{RT}}$ a synchronization scheme with complete temporal and logical isolation to enable safe and secure resource sharing among an unknown number of mutually untrusting tasks, a core requirement of practical mixed-criticality systems.[5]

Future work. As projects in each of the three summarized topic areas are still very much in progress, the group plans to continue its research in the next two years along the current directions.

[5] *Mixed-criticality systems* are systems in which, for cost and efficiency reasons, safety-critical components subject to strict certification requirements are co-hosted with less trusted, less critical components — a topic of rapidly growing practical relevance.

With regard to the topic of real-time synchronization, the group plans to explore algorithmic questions along two principal directions. First, to date there still does not exist a practical WCBA for common lock types on multiprocessors in the presence of nested critical sections, an obvious shortcoming of the state of the art that ongoing work is attempting to address. Second, the group plans to revisit migratory priority inheritance, a technique to ensure lock-holder progress without negatively affecting the scheduling latency incurred by higher-priority tasks. Migratory priority inheritance has been analyzed previously in a suspension-oblivious setting [53]; the more challenging suspension-aware case remains open, however.

Concerning APA scheduling, the project is expected to shift to a more systems-heavy focus, with specific plans to investigate (in LITMUS$^{\mathrm{RT}}$) implementations of weak and strong APA schedulers, as well as partitioned and semi-partitioned alternatives.

Further system building work is planned in the context of the collaboration with SYSGO AG. On the basis of SYSGO's PikeOS, a certified commercial RTOS for avionics, the goal is to investigate various approaches for implementing earliest-deadline first (EDF) scheduling, and to evaluate and contrast different implementations from the point of view of "certifiability."

In another ongoing project with a systems focus, the group is looking at the scheduling of virtual machines in cloud environments such as Amazon's EC2 from a hybrid real-time and datacenter point of view, continuing the theme of reconciling good average-case performance (i.e., high throughput) with predictable worst-case guarantees (i.e., strict latency bounds).

3 The Foundations of Programming Group

3.1 Overview

This report covers the period November 2013 – July 2015. The research of this group focuses on the design, semantics, verification and implementation of modern programming languages and systems, with a particular emphasis on the importance of modularity in designing and reasoning about programs. The three major areas we have focused on in the current review period are (1) developing program logics for modularly verifying challenging concurrent code, (2) verifying compiler correctness compositionally, and (3) designing a new package system for Haskell.

Personnel. The group is led by **Derek Dreyer**, who joined the institute in January 2008 and received tenure in 2013. It currently includes five PhD students: **Georg Neis**, **Scott Kilpatrick**, **David Swasey** (co-advised by Deepak Garg), **Ralf Jung**, and **Jan-Oliver Kaiser**. During the review period, the group also included PhD student **Beta Ziliani**, postdoc **Aaron Turon**, and intern **Joseph Tassarotti** (currently a PhD student at CMU). In March 2015, Ziliani successfully defended his PhD thesis, *Interactive Typed Tactic Programming in the Coq Proof Assistant* [191]; he is now a postdoctoral researcher at the National University of Cordoba in Argentina. In April 2014, Turon left for a position in the core team developing the Rust programming language at Mozilla Research, where he has since been promoted to Research Engineering Manager. In September 2015, Neis will be joining Google Munich as a software engineer on the V8 team.

Collaborations. During the review period, the group has engaged in successful collaborations with a number of leading researchers in Europe, the U.S., and Asia, including: Lars Birkedal, Filip Sieczkowski, and Kasper Svendsen (Aarhus University), Lindsey Kuper and Ryan Newton (Indiana University), Simon Peyton Jones and Simon Marlow (Microsoft Research and Facebook), and Chung-Kil Hur, Jeehoon Kang, and Yoonseung Kim (Seoul National University). The group also collaborates actively with fellow MPI-SWS faculty Viktor Vafeiadis and Deepak Garg.

Publications. The group publishes regularly in the top conferences and journals in the field of programming languages. During the review period, group members have co-authored five conference papers (three POPL [125, 132, 121], one PLDI [172], and one OOPSLA [176]) and one article in the Journal of Functional Programming (JFP) [163]. Group members also have two papers accepted to the upcoming ICFP 2015 conference [151, 193], and one paper to appear shortly in a special issue of JFP [192].

Awards and Honors. Turon received the 2014 ACM SIGPLAN John C. Reynolds Doctoral Dissertation Award for his PhD thesis, *Understanding and Expressing Scalable Concurrency* [173]. Part II of his thesis concerns our joint work published in POPL'13 [175] and ICFP'13 [174].

Software. The group has published two papers on concurrent program logics (GPS [176] and Iris [121]), both of which come equipped with soundness proofs for the logics fully mechanized in the Coq proof assistant. These are freely available from the MPI-SWS PLV (programming languages and verification) group web page (`http://plv.mpi-sws.org`). In addition, in joint work with Jeehoon Kang, Yoonseung Kim, and Chung-Kil Hur at Seoul National University, we have developed SepCompCert, an extension to the verified CompCert C compiler which verifies correctness of separate compilation. We have submitted a paper about SepCompCert to POPL'16 [123]; the Coq code is available online at `http://sf.snu.ac.kr/sepcompcert`.

Teaching. In Winter 2014-15, Dreyer led a graduate *privatissimum* (like a small seminar) on "Categorical Logic". Ziliani also taught an informal course on Coq using Ssreflect, and Jung is currently teaching an informal course on the Rust programming language. In the upcoming Winter 2015-16 semester, Dreyer will be co-teaching the core course on "Semantics" with Prof. Gert Smolka of Saarland University.

External funding. The research of the group is partially funded by fellowships from Google and Microsoft Research. Georg Neis was awarded a 2012 Google European Doctoral Fellowship for his thesis project on "Compositional Multi-Language Reasoning". Dreyer was awarded a 2013 Microsoft Research PhD Scholarship for a project on "Compositional Verification of Scalable Joins by Protocol-Based Refinement", which is being used to fund David Swasey.

Dreyer has applied for a 2015 ERC Consolidator Grant, with the proposal "RustBelt: Logical Foundations for the Future of Safe Systems Programming". The proposal concerns the development of solid formal foundations for the safety of the Rust programming language and its standard libraries. It has made it past Step 1 of the two-step ERC evaluation process; Step 2 involves an in-person interview in Brussels on October 22, 2015.

Invited talks. Dreyer was an invited speaker at PLMW 2014 (the Programming Languages Mentoring Workshop at POPL), OPLSS 2014 (the Oregon Programming Languages Summer School, and the 2014 Workshop on Certification of High-Level and Low-Level Programs at the Institut Henri Poincaré in Paris.

Service. Internally, Dreyer served as chair of the MPI-SWS faculty recruiting committee in 2015.

Externally, Dreyer was elected and served as "awards czar" on the ACM SIGPLAN Executive Committee from July 2012 to June 2015. This entailed overseeing the management of all SIGPLAN awards, and chairing the committees for test-of-time awards for the four major SIGPLAN conferences.

Dreyer was lead organizer of the Dagstuhl Seminar on *Compositional Verification Methods for Next-Generation Concurrency*, which took place in May 2015. He also served as co-chair for the 2014 Coq Workshop and co-organizer of the 2015 PLMW workshop. He served on the program committees for ICFP 2014, ESOP 2015, and MFPS 2015, as well as the external review committees for POPL 2015 and PLDI 2015. He is currently serving as co-organizer of the 2016 PLMW workshop, as well as on the selection committee for the ICFP 2015 student research competition and the external review committee for POPL 2016.

In 2014, Dreyer was invited to join the editorial board of the Journal of Functional Programming, as well as the IFIP Working Group 2.8 on Functional Programming.

3.2 Research agenda

During the review period, the group has continued to pursue several distinct lines of work under the overarching theme of modular programming and verification. These include:

- **Backpack** (POPL'14) [125]: a new package system for Haskell building on our previous work on mixin modules
- **Iris** (POPL'15) [121]: a new unifying foundation for concurrent separation logics and reasoning about logical atomicity
- **GPS** (OOPSLA'14, PLDI'15) [176, 172]: a modern concurrent separation logic for weak memory, in particular the release-acquire semantics provided by the C/C++11 memory model
- **Pilsner** (ICFP'15) [151]: the first multi-pass compiler for an ML-like language to be compositionally verified (in Coq)
- **SepCompCert** (under submission) [123]: a lightweight approach to verifying separate compilation, which we have applied successfully to the full CompCert 2.4 compiler

We briefly describe these projects here, and then conclude by describing our initial work on the **RustBelt** project that is the subject of our ERC grant proposal (currently under review).

Backpack: Retrofitting Haskell with strong modularity. Most prior work on module system design has taken a clean-slate approach, ignoring the practical concern of how to integrate stronger support for modular programming into weakly modular languages.

In this project, which is joint work with Simon Peyton Jones and Simon Marlow (lead developers of GHC, the leading Haskell compiler), and which we described in the previous institute report, we develop Backpack [125], a new language for retrofitting a weak module system like Haskell's with separately typecheckable *packages*. The design of Backpack is inspired by the MixML module calculus [162] that members of our group developed in earlier work, but differs significantly in detail. In particular, Backpack is motivated less by foundational concerns (developing a minimal core calculus of mixin-modular constructs) and more by the practical problem of integrating support for separate modular package development into Haskell.

Ongoing work. Our initial paper on the design of Backpack [125] will serve as the foundation of Scott Kilpatrick's PhD thesis, which is expected to be completed by the end of 2015. The thesis will also include ongoing work on extending Backpack with support for type classes, but it is essentially a theoretical thesis on the formal foundations of the Backpack design. We (Kilpatrick, Dreyer, and Peyton Jones) have also been supervising Edward Yang (a PhD student of David Mazières at Stanford) on a separate project to develop an actual working implementation of Backpack for GHC.

Iris: Monoids and invariants as an orthogonal basis for concurrent reasoning. The last decade has seen a proliferation of formal tools for grappling with the complexity of concurrent programs. In terms of compositional verification in particular, many *separation logics* have been proposed, but without a clear understanding of how they relate to one another.

In this project, which is led by PhD students Ralf Jung and Dave Swasey, and joint with Lars Birkedal's group at Aarhus, we develop Iris [121], a concurrent separation logic that unifies and strengthens the proof principles of existing logics. The core idea of Iris is to decompose concurrent reasoning into an "orthogonal basis", consisting of: (1) *partial commutative monoids*, which are used to represent various forms of "ghost state", *e.g.,* for describing permissions or trace information, and (2) *invariants* on shared state.

We show how these two mechanisms can be fruitfully combined to express shared-state *protocols*, which have shown up in many different guises in previous logics as a way of controlling interference between threads, and to encode the inference rules of those previous logics as derived rules within Iris. We also show how monoids and invariants can be used to encode a notion of *logical atomicity*, thus enabling one to internalize linearizability within the logic.

Ongoing work. Our ongoing work with Iris primarily concerns its use in the RustBelt project, described at the end of this section.

GPS: A modern concurrency logic for C/C++11. Nearly all prior work on concurrency verification (including Iris) assumes a sequentially consistent (SC) model of concurrency. This is unrealistic: modern architectures support weak (relaxed) memory models, and languages like C/C++11 and Java expose these memory models to programmers so that they can write more efficient concurrent code.

In a project led by postdoc Aaron Turon, and joint with Viktor Vafeiadis, we develop GPS [176], one of the very first separation logics (or formal tools of any kind) for compositionally verifying concurrent programs written using the "release-acquire atomics" provided by the C/C++11 weak memory model. GPS improves on earlier work of Vafeiadis and Narayan on RSL [178] by incorporating support for several "modern" concurrency-logic features that are crucial for verifying more complex programs. In particular, like Iris and other advanced concurrency logics, GPS supports *ghost state* via partial commutative monoids, as well as *protocols* on shared state. GPS differs from prior concurrency logics, however, in that its protocols are restricted to govern only a single memory location at a time. Intuitively, this restriction makes sense because under weak memory there is not a single globally consistent view of the machine state, so one cannot in general impose invariants on the state of multiple locations simultaneously; cache coherence does, however, ensure soundness of single-location protocols.

We have applied GPS to a number of challenging case studies, the most significant being the first verification of a weak-memory implementation of RCU (read-copy-update), a key synchronization mechanism deployed heavily in the Linux kernel [172].

Ongoing work. GPS is only a first step toward building usable tools for compositional verification of C11 programs. One direction of ongoing work is to develop a more canonical or general set of proof rules, which we could use to give general modular (*e.g.,* linearizability-like) specifications for higher-

level weakly consistent data structures. Another direction is to reorient GPS to target Vafeiadis *et al.*'s recent SRA (strong release-acquire) model (currently under submission). Unlike C11's release-acquire, SRA admits a simple operational definition while still remaining efficiently implementable. We believe it should be possible to give a much simpler and more easily adaptable soundness proof for GPS against such an operational model. Jan-Oliver Kaiser, PhD student in our group, is exploring these directions.

Pilsner: Compositional compiler verification for ML-like languages. Compiler verification is essential for the construction of fully verified software, but most prior work (such as Leroy's CompCert compiler [137]) has focused on verifying whole-program compilers. To support separate compilation and to enable linking of results from different verified compilers, it is important to develop a compositional notion of compiler correctness that is *modular* (preserved under linking), *transitive* (supports multi-pass compilation), and *flexible* (applicable to compilers that use different intermediate languages or employ non-standard program transformations).

In this work, which is led by PhD student Georg Neis, and which is joint work with Chung-Kil Hur (former postdoc in our group) and Viktor Vafeiadis, we develop Pilsner [151], the first multi-pass compiler for an ML-like (higher-order imperative) language to be compositionally verified (in Coq). The key building block in our verification of Pilsner is a novel simulation technique called *parametric inter-language simulations (PILS)*, which builds on our earlier work on parametric bisimulations published in POPL'12 [118]. PILS inherit many of the benefits of the Kripke logical-relations methods we have studied previously [81, 117] (namely, their modularity and flexibility), while in addition being transitive.

The Pilsner verification was an extensive, multi-person-year effort, involving 55K lines of Coq code. It will be the cornerstone of Georg Neis's PhD thesis, which we expect to be completed in the early fall of 2015.

SepCompCert: Lightweight verification of separate compilation. Most existing work on compositional compiler verification has focused on establishing strong compositionality guarantees—*e.g.,* the ability to link results of *different* verified compilers, possibly together with hand-written assembly code. However, establishing such strong guarantees seems to require a large and complex proof effort. This is certainly true for our own Pilsner verification, as well as for recent work by Appel's team at Princeton on Compositional CompCert [169], a compositional version of (part of) the

CompCert compiler.

In this work, which is joint with Chung-Kil Hur's team at Seoul National University as well as Viktor Vafeiadis, we show that if we aim a little lower, we can do a lot better. That is, if we target a somewhat more restrictive notion of compositional compiler correctness—namely, in which we only consider the linking of results from variants of the *same* verified compiler—then we can port an existing proof of correctness for whole-program compilation to a compositional proof of correctness for separate compilation, and we can do so with relatively minimal effort. In particular, we develop SepCompCert [123], a compositional version of the CompCert 2.4 compiler. The key idea behind SepCompCert is that, when all modules in a program are compiled by the same compiler, one can recast the compositional verification problem in terms of whole-program simulations, thus enabling reuse of existing whole-program compiler proof efforts. Our porting of CompCert to SepCompCert took only two person-months to complete, it only extended the size of the CompCert verification by 3%, and it accounts for the entire CompCert compiler. The result is the first verification of separate compilation for the full CompCert compiler.

Ongoing work. We are currently investigating whether the approach taken by SepCompCert is even more general than we originally thought. We believe it may in fact be possible to strengthen SepCompCert's verification to support the same level of compositionality offered by Compositional CompCert, but with a much lower verification effort.

RustBelt: Logical foundations for the Rust programming language. This section briefly describes the RustBelt project, which we have just begun and which is the subject of our ERC grant proposal (currently under review).

A longstanding question in the design of programming languages is how to balance *safety* and *control*. C-like languages give programmers low-level control over resource management at the expense of safety, whereas Java-like languages give programmers safe high-level abstractions at the expense of control.

Rust is a new language developed at Mozilla Research that marries together the low-level flexibility of modern C++ with a strong "ownership-based" type system guaranteeing type safety, memory safety, *and* data race freedom. As such, Rust has the potential to revolutionize systems programming, making it possible to build software systems that are safe by construction, without having to give up low-level control over performance.

Unfortunately, none of Rust's safety claims have been formally investigated, and it is not at all clear that they hold. (In fact, in the past few months since we began this project, several soundness flaws in the Rust type system have been uncovered.) To rule out data races and other common programming errors, Rust's core type system prohibits the aliasing of mutable state, but this is too restrictive for implementing some low-level data structures. Consequently, Rust's standard libraries make widespread internal use of `unsafe` blocks, which enable them to opt out of the type system when necessary. The hope is that such `unsafe` code is properly encapsulated behind safe interfaces, so that Rust's language-level safety guarantees are preserved. But due to Rust's reliance on a weak memory model of concurrency, along with its bleeding-edge type system, verifying that Rust and its libraries are actually safe will require fundamental advances to the state of the art.

In the RustBelt project, which we are pursuing in collaboration with Aaron Turon and Niko Matsakis at Mozilla Research, we aim to develop the first formal tools for verifying safe encapsulation of `unsafe` Rust code. To achieve this goal, we will build on our prior work on concurrent program logics and semantic models of type systems. In particular, PhD student Ralf Jung has already begun to define λ_{Rust}, a core calculus that captures the key elements of the Rust type system, and to define a Kripke model of λ_{Rust} types which interprets them as predicates in our recently developed Iris logic. We will then use this model to (a) verify soundness of the safe fragment of λ_{Rust}, and (b) verify that `unsafe` Rust libraries (hand-ported to λ_{Rust}) inhabit the semantic interpretation of their interfaces. Eventually, a key component of the project will be to retarget this semantic model to a program logic combining aspects of Iris and GPS, so that we can validate safety of the essential Rust libraries (like `Arc` and `Channel`) that make use of C11-style relaxed memory accesses.

4 The Distributed Systems Group

4.1 Overview

This report covers the period from Nov 2013–July 2015. The group's research during this period has focused on the following areas: 1) Compliance with data data confidentiality and integrity policies in distributed data processing environments (in collaboration with Deepak Garg's group); 2) privacy for mobile social applications and private digital capture (in collaboration with Bobby Bhattacharjee at the University of Maryland); 3) traffic-analysis resistant anonymity for VoIP calls; and 4) studying active attacks.

Personnel. The group is led by Peter Druschel and currently has six graduate students (Paarijaat Aditya, Eslam Elnikety, Viktor Erdelyi, Raul Herbster, Aastha Mehta, Anjo Vahldiek). Peter co-advises post-doc Rijurekha Sen with Krishna Gummadi. Stevens Le Blond (previously a post-doc in Paul Francis' group) joined the group as a research scientist in Feb 2012. Cedric Gilbert has joined the group as a research support engineer, and William Caldwell continues to work for the group as a research support engineer on a freelance basis. Prof. Lorenzo Alvisi (University of Texas at Austin) has visited the group (and the institute) again in the summers of 2014 and 2015; he has been supported by a Humboldt Research Award. Prof. Bobby Bhattacharjee (University of Maryland, College Park) visited the group and institute in July 2015.

Collaborations. Internally, the group has collaborated with the groups of Michael Backes, Rodrigo Rodrigues, Deepak Garg, Krishna Gummadi, and Paul Francis. Externally, we have worked with colleagues at the University of Maryland, Duke University, Yale University, the University of Pennsylvania, the University of Washington, Northeastern University, Microsoft Research Cambridge, the Max Planck Institute for Informatics, and Saarland University.

Publications. Group members have co-authored papers that appeared in SIGCOMM [135], Mobisys [10], Usenix Security[136, 50], ACNS [23, 167], and Eurosys [179]. Workshop papers appeared in SPME 2014 [8] (reprinted in Mobile Computing and Communications Review [9]), and USEC 2013 [147]. Two papers are under submission [83, 11].

Teaching. Peter Druschel taught the core course on Operating Systems in 2013 (jointly with Bjoern Brandenburg) at Saarland University, and taught a core course on Distributed Systems in 2014/2015 (jointly with Paul Francis) at both Saarland University and the TU Kaiserslautern.

External funding. Druschel is a co-PI in the successful renewal of both Saarland University's MMCI Cluster of Excellence and the Saarbrücken Graduate School in Computer Science, funded by the German National Science Foundation (DFG). He is also a co-PI and assistant director of the Center for Information Security, Privacy and Trust, funded by the German ministry of science (BMBF). Peter Druschel, Rupak Majumdar, Michael Backes, and Gerhard Weikum (MPI for Informatics) are co-PIs on an ERC Synergy Grant on Privacy, Accountability, Compliance, and Trust in the Internet (EUR9.25M, 2015–2021). Jointly with Deepak Garg, Druschel won a Google Research Award in 2014.

Awards and invited talks. Druschel gave invited/keynote talks at the Technicolor/INRIA Workshop on Storage and Cloud Computing in Rennes, Nov 2013, at the 2nd Workshop on Accountability: Science, Technology, and Policy at MIT, Jan 2014, at the NICTA Software Systems Summer School in Sydney, Feb 2014, at the Microsoft Research Devices and Networking Conference in Paris, May 2015, at the 8th ACM International Systems and Storage Conference (SYSTOR) in Haifa, May 2015, and at the Technion Computer Engineering Center Conference in Haifa, June 2015.

Service. Within the MPS, Peter Druschel continues to serve on the strategy committee (Perspektivenkommission) of the Chemistry, Physics, and Technology section, and the selection panel of the joint Fraunhofer/Max Planck research program. He served on the MPS Committee on Information Technology (BAR) through June 2014. Currently, Druschel also serves on two presidential committees of the MPS: The committee on the Support of Junior Scientists, and the committee on IT Security. Lastly, Druschel co-organized a Symposium on Foundations of Security and Privacy for the CPT Section of the MPS in July 2015, part of an initiative that could lead to the creation of a new MPI on this subject.

Peter Druschel contines to serve on the Technical Advisory Boards of Microsoft Research, Cambridge and Microsoft Research, India, the scientific committee of the Laboratory on Information, Networking and Communication Sciences (LINCS), Paris, and on the steering committee of the Eu-

roSys/INRIA Winter School on Hot Topics in Distributed Systems (HTDC). He served on the steering committee of the ACM SIGOPS Asia-Pacific Workshop on Systems (APSys) through 2014.

Peter serves on the editorial boards of the ACM Communications of the ACM (CACM) and the Royal Society Open Science Journal, and he co-chaired the PC of EuroSys 2014. He also chaired the selection committee for the ACM SIGOPS Hall of Fame Award in 2011 and 2012, and continues to serve on the committee. Peter also served on the program committees of EuroSys 2014, W+PIN+NetEcon 2014, ASPLOS 2015, MobiSys 2015, SYSTOR 2015, APSys 2015, and HotOS 2015.

4.2 Research agenda

The group's research takes an empirical approach towards realizing the potential of emerging distributed and mobile systems while ensuring security and privacy.

During the reporting period, the group's work has focused on policy compliance in large-scale data processing systems, and on mobile privacy (jointly with researchers at Maryland, Saarland University, and MPI-INF). Additionally, we have collaborated with researchers at NEU on traffic-analysis resistant anonymous voice-over-IP [135], with researchers at Saarland University on accountable anonymity [23], and on mitigating privacy leaks by controlling the discoverability of information [167]. With Krishna Gummadi's group, we continue to collaborate on online privacy [147].

4.2.1 Data policy compliance

In this project, we study methods to enforce declarative data integrity and confidentiality policies in data processing systems, while relying on a small trusted computing base. This work is done in collaboration with Deepak Garg's and Rodrigo Rodrigues' group, as well as researchers at Cornell and Google.

Past work: Enforcing data policies at the storage layer In today's data processing systems, both the policies protecting stored data and the mechanisms for their enforcement are spread over many software components and configuration files, increasing the risk of policy violation due to bugs, vulnerabilities and misconfigurations. Guardat [179] addresses this problem. Users, developers and administrators specify file protection policies declaratively, concisely and separate from code, and Guardat enforces

these policies by mediating I/O in the storage layer. Policy enforcement relies only on the integrity of the Guardat controller and any external policy dependencies. The semantic gap between the storage layer enforcement and per-file policies is bridged using cryptographic attestations from Guardat. We designed and prototyped Guardat, demonstrated how to enforce relevant policies in a Web server, and showed experimentally that Guardat's overhead is low.

Current work: Ensuring compliance in a data retrieval system In recent work, we have designed *Thoth*, a system that extends Guardat by mediating not only storage accesses, but all inter-process data flow in a distributed, parallel computation. Therefore, Thoth can enforce data confidentiality and integrity policies in a complex data retrieval system. For this purpose, we have extended the policy language to support provenance and declassification policies, which control the upstream and downstream data flows, respectively. A poster on Thoth appeared at SOSP 2013 [84] and a paper is currently under submission [83].

Thoth is motivated by the fact that data retrieval systems process data from many sources, each subject to its own data use policy. Ensuring compliance with these policies despite bugs, misconfiguration, or operator error in a large, complex, and fast evolving system is a major challenge. Thoth provides an efficient, OS-level compliance layer for data use policies. Declarative policies are attached to the systems' input and output files, key-value tuples, and network connections, and specify the data's integrity, confidentiality, provenance and disclosure requirements. Thoth tracks the flow of data through the system, and enforces policy at process boundaries, regardless of bugs and misconfigurations outside the OS, or errors by unprivileged operators. Thoth requires minimal changes to an existing system and has modest overhead (less than 4% throughput reduction), as we show using a prototype Thoth-enabled distributed search engine based on the popular Apache Lucene.

4.2.2 Mobile privacy

In this project, we study practical methods to ensure user privacy in mobile systems. This work is done in collaboration with Bobby Bhattacharjee and Elaine Shi at the University of Maryland.

SDDR: Light-Weight Cryptographic Discovery for Mobile Encounters Many emerging mobile social applications use short-range radio to

discover and share content with nearby users. The discovery protocol used to locate other users' devices must preserve user privacy (users cannot be tracked by third-parties), while providing selective linkability (users can recognize friends when strangers cannot) and efficient silent revocation (users can unfriend without permission and without rekeying their entire friend set). Moreover, a good discovery protocol used by such applications must be extremely power efficient since it runs continuously in the background.

In [136], we introduced SDDR (Secure Device Discovery and Recognition), a device discovery protocol that simultaneously satisfies all of the privacy (selective linkability and efficient silent revocation) requirements, while being highly power efficient. Whenever two devices meet, a *secure encounter* is formed, which includes a unique id and an associated cryptographic shared between the encounter peer. The encounter can be used by the peers to address, authenticate and securely communicate, without requiring the exchange of any linkable information. We formally prove the correctness of SDDR, present a prototype implementation over Bluetooth and show how existing frameworks, such as Haggle, can directly use SDDR. Our results show that our SDDR discovery implementation, run continuously over a day, uses only 10% of the battery capacity of a typical smartphone. This level of power consumption is four orders of magnitude more efficient than prior cryptographic protocols with proven security, and one order of magnitude more efficient than prior (unproven) protocols designed specifically for power-constrained devices.

EnCore: Context-based Communication for Mobile Social Apps Mobile social apps provide sharing and networking opportunities based on a user's location, activity, and set of nearby users. A platform for these apps must meet a wide range of communication needs while ensuring users' control over their privacy. In [10], we introduced EnCore, a mobile platform that builds on *secure encounters* between pairs of devices as a foundation for privacy-preserving communication. An encounter occurs whenever two devices are within Bluetooth radio range of each other, and generates a unique encounter ID and associated shared key. EnCore detects nearby users and resources, bootstraps named communication abstractions called *events* for groups of proximal users, and enables communication and sharing among event participants, while relying on existing network, storage and online social network services. At the same time, EnCore puts users in control of their privacy and the confidentiality of the information they share. Using an Android implementation of EnCore and an app for event-based com-

munication and sharing, we evaluated EnCore's utility using a live testbed deployment with 35 users.

iPic: Digital capture privacy Ubiquity of portable mobile devices equipped with built-in cameras have led to a transformation in how and when digital images are captured, shared, and archived. Photographs and videos from social gatherings, public events, and even crime scenes are commonplace online. While the spontaneity afforded by these devices have led to new personal and creative outlets, privacy concerns of bystanders (and indeed, in some cases, unwilling subjects) have remained largely unaddressed.

We have designed iPic [11], a trusted software platform that integrates digital capture with user-defined privacy. In iPic, users choose a level of privacy (e.g., image capture allowed or not) based upon social context (e.g., out in public vs. with friends vs. at work). The privacy choices of nearby users are advertised via short-range radio, and iPic-compliant capture platforms generate media *edited* to conform to the privacy choices of image subjects.

iPic uses secure multiparty computation to ensure that users' visual features and privacy choices are not revealed publicly, regardless of whether they are the subjects of an image capture. Just as importanty, iPic preserves the ease-of-use and spontaneous nature of capture and sharing between trusted users. Our initial evaluation of iPic shows that a practical, energy-efficient system that conforms to the privay choices of image subjects can be built and deployed using current hardware.

Current work: In ongoing work, we are collaborating with vision researchers from the MPI for Informatics on improving the accuracy and robustness of person recognition in iPic, seeking to adapt state-of-the-art recognition algorithms to the mobile environment. We are plannning to extend the system to video and audio recordings. Finally, we are planning to integrate the system into the Android platform, as well as deploy a demo application.

4.2.3 Counter Surveillance

The goal of the following projects is to inform the design, implementation, and deployment of defenses against powerful adversaries such as nation states. It is done in collaboration with Prateek Saxena from National University of Singapore, David Choffnes and Engin Kirda from Northeastern University, Kirill Levchenko from UC San Diego, and Damon McCoy from UC Berkeley.

Herd: A Traffic-analysis Resistant Anonymity Network for VoIP Systems. *Herd* prevents the tracing of VoIP calls, even in the presence of global adversaries able to eavesdrop on all users of a VoIP system [135]. Herd scales to very large numbers (potentially millions) of users and resists traffic analysis while providing latency of at most a few hundred milliseconds between any two pairs of clients. To achieve these properties, Herd combines infrastructure and peer-to-peer resources with the ability to avoid untrusted jurisdictions. Finally, Herd reduces load on its trusted infrastructure using untrusted superpeers, from which call patterns are hidden using a novel design based on network coding.

A Look at Targeted Attacks Through the Lense of an NGO. The goal of this project is to measure and analyze *targeted attacks*, i.e., malicious communication that seeks to compromise the devices of specific individuals, organizations, or communities. To achieve this goal, we contacted over 100 NGOs with a presence in countries where targeted attacks had been reported in the wild [50]. We came into contact with two members of a human-rights NGO representing an ethnic minority living in China who shared $1,493$ suspicious emails that they had collected over a four-year period, $1,176$ of which containing malware. We used this dataset to perform an empirical analysis of targeted attacks in the wild and found that the social engineering was sophisticated (e.g., contact impersonation was frequent), that most attacks used recent, known vulnerabilities that tended to bypass common defenses. Moreover, a quarter of the malware appears to originate from entities reported to have carried out other targeted attacks.

Future work As part of future work, we plan to consider anonymous network access services for additional applications (e.g., Web access) and computing environments (e.g., mobile clients). We will also explore synergies between application dependent, anonymous channels and defenses against targeted attacks, e.g., routing over an anonymity network mitigates packet injection because all traffic is encrypted by default.

5 The Large Scale Internet Systems Group

5.1 Overview

The report covers the period from Nov. 2013 – Jul. 2015. The group's research has focused on the problem of online privacy, particularly problems stemming from the widespread gathering and sharing of user information. The group takes a practical systems approach to this problem. Specifically, it is designing and building systems that accomplish the primary goals of gathering and analyzing user data while protecting user privacy. The primary focus during this period has been on anonymous analytics, which we are pursuing in close cooperation with the spinoff Aircloak (www.aircloak.com).

Personnel. The group is led by Paul Francis. It currently has two graduate students (Istemi Ekin Akkus and Reinhard Munz). Ekin is finishing up and will graduate early fall. He will be working for Bell Labs in Stuttgart. During most of the reporting period, I had one postdoc, Stevens Le Blond, but he has been working independently and collaborating more with Peter Druschel than myself. He recently took a multi-year contract with Peter. Aircloak has four full-time developers—Sebastian Probst Eide, Matthias Kretschmer, Sasa Juric, and Cristian Berneanu—and one full-time sys admin, Sven Knohsalla.

One student left the group during this reporting period: Alexey Reznichenko, who is working with Microsoft's advanced development lab in Munich.

Collaborations. During this period, Paul's group has or had collaborations with Columbia University, Colorado State University, Humboldt University Berlin, Telefonica Research, Ludwig-Maximilians-University Munich (Faculty of Medicine, Eye Clinic), Warwick University, Cisco, and InnoZ (Innovation Centre for Mobility and Societal Change, Berlin).

Publications. During the reporting paper, the group published only a single paper, in ACM CCS [160].

Invited Talks.

- Keynote at the P-Medicine 9th Progress Meeting.
- Presentation at the Telefonica Network Disruptive Council, an executive level think tank focusing on future disruptive technologies.

- Talk at the Microsoft Cloud Computing Research Centre (MCCRC) First Annual Symposium "A Cloud for Europe? Feasibility and Implications of Internet Regionalisation."
- Talks at Bell Labs Research in Stuttgart, and Telefonica Research in Barcelona.

Patents and Technology Transfer. Paul's group has submitted three patent applications. The spinoff company Aircloak became a GmbH (limited liability company). Aircloak continues to work closely with Paul's research program. Aircloak has five employees, and is currently seeking an initial round of angel financing.

External funding.

- EXIST Forschungstransfer (Technology Transfer) Phase I: EUR 394,000. This grant is awarded to MPI-SWS for funding Aircloak. (Aircloak subsequently received Phase II funding of EUR 180,000.)
- Cisco IoT Security Grand Challenge: $75,000. This money was split between MPI-SWS and Aircloak. The MPI-SWS portion was used in support of the startup.

Service Paul served on the following PCs:

- Sigcomm 2014 PC Light
- PETS 2014 PC
- Telefonica Data Transparency Lab Grants Program 2015
- NSDI 2015 PC

Paul also gave a lecture at the 5th PhD School on Traffic Monitoring and Analysis (TMA), 2015.

Awards and media Paul's student Alexey Reznichenko received a Best Student Paper Award at CCS 2014 for his paper "Private-by-Design Advertising Meets the Real World".

The cloaked database system received a TÜViT Trusted Process certification for generalized anonymized analytics. This is to our knowledge the

first such certification given for a general-purpose anonymizing analytics system (i.e. independent of use case).

MPI-SWS Aircloak jointly won the 2014 Cisco IoT Security Grand Challenge competition.

Paul presented at one of the "Impulse aus der Zukunft" open public lectures in Berlin sponsored by Technologie Stiftung Berlin. The event was on the topic "Rethinking online tracking," and included a lecture by Paul and a panel discussion with Dr. Christoph Peylo, VP of Duetsche Telekom Innovation Lab in Berlin (`https://www.technologiestiftung-berlin.de/de/aktuelles/veranstaltungen/beitrag/rethinking-online-tracking/`).

Systems and Tools Aircloak Anonymized Analytics system (www.aircloak.com).

Teaching Paul co-taught the Distributed Systems course with Peter Druschel, winter semester 2015, offered jointly at UdS and TU-KL.

5.2 Research agenda

The almost exclusive focus of my research group during this period is a project to design and build a practical anonymized analytics (database) system. The problem itself is decades old, with the first papers being written in the 1970's, not long after the first relational databases began to appear. In spite of the thousands of papers written on this problem, I think that there is still low-hanging fruit to pick.

The majority of work in this space has focused on the perturbation of simple static data. This includes techniques originating from the statistics community like cell swapping, cell suppression, rounding, and aggregation, as well as techniques originating from the computer science database community like k-anonymity and l-diversity. As a rule, these techniques don't work on more complex data like time-series data (location trace) or for instance a DNA sequence.

A broad alternative approach is *output perturbation*, where the data itself is kept intact, and answers to queries are perturbed in some way. Output perturbation opens the way to dealing with more complex data, since the raw data remains available. A number of different output perturbation approaches were proposed in the early 1980's, but were found to be quite difficult relative to data perturbation, and so the approach didn't get much attention subsequently. Notably, however, the approach was picked up again in the mid-2000's by the cryptography community in the form of differential privacy.

This latter has received an extraordinary amount of attention, especially given that it is a very pessimistic model of privacy with a weak privacy mechanism—the addition of zero-mean random noise. Indeed this very mechanism was called "insecure" by Denning as early as 1980. I am convinced that differential privacy is a non-starter, practically speaking, and the absence of a single broadly usable implementation after 10 years and hundreds of papers testifies to this.

Nevertheless, differential privacy research has shown that counting-queries alone can serve as the basis for a wide variety of data analyses, and so we find counting-queries, with some kind of perturbation, to be a compelling starting point. The low-hanging fruit, then, is to borrow noisy counting queries from differential privacy, but to strengthen the anonymity mechanisms with an empirical systems approach—an approach largely abandoned over 30 years ago. We believe that considerable progress can be made this way.

This begs the question, how do you define progress in this space? At a high level, this is an easy question to answer. The status quo today for dealing with complex data in practice is *pseudonymization*; essentially replacing Personally Identifiable Information like names and account numbers with random values, but otherwise keeping the raw data. An example of this are the HIPAA privacy standards for the sharing of medical data for research purposes. From an anonymity point of view, the HIPAA standards are weak, not because HIPAA doesn't care about privacy, but rather because this is the best that can be done without unduly degrading *utility*. In other words, *utility is more important than anonymity*. Typically contractual means are used to mitigate the weak anonymity. For instance, the HCUP medical database requires analysts to answer questions about their contractual obligations to protect the data, and to acknowledge that they are liable to large personal fines and imprisonment for violations.

So one can make progress in this space with anonymity mechanisms that substantially raise the privacy bar relative to mere pseudonymization, but with almost no degradation in utility. Note that the anonymity mechanisms do not need to be perfect, and do not require guarantees, as evidenced by the fact that people work with shared private data all the time without these guarantees, albeit nervously.

The real questions, then, are: what is the appropriate setting for the utility/privacy tuning knob that would compel its use among practitioners, and is there a mechanism that provides that setting? I believe that the only way to really answer these questions is to provide practitioners with a working implementation, and see if they use it or not.

Towards this end, we created a startup, Aircloak, at the tail end of the last review period. Besides providing a commercial-quality implementation of an anonymizing analytics system, Aircloak is a vehicle by which we can engage the business, regulatory, advocacy, and certification communities. Aircloak has progressed well, having grown to five full-time employees (not including my participation and that of a sys-admin here at MPI-SWS) and, until July of this year, primarily funded by the German EXIST program.

These last two years have produced many insights, and an implementation that is suitable for major trials, a number of which are finally now getting underway. One of these is with a university eye clinic in Munich that will give us eye data for over 3 million patients. We have not yet, however, produced any academic papers. We do expect a number of good papers in the next review period (all of the trials are research trials, with no restrictions on open publication).

A key accomplishment of Aircloak was to obtain a certification from the German data protection certification organization TÜViT which judges the Aircloak service to be suitable for any anonymized analytics application. Such a certification is very important to analytics practitioners. The value of this certificate should not be under-estimated: it was very hard to get, and the certifiers requested a number of changes to our system before they were satisfied. By the same token, the value of the certificate should not be over-estimated: TÜViT by no means did a complete evaluation of our system, not even the anonymizing mechanisms. It is ultimately up to us, and to some extent other academics, to provide a more complete evaluation of the system.

A key lesson learned in these two years is the extent to which practitioners are unwilling to give up on utility. Going into this I had certainly assumed that utility was more important than the academic community normally supposes, but I'm nevertheless continuously surprised at just how important. We started with a set of anonymization mechanisms that included Gaussian noise with a standard deviation of around 30. In the process of interacting with literally dozens of potential users, we have whittled this number down to under 3 (which we typically explain as "plus or minus five"). This has forced us to rethink how to do anonymization, and much of my energy this past two years has been spent doing just that.

The system we've built has two key aspects. One is its data protection, and the other is its anonymization. The data protection is necessary because we store raw data. The level of data protection is such that nobody, including system administrators, has access to the stored data. The servers, called cloaks, operate as super-hardened black boxes. Anything that would

allow someone to log into the box (ssh, console, telnet) is disabled, enforced by SELinux.

We can of course upgrade the software on a cloak, but a trusted third party has oversight of the upgrade that is enforced by the trusted computing base (TPM sealing function). The crypto-disk key is sealed based on a measurement of the boot process up to and including a function that checks a third-party signature of the full cloak image. If the boot sequence or checking function is modified, the TPM will refuse to unseal the crypto-disk key. If the third-party signature does not match, the checking function will halt system operation. Using the TPM in this fashion is not a new idea per se, but we are the first to put it into production.

The anonymization itself takes a defense-in-depth approach. All query output is in the form of user counts, and the anonymization takes place on these counts (similar to differential privacy). Analysts can write arbitrary untrusted code as part of their queries. To prevent this code from defeating the noise by over-counting users, the code is run in a sandbox that limits the code to operating on one user's data at a time, and to outputting only simple count/don't count directives (similar to Aravat).

The first line of defense is noise addition. We add two layers of noise, called random noise and fixed noise, both Gaussian with standard deviation of 2. The random noise is just that. The fixed noise is random, but based on a PRNG seeded with a hash of the user IDs of the counted users. This way, repeats of the same answer can't be averaged to produce the true count.

The system, however, is still subject to a class of attacks that we call *difference attacks*. Here the attacker can generate multiple pairs of answers, whereby the counted users in the pairs are identical except for the possible presence of a victim user in one of the answers. The attacker can create multiple such pairs, each with different user sets, so that there is never a repeated answer. The attacker compares the average value of the sum of each pair half. If they are (nearly) the same, the victim was not in the answer. If they are different, the victim was in the answer, thus revealing the presence or absence of one user attribute.

A second layer of defense deals with this by measuring overlap between pairs of answers, detecting when the attack is taking place (usually), and suppressing the victim from the answer. This approach requires that all previous answers be stored and compared. We have developed efficient implementations of this based on storing bloom filters of answers instead of complete user lists, and on limited what answers need to be compared based on the query input parameters.

This still does not detect all such attacks. We plan to use a third layer

of defense that generates a "signature" of analyst activity, and distinguishes between a normal (non-attack) signature and an attack signature. When an attack signature is detected, then the system operators are informed, and can manually inspect the queries.

Our methodology for testing all of this is empirical. We implement the attacks and defenses, and essentially measure the confidence that the attacker has in guessing the presence or absence of the targeted user attribute for different database settings. The advantage of this approach is that we are free to use defense mechanisms that cannot be formally modeled. The disadvantage, of course, is that we can only test the attacks that we think of. There may be other attacks that we are unaware of.

It is worth mentioning that we have not been particularly successful at integrating students' research into the Aircloak activity. I think this is mostly due to the fast pace of innovation and development in Aircloak, much of which is not suitable for publishing anyway. I'm optimistic that once we are further into the trials and gathering operational data, it will be easier for students to participate and add value.

Deploying Privad and PDDP PhD student Alexey Reznichenko finished his large-scale deployment of Privad (private advertising system) and PDDP (private analytics system). Both of these technologies use a client-centric approach to anonymization, whereby user data is kept at the user client, and an honest-but-curious proxy is used to broker between the client and the back-end (advertising or analytics) system. Alexey managed to get around 80,000 installs, of which 15% opted-in. During the experimental run-time of roughly two months, he had 4500 average daily users, with a typical daily peak of around 2000 online users. All told, these users generated 1.1M ad requests, and were delivered 9.5M ads of which 790K were displayed. There were 417 clicks (which in Alexey's system led to real shopping sites), and ultimately 4 products were purchased. This click-through rate was comparable to those of Google display ads (text) for the same set of users.

All of the experimental analysis of the system took place through the PDDP private analytics system itself. In total 159 distinct queries were made, producing 790K distinct differentially private answers (buckets) across 9400 unique clients.

The results of this experiment were mixed. On the positive side, Alexey demonstrated that a private advertising system can work, and indeed could compete with Google's non-targeted text ads. This is a remarkable result

for a single student! On the negative side, we found that differential privacy is a poor model for understanding privacy loss in practice. Based on the number of queries Alexey made, differential privacy tells us that some loss of privacy could have taken place. In practice, however, given the nature of our data, what we knew of the user population (nothing), and the types of queries we made, no privacy loss was possible.

Alexey won a best student paper at CCS14 [160] for this work. I was gratified to see this, because this kind of experience paper requires a lot of work relative to the number of publications one can get, and so is rarely done. It is nice to see this kind of work rewarded. Alexey graduated with a position at Microsoft's advanced technology lab in Munich.

6 The Foundations of Computer Security Group

6.1 Overview

This reports covers the group's activities in the period November 2013 – July 2015. The group is broadly interested in security and privacy problems, both foundational and systems-oriented. During the reporting period, the group's research has primarily focused on methods for policy compliance and applied type systems.

Personnel. Deepak Garg (faculty); Ezgi Cicek, Vineet Rajani (students); Eslam Elnikety, Aastha Mehta, Anjo Vahldiek-Oberwagner (students co-advised with Peter Druschel); Paul David Swasey (student co-advised with Derek Dreyer).

Collaborations. Internally, the group has collaborations with Peter Druschel and Derek Dreyer's groups. Externally, the group collaborated with researchers at the Saarland University (Christian Hammer and Aniket Kate), Carnegie Mellon University (Anupam Datta, Limin Jia and Umut Acar), Chalmers University (Andrei Sabelfeld), ICSI Berkeley (Michael Tschantz), Purdue (Omar Chowdhury), INRIA (Tamara Rezk) and the Indian Institute of Science (Sanjit Chatterjee) during the reporting period.

Publications. During the reporting period, the group's activities have resulted in conference papers at CAV [68], CSF [157, 78, 119], CCS [67], ESOP [69], EuroSys [180], POST [48] and CODASPY [154], and a workshop paper at PLAS [47].

Teaching. Deepak taught a graduate course on Logics in Security in the winter of 2014/15 and a seminar on Language-based Security in the summer of 2015. The former was rated among the top 5 courses at Saarland University based on student feedback.

External funding. The group's research on security in web browsers is funded partially by the German Research Foundation, Deutsche Forschungsgemeinschaft (DFG), under a priority program titled RS3. The funding has been secured jointly with Christian Hammer from Saarland University. The funding lasts from October 2012 to October 2016. Deepak and Peter Druschel jointly won a Google Faculty Research award in 2014. Deepak, Aniket Kate (MMCI), Sanjit Chatterjee (IISc Bangalore) and Mridul Nandi

(ISI Kolkata) were jointly awarded an Indo-German Max Planck Center for Computer Science (IMPECS) grant for a collaborative project.

Awards and invited talks. Deepak was awarded a Google Faculty Research Award, jointly with Peter Druschel in 2014. Deepak was an invited speaker at the 2015 DICE-FOPARA workshops.

Service. Internally at MPI-SWS, Deepak served on the faculty recruiting committee in 2014 and 2015, organized the internship program in 2015 and the annual institute retreat in 2014. Externally, Deepak co-chaired the workshop on Foundations of Computer Security (FCS) in 2015. He served on the program committees of IEEE S&P 2014 and 2015, POPL 2015, CCS 2014 and 2015, CSF 2014 and 2015 and SACMAT 2014. He has also been the Publications Chair of CSF every year since 2012. Ezgi, a graduate student in the group, is the elected student representative of the institute.

6.2 Research agenda

During the reporting period, the group conducted research on several aspects of policy compliance (information flow control in web browsers, auditing for privacy policy violations, policy compliance in distributed data processing systems, and cause analysis) and applied type systems (for incremental computational complexity and for verification of systems that use code sandboxing).

6.2.1 Policy compliance

IFC4BC: Information flow control for web browser clients Web applications routinely rely on untrusted, third-party libraries. However, current browser security models only support access control at the granularity of entire scripts, which allows third-party scripts to use available sensitive data (passwords, credit card numbers, browsing history) beyond their intended purposes. Fine-grained information flow control (IFC) is a promising solution to this problem, but IFC is difficult to design and implement for browsers because the complex semantics of JavaScript and the browser page state (the Document Object Model or DOM) make it easy to miss exploitable flows.

In the last bi-annual report, we explained our work towards a practical, fine-grained IFC solution for JavaScript. Back then, we instrumented the *bytecode* interpreter in WebKit, an open source browser engine. By working

with bytecode (instead of source), we leverage several years of industrial performance optimizations in both the source-to-bytecode compiler and the interpreter itself. We have moderate overheads (tens of percents on the most intensive benchmarks), but that overhead is several orders of magnitude less than that of prior work. Our instrumentation, and a formal proof of its soundness, have since been published [48, 47].

In the last two years, we have expanded our IFC instrumentation to two critical components of the browser [157]. First, we have built the first reasonably complete model of the brower's shared-memory APIs (the so-called DOM APIs), and expanded our IFC instrumentation to them. Second, we have built the first formal model of the browser's event handling logic, which is surprisingly complex and subtle. Based on this model, we observed that the event handling logic is an easily exploitable channel of information leaks. We have proposed solutions and have expanded our formal model and soundness proof to cover these channels.

Our ongoing work is aimed at scaling IFC on JavaScript. Specifically, we are conducting empirical studies to understand IFC's false positives in web applications. We are also developing a practical declassification model for web scripts. (IFC4BC is carried out jointly with the group of Christian Hammer at UdS, and is funded by the DFG. In somewhat separate joint work with Andrei Sabelfeld at Chalmers and Willard Rafnsson at CMU, we are looking at extending our methods to prevent leaks through program progress and termination.)

Audit for privacy policies Service providers like hospitals, banks and online social networks collect sensitive client data to optimize their services. To detect misuse, these providers record detailed logs of internal accesses to data (and associated metadata), and audit these logs for policy violations. In the last reporting period, we explained our work on a syntax-directed *offline* audit algorithm for policies written in linear temporal logic (LTL). LTL is known to be a good foundation for expressing many data-use policies. That work continued in two different projects, Eunomia and précis, during the reporting period. We explain both projects briefly.

Eunomia: In practice, logs are large and audit is computationally expensive, so small- and medium-scale enterprises may outsource log storage and audit to third-party clouds. However, the log has sensitive data, which must be encrypted before the log is sent to a cloud. Standard encryption cannot be used because it would render the log useless for audit (audit runs SQL-like select, project and join queries on the log). Eunomia [67]

is a database encryption system that supports queries with pre-determined operations (select, project, join) and, importantly, provides cryptographically provable upper-bounds on information leaked (unlike prior work like cryptDB). As one specific application, we show that policy audit runs correctly on Eunomia with only moderate overhead.

Précis: We have also developed an alternate, *online* audit algorithm called précis [68] that can detect violations in real time. A novel, flow-sensitive static analysis of policies guides caching of only relevant log information, thus optimizing space utilization. (Eunomia and précis were developed jointly with Omar Chowdhury from Purdue, and Limin Jia and Anupam Datta from CMU.)

Policy compliance in distributed data processing systems In collaboration with Peter Druschel's group, we are working on two systems — Guardat [180] and Thoth (in submission) — that enforce confidentiality and integrity policies in distributed data processing systems. See Peter Druschel's section of the report for details of these systems.

Cause analysis A security protocol may go wrong due to bugs and misconfigurations and human deviance from prescribed norms. When this happens, it may be important to hold accountable and possibly even punish responsible human agents and to fix bugs and misconfigurations. A common requirement for accountability and blame analysis is the determination of *cause* — the agents, software components, code lines or inputs — that caused the error. Even defining cause is known to be difficult. The philosophy community (e.g., Hume, Pearl, and Halpern) has debated many definitions of cause over several centuries, without reaching a clear consensus. Specifically, these definitions are based on counterfactual considerations: very approximately, event A is a cause of event B if had A not happened, B would also not have happened.

In recent work, we have defined a notion of cause in process calculi [78]. Process calculi are often used to model security protocols and are, therefore, an appropriate foundation for our work. Our definition determines which subexpressions of a program are responsible for the violation of a safety property. It refines the Lamport cause (all events in the happens-before relation of a property-violating event) by eliminating subexpressions on which the outcome does not depend. Some notable features of our definition are that it determines both joint and independent causes, it is semantic, and it captures *minimal sufficiency* — the determined causative set of subexpres-

sions suffices to cause the violation and among all such sets, it is minimal. Minimal sufficiency seems to be an appropriate notion of cause for debugging protocols.

In ongoing work, we are examining more foundational questions in cause analysis. How does language abstraction affect cause determination? Which counterfactual scenarios should be considered in determining cause? Is there a connection between specific causation and general causation? (Most of this project was executed jointly with Divya Sharma, Dilsun Kaynar and Anupam Datta from CMU. Recently, we have started collaborating with Michael Tschantz from ICSI, Berkeley.)

6.2.2 Applied type systems

Code sandboxing and object capabilities Code sandboxing is widely used to enforce security properties on untrusted code. However, support for statically verifying properties of programs that use code sandboxing is limited. One difficulty is that the adversary code (sandboxed code) is usually not available for analysis ahead of time. We have recently developed System M [119], a type-theory for proving safety properties in systems that use code sandboxing. The central reasoning principle in System M is that sandboxed code will *always* adhere to the *intersection* of the invariants of all interfaces it can access. Hence, to reason about the properties of a sandbox, the code inside it is not needed. Technically, System M is a refinement type system that supports higher-order programs and loosely builds on Hoare type-theory (HTT). (System M was developed jointly with Limin Jia and Anupam Datta from CMU.)

Whereas code sandboxing is coarse-grained, object capabilities are widely used to enforce fine-grained access control. For instance, the Firefox web browser uses object capabilities extensively to limit cross-domain access in web apps. Despite decades of *use* of object capabilities, formal foundations for understanding and reasoning about properties obtained from object capabilities are lacking. Within this space, we are pursuing two lines of work. First, in collaboration with Derek Dreyer's group, we are developing program logics for verifying programs that use object capabilities (our current case study is Firefox). One key observation is that object capabilities are pointers, and methods like separation logics and Kripke logical relations that were developed originally for reasoning about memory safety can be applied to prove safety properties of capability systems. Specifically, we are building on Iris, a verification framework recently developed in Derek Dreyer's group. Second, in collaboration with Tamara Rezk (INRIA, Sophia-Antipolis), we

are looking at formalizing the object capability model abstractly and examining fundamental differences in expressiveness from access control and information flow control.

CostIt: Types for incremental computational complexity. This project applies techniques from the security community (specifically, IFC) to a non-security domain. Incremental computation seeks to speed up repeated executions of a program with updated inputs. During the first run of the program, a graph recording the dependencies between inputs, intermediate values and outputs is built. In a subsequent run, changes are *propagated* through the dependency graph, starting from modified inputs to the outputs. If the graph is wide and shallow (i.e., the computation is data parallel) and few inputs change, then change propagation is asymptotically sub-linear in the size of the graph. This speedup comes at the cost of storage for the dependency graph but it pays off when reduction of computational latency is paramount; a classic example is the utility `make`, which retains intermediate compiled objects to often cut incremental build latencies from several hours to a few seconds or minutes. Recent research in incremental computation has developed language-based techniques to automatically compile programs to their incremental counterparts. However, despite nearly a decade of work, there is no systematic method (even theoretical) to help the programmer determine the actual benefit of incremental programming, i.e., to explore the space-time trade-off of incrementalization for a given program.

To fill this gap, we have developed a type-theory, CostIt, that enables formal proofs of the asymptotic complexity of incremental computations [69]. For instance, one can show in CostIt that standard merge-sort on a list of length n with constant number of changes can be re-computed in time $O(n)$ (the worse-case asymptotic complexity of merge-sort is $O(n \cdot \mathsf{log}(n))$, so this is a provable speedup). CostIt works particularly well on divide-and-conquer algorithms, which form the basis of many data parallel computations like those running on MapReduce. CostIt relies heavily on refinement types and dependency analysis, and on IFC annotations for static approximations of what can change and what cannot change during an incremental run. Foundationally, CostIt's soundness is based on a novel model of relational side-effects. CostIt can be simplified to reason about the size of the dependency graph and, therefore, it encompasses both axes of the aforementioned trade-off. In ongoing work, we are implementing CostIt using a combination of algorithmic type-checking and state-of-the-art SMT solvers, and we describe how standard programs can be (provably) correctly compiled to

incremental programs. In the future, we intend to examine incrementalization for lazy computation. (Early iterations of CostIt were developed in collaboration with Umut Acar from CMU. In the last report, a preliminary development of this work was reported under the title "Trace continuity".)

7 The Networked Systems Group

7.1 Overview

This section describes the activities of the Networked Systems group between November 2013 and July 2014. The group's research interests are in measurement, analysis, design, and evaluation of complex Internet-scale systems. Recently the group's projects have focused on understanding and building social computing systems. Specifically, they tackle the challenges associated with (i) reasoning about the trustworthiness of user identities, (ii) protecting the privacy of users sharing personal data, (iii) understanding and improving the dependability of information exchange mechanisms in social media, and (iv) defining the notions of fairness, discrimination, and transparency in data-driven decision making and developing mechanisms for enforcing them.

Personnel. The group is led by Krishna Gummadi. It is currently comprised of five graduate students (Mainack Mondal from October 2010, Juhi Kulshretha from April 2011, Muhammad Bilal Zafar from October 2012, Giridhari Venkatadri from October 2013, Reza Babaei from April 2014) and four postdoctoral researchers (Przemyslaw Grabowicz from October 2013, Rijurekha Sen from June 2014, Oana Goga from July 2015, Denzil Correa from March 2015). Rijurekha Sen is co-advised with Peter Druschel.

Bimal Viswanath will be graduating with a PhD in October 2015. He will be joining Bell Labs as a researcher. Saptarshi Ghosh, a postdoctoral researcher, will resume his position as an Assistant Professor at the National Institute of Technology, Shibpur, India.

Collaborations. Internally, the group members have close collaborations with the distributed, social information systems, and machine learning groups led by Peter Druschel, Cristian Danescu-Niculescu-Mizil, and Manuel Gomez-Rodriguez respectively.

Locally, the group members collaborate with security, cryptography, and databases groups at University of Saarland and MPI-INF led by Michael Backes, Aniket Kate, and Gerhard Weikum respectively.

External collaborators include researchers from Boston University (Mark Crovella), Microsoft Research (Saikat Guha), UFMG (Fabricio Benevenuto), Northeastern University (Alan Mislove), Cornell University (Michael Macy), AT&T (Balachander Krishnamurthy), KAIST (Meeyoung Cha), Facebook (Winter Mason), INRIA (Renata Teixeira), ICSI Berkeley (Robin Summer),

Telefonica Research (Pablo Rodriguez), Kings College London (Nishanth Sastry), QCRI (Ingmar Weber), and IIT Kharagpur (Niloy Ganguly).

Publications. Due to the inter-disciplinary nature of their work, the group members regularly publish their research in the top conferences, journals, and workshops in different sub-areas of computer science. Specifically, during the reporting period, group members have co-authored papers in:
1. Security and Privacy: Usenix Security [181], SOUPS [150], and USEC (an NDSS workshop) [148].
2. Datamining, Maching Learning and Information Retrieval: KDD [104], FATML (an ICML workshop) [189], RecSys [46], and CIKM [164].
3. Web and Social Media: IMC [105], Transactions on the Web [188], COSN [182], ICWSM [131, 70, 107], and ASONAM [95].
4. Human-Computer Interaction: CSCW [45] and CHI [93].

The group's past publications in the area of social computing systems have been highly cited. Two of the group's papers in AAAI's International Conference on Web and Social Media (ICWSM) 2010 and Usenix/ACM Internet Measurement Conference (IMC) 2007 are the highest cited papers in the history of the conferences with over 1500 and 2000 citations (according to Google Scholar), respectively. Eleven other papers in the area have been cited over 100 times each.

Software, data, and technology transfer. The group strives to make its software and data sets available to others to the extent possible. To date, over 1000 research groups at universities and research labs worldwide have used the data sets we gathered as part of our measurement studies of online social networks.

During the reporting period, the group members have implemented and publicly deployed several proof-of-concept systems on the popular Twitter social media site: (a) to detect manipulation (tampering) of popularity of content [4] and users [3], (b) to help users measure and manage their information diets [6], (c) to find relevant and trustworthy topical content from the site [5]. These systems have attracted the attention of developers at Twitter, Facebook, and Yelp, who have acknowledged (in personal communication) the influence of our system designs on the design of their official trust assessment and search systems.

Amongst the systems that we have publicly deployed in the past: (i) Glasnost software [1] designed to test the traffic management policies of their access ISPs (e.g., cable and DSL providers), continues to attract large num-

bers of users (over 500,000) during the reporting period, with the data gathered being used by multiple telecom regulators world-wide, (ii) FriendList Manager app [2] to help Facebook users manage their privacy settings better has acquired than 2000 users, and (iii) our Twitter-based system for finding trustworthy topical experts [7] has been used by a few hundred research groups world-wide.

External funding. The research of the group has been partially funded by Humboldt and IMPECS fellowships. Saptarshi Ghosh and Rijurekha Sen have won both Humboldt and IMPECS postdoctoral fellowships, while Denzil Correa has won an IMPECS postdoctoral fellowship.

The group also received an unrestricted grant of 25K US dollars from AT&T for studying data quality loss resulting from data deletions in social media.

Invited talks and awards. Krishna gave keynote talks at the 15th International Conference on Web Information System Engineering (WISE 2014) and at the 2nd Annual Security and Privacy Symposium organized by IIT Kanpur, India.

Mainack Mondal, Bimal Viswanath, and Krishna Gummadi also received the Distinguished Paper Award at the Symposium on Usable Security and Privacy (SOUPS) 2014 conference for their work on understanding and specifying social access control lists.

Service. Internally to MPI-SWS, Krishna served as the chair of the Faculty Recruiting Committee for the 2014 hiring season, and is currently leading the effort to redesign MPI-SWS website.

Externally, Krishna has served as a program co-chair for ACM's COSN 2014 and IW3C2's WWW 2015. He has also served on the program committees of Oakland S&P 2014, SIGCOMM 2014, VLDB 2014, ICWSM 2015, KDD 2015, COSN 2015, WSDM 2015.

Krishna currently serves as an associate editor for Transactions on the Web, a steering committee member of Measurement Lab and as the chair of the technical advisory board of the ACM conference on online social networks (COSN).

Krishna also co-chaired the first Data Transparency Lab (DTL) grants' committee. DTL is a community of technologists, researchers, policy makers, and industry representatives working to advance online personal data transparency through scientific research.

7.2 Research agenda: Social computing systems

Social computing systems are an emerging class of *societal-scale human-computer systems* that facilitate interactions and knowledge exchange between individuals, organizations, and governments in our society. Examples include social networking sites like Facebook and Google+, blogging and microblogging sites like Twitter and LiveJournal, content sharing sites like YouTube and Instagram, social bookmarking sites like Reddit and Pintrest, crowdsourced opinion sites like Yelp and eBay seller ratings, and social peer production sites like Wikipedia and Amazon's Mechanical Turk.

From a computer scientist's perspective, social computing systems are *human-centric distributed systems*, where humans (and their psychological behaviors and social interaction networks) are as much a key component of the distributed system as are the software systems interconnecting them. From a social scientist's perspective, they are *computer systems-assisted human interactions*. Our studies are motivated by both these perspectives.

The high-level goals of our studies of social computing systems are to *understand, predict, and control* the behaviors of their constituent human users and computer systems. We leverage a variety of inter-disciplinary methods in our studies:

1. User-centric Studies: We conduct empirical studies of user behaviors and interactions in social computing systems using large-scale observational studies of deployed systems as well as smaller-scale studies of users recruited on the Web to participate in experimental systems and surveys.

2. Data-centric Studies: We construct and analyze computational models of user and system behaviors from the data we gather using appropriate data mining (e.g., graph analysis, data clustering and dimensionality reduction), statistical learning (e.g., supervised learning and convex optimization) and NLP (e.g., statistical language and topic models) techniques.

3. Systems-centric Studies: We leverage insights from our empirical and analytical studies above to design, implement and deploy useful systems and services in practice (e.g., fraud detection services, privacy management tools, diversity-preserving recommender systems).

In the past, we have conducted some of the earliest studies of the structure and growth of large-scale real-world social networks and evolution of user interactions over them. Our recent studies are focussed on four themes:

1. Trustworthiness, Reputation, and Accountability of Social Identities

2. Privacy, Anonymity, and Exposure in Crowdsourcing Systems

3. Information Dissemination and Retrieval in Social Media Systems

4. Fairness, Bias, and Transparency in Data Driven Decision Systems

7.2.1 Trustworthiness, reputation, & accountability of social ids

Background and Past Work: Social computing systems need mechanisms to assess the *reputation* of their users, reason about their *trustworthiness* of their actions, and hold them *accountable* for their actions. Unfortunately, user identity infrastructures are the achilles heel of today's social computing systems. Most systems allow users to operate behind *weak identities*, i.e., online accounts that can be created users without showing any proof or certification of their offline / physical identities. While they offer a certain level of user anonymity, weak identities leave current systems vulnerable to a variety of attacks using fake (Sybil) identities.

Our past work explored ways to exploit the decentralized social trust implicit in social links created by user identities within a social computing system. Our studies not only identified fundamental weaknesses with existing schemes, but also proposed a new approach, *Sybil tolerance*, that uses social networks as credit networks to bound the impact of fake identities.

Recent Work – during the review period: Our work in this period is motivated by the following fundamental question: *what other information (beyond network links) within a social computing system could be leveraged to assess reputation and trustworthiness of its weak identites?* One possibility is to use the past behavioral activity of the identities within the system. But existing approaches have a fundamental limitation: they rely on checking whether an identity has exhibited some *a priori* known pattern of misbehavior (using supervised learning). In practice, this sets up a arms-race with attackers constantly adapting their activity to avoid detection.

In our work, published at Usenix Security 2014 [181], we circumvent this arms race by capturing all normal patterns of user behavior so that any identity whose behavior deviates significantly from the normal behavior could be declared as anomalous and inspected further to detect potential misbehavior (unsupervised learning approach). The key challenge with our approach lies in capturing the wide variety of normal behaviors of identities. Our insight is that in practice, identities do not behave randomly, i.e., even as their behaviors span a high dimensional space, the common or normal patterns could be succinctly mapped to a lower-dimensional sub-space using dimensionality reduction techniques, such as principal component analysis. We showed the viability of the idea using real traces of identity behaviors from sites such as Facebook, Yelp and Twitter. Our study of Facebook ad-clicks revealed wide-spread prevalence of click-fraud in certain types of Facebook advertisements. We also discussed with Facebook security team about the feasibility of deloying our techniques on the Facebook platform.

Our above approach, while effective in many scenarios, cannot be applied for newly created identities with no prior activity. So resourceful attackers can create a large number of new identities and use them immediately to attack the system. published at We address this limitation in our COSN 2015 [182].Our key insight is that fake identities cannot forge the timestamps of their activities as they are recorded by system operators. So we can analyze the statistical distributions of timestamps of identities activities to robustly detect tampering of a crowd of identities. We have successfully applied our approach to detect tens of thousands of previously undetected tampered crowd promotions in Yelp and Twitter. The Yelp security team is interested in trying out the techniques over Yelp platform.

Finally, in our IMC 2015 [105] work, we attempted to detect the most challenging type of identity fraud, *impersonation* – where a user pretends to be some other real-user and simply copies their profiles and activity from another site (e.g., an attacker creates an identity on Twitter imitating the public posts of a real user on Facebook). Impersonation can cause serious damage to the victim's reputation. Our study shows that while the problem is hard in the general case, impersonators could be identified when the victim also has an account on the same site. We detected thousands of impersonating accounts in Twitter using a technique we proposed.

7.2.2 Privacy, anonymity, & exposure in crowdsouring systems

Background and Past Work: The growing popularity of crowdsourcing websites, where users share unprecendented amounts of personal information about their activities and preferences, raises serious concerns about users' privacy. We believe that these concerns are being exacerbated by the lack of meaningful and intuitive privacy controls and abstractions. Instead of making the task of privacy management easier, the most common controls today are not intuitive and require users to expend significant mental effort.

In our past work [140], we quantified the magnitude of the privacy management problem by measuring (via user surveys) the disparity between the desired and actual privacy settings of users in Facebook. We also implemented and deployed a Facebook app called FriendList Manager to help simplify the complex task of configuring privacy settings of users [141].

Recent Work – during the review period:

1. Access controls: To protect their personal content from being exposed to the wrong audience, today, most crowdsourcing systems provide fine-grained mechanisms for specifying *social access control lists* (social ACLs,

or SACLs), allowing users to restrict their sensitive content to a select subset of their friends. To design better privacy management tools for users, we need to first understand the usage and complexity of SACLs specified by users. In a study published at SOUPS 2015 [150] (that won the distinguished paper award), we presented the first large-scale study of fine-grained privacy preferences of over 1,000 users on Facebook, providing us with the first ground-truth information on how users specify SACLs on a social networking service. Our key take-aways are that (i) a surprisingly large fraction (17.6%) of content is shared with SACLs; (ii) SACL membership shows little correlation with either profile information or social network links; as a result, it is difficult to predict the subset of a users friends likely to appear in a SACL; and (iii) SACLs are often reused, suggesting that simply making recent SACLs available to users is likely to significantly reduce the burden of privacy management on users (an option that sites like Facebook do not provide today).

2. Exposure controls: Having analyzed the challenges with configuring access control settings, in our work published at USEC 2014 [148], we argued that traditional access control models are fundamentally inadequate for today's online world. First, with access control, users must a priori specify precisely who can or cannot access information by enumerating users, groups, or roles, which is difficult to get right. Second, access control fails to separate who can access information from who actually *does*, because it ignores the difficulty of *finding* information. Third, access control does not capture if and how a person who has access to some information redistrbutes that information. Lastly, access control fails to account for information that can be inferred from other public information. We proposed an alternate model for information privacy called *exposure*, which captures the set of people expected to learn an item of information eventually. We are currently investigating mechanisms to enable users to control the exposure of their personal data.

3. Anonymity of users and content: Recently, a number of online services such as the "people search" engine Spokeo are offering to aggregate and link data about individual users from multiple online social networking sites. Such matching of user profiles across different sites raises serious privacy concerns. In our work published at KDD [104], we studied the extent to which profiles can be matched reliably across real-world social networking sites with hundreds of millions to billions of profiles. Our key finding was that, when matching profiles at large-scales, the potential for false matches is very high using existing profile matching schemes. We show that to avoid false matches one should pay a significant cost in terms of reduced recall.

Our study raises serious concerns about the potential for falsely linked data in information retrieved via today's people search engines.

Recently, there has been a significant increase in the use of anonymous crowdsourcing sites like Whisper and Secret, where posts are not associated with well-defined user identities or profiles. In our study published in ICWSM 2015 [70], we investigated the differences between content posted on anonymous and non-anonymous crowdsourcing sites like Twitter. We introduced the notion of *anonymity sensitivity* of a piece of content, which captures the extent to which users think the content should be anonymous. We proposed a human annotator based methodology to measure the same for Whisper and Twitter posts. Our analysis reveals that anonymity sensitivity of most whispers (unlike tweets) is not binary, with different whispers being marked as sensitive by different fractions of users. Our findings shed light on human behavior in anonymous media systems that lack a well-defined notion of an identity.

7.2.3 Information dissemination & consumption in social media

Background and Past Work: With the widespread adoption of social media sites like Twitter and Facebook, there has been a shift in the way information is produced, disseminated, and consumed in our society. Earlier, the only producers of information were traditional news organizations, which broadcast the same carefully edited information to all consumers over mass media channels. Whereas, now, in online social media, any user can be a producer of information, every user can select which other users she connects to, and all users can selectively propagate information they like to their social connections.So an important research challenge lies in understanding and improving existing mechanisms for information exchange in social media.

In the past, we have conducted some of the earliest studies of information dissemination in online social networks by (i) analyzing photo dissemination in the Flickr social network and URL dissemination in the Twitter social network (ii) studying the emergence and adoption of social (linguistic) conventions over the Twitter social network, (iii) studying the role or influence of individual users in propagating information in the Twitter social network.

Recent Work – during the review period:

1. Information dissemination: When users in social media systems are overloaded with information, they might overlook much of the information they receive, which could severely impact the spread of information by word-of-mouth. In a collaborative work with Manuel Gomez-Rodriguez [107] published at ICWSM 2014, we conducted a large scale quantitative study of

information overload and evaluate its impact on information dissemination in the Twitter social media site. Our analysis, modeling users as information processing systems, provides first empirical evidence of information processing limits for social media users and the prevalence of information overloading. We find that the rate at which users receive information impacts their processing behavior, including how they prioritize information from different sources, how much information they process, and how quickly they process information. We show that the susceptibility of a social media user to social contagions depends crucially on the rate at which she receives information. Our findings make the case for new models of viral propagation that take into account the extent to which users are overloaded.

2. Information consumption: To quantify and understand the impact of exchange mechanisms on information consumed by users in social media, we introduced the concept of information diet in our work published at ICWSM 2015 [131]. We propose to measure the diet of a given set of information items (e.g., tweets) by the statistical distribution of topics assigned to the information items. Our analysis of information diets of Twitter social media users showed that (i) popular users mostly produce very specialized diets focusing on only a few topics; in fact, news organizations (e.g., NYTimes) produce much more focused diets on social media as compared to their mass media broadcasts, (ii) most users consumption diets are primarily focused towards one or two topics of their interest, and (iii) the personalized recommendations provided by Twitter help to mitigate some of the topical imbalances in the users consumption diets, by adding information on diverse topics apart from the users primary topics of interest.

3. Information assimilation: Not all information consumed by a social media user impacts or affects the user's thinking or opinions in a similar way. In social psychology it is well-known that users have a variety of cognitive biases, e.g., confirmation bias, where users tend to accept information that confirms their prior beliefs and reject those that contradict their beliefs. Against this background, an interesting research challenge lies in understanding the extent to which different factors affect the popularity of a content in social media. In our CHI 2014 study [93], we investigated the extent to which content by itself (i.e., the interestingness, topicality, or quality of the information as perceived by users) determines the popularity of YouTube videos. Using mechanical turk as experimental platform, we asked users to evaluate pairs of videos, and compared users relative perception of the videos content against their relative popularity reported by YouTube. We found that in most evaluations users could not reach consensus on which video had better content as their perceptions tend to be very subjective.

Nevertheless, when consensus was reached, the video with preferred content almost always achieved greater popularity on YouTube. Our study shows that while high quality content has a high chance of becoming popular in social media, most popular content is not necessarily of high quality.

7.2.4 Fairness, bias, & transparency in data-driven decisions

Background and Motivation: Decision making processes in online services have become increasingly automated and data-driven. For example, algorithms trained on past data about users are increasingly being used to determine what news and information users get to see, who they meet online, and what prices they are offered. The role of algorithmic (data-driven) decision making in predicting and influencing user behaviors in these systems raises fundamental questions about their transparency, fairness, bias, and potential for discrimination – concepts that have been explored previously in social, cultural, economic, and legal frameworks, but are yet to be defined in computational learning frameworks. A grand challenge here is to formally define these concepts in learning frameworks and propose mechanisms to enforce them in practice.

Recent Work – during the review period: In a recent collaboration with Manuel Gomez-Rodriguez's group (published at FATML, an ICML workshop [189]), we focussed on the design of classifiers with formally defined fairness guarantees. Specifically, we introduced the idea of fairness constraints, which prevent a classifier from making predictions / decisions that are correlated with certain sensitive attributes in the data (e.g., gender or race). Our framework i) is readily generalizable to a variety of classifiers; ii) does not utilize sensitive attributes during tesing but only during training; iii) supports both continuous and categorical sensitive attributes; and iv) provides clear mechanisms to trade-off fairness and accuracy. Our proposed approach is guranteed to achieve optimal classification accuracy under the fairness constraints for a variety of classifiers. We also validated our approach via experiments over real-world census datasets.

Ongoing and Future Work: Social media users are increasingly relying on recommendation algorithms to cope with the deluge of user-generated content shared on these systems. However, social media systems today reveal very little publicly about how their recommender services personalize information access to users, making it hard to check whether the information they recommend to their users is biased in some way. In ongoing work, we are focussing on making recommender systems more transparent.

8 The Rigorous Software Engineering Group

8.1 Overview

The report covers the period from November 2013 – July 2015. The group's research interests are in the foundational principles of software engineering (models of computation, analysis algorithms) and applications of these principles to programmer productivity tools. Major research topics are in the verification and control of reactive, real-time, and hybrid systems, software verification and program analysis, logic, and automata theory.

Personnel. The group is led by Rupak Majumdar and currently has four graduate students (Zilong Wang, Johannes Kloos, Filip Niksic, and Susanne van den Elsen), and four postdoctoral researchers (Dmitry Chistikov, Rayna Dimitrova, Vinayak Prabhu, and Anne-Kathrin Schmuck).

An additional postdoctoral researcher (Sadegh Soudjani) worked at the institute for 4 months before moving to Oxford.

Collaborations. The group has joint publications with Viktor Vafeiadis and Ruzica Piskac. We have a joint ERC project together with Michael Backes, Peter Druschel, and Gerhard Weikum.

Externally, the group has had collaborations with Alessandro Abate (Oxford), Sanjoy Baruah (UNC), Krishnendu Chatterjee (IST Austria), Jyo Deshmukh (Toyota), Michael Emmi (IMDEA Madrid), Javier Esparza (TU Munich), Pierre Ganty (IMDEA Madrid), Milos Gligoric (Univ of Texas at Austin), Stefan Göller (ENS Cachan), Sumit Gulwani (Microsoft Research), Lei He (UCLA), Holger Hermanns (Saarland University), Aditya Kanade (IISc Bangalore), James Kapinski (Toyota), Stefan Kiefer (Oxford), Viktor Kuncak (EPFL), Jerome Leroux (LABRI Bordeaux), John Lygeros (ETH Zurich), Darko Marinov (UIUC), Roland Meyer (University of Kaiserslautern), Todd Millstein (UCLA), Supratik Mukhopadhyay (Louisiana State), Joel Ouaknine (Oxford), Andreas Podelski (University of Freiburg), Shaz Qadeer (Microsoft Research), Jeffrey Shallit (University of Waterloo, Canada), Paulo Tabuada (UCLA), Sai Deep Tetali (UCLA), Ufuk Topcu (University of Pennsylvania), James Worrell (Oxford), Thomas Wies (NYU), Matthias Woehrle (Bosch), Hongseok Yang (Oxford), and Majid Zamani (TU Munich). Many of these collaborations are ongoing.

Publications. The publications of the group have broadly been in three areas: embedded systems and control theory, formal verification and soft-

ware engineering, and logic and foundations.

In embedded systems and control, the group has published 3 papers in HSCC [62, 143, 165], 2 papers in IEEE Transactions on Automatic Control [171, 190], 2 papers in EMSOFT [77, 80], and 1 in CAV [79].

In formal verification and software engineering, the group has published 3 papers in CAV [89, 101, 146], 1 each in PLDI [142], TACAS [65], CONCUR [96], ECOOP [126], ESOP [85], FSTTCS [58], FASE [86], and FMCAD [144].

In logic and foundations, the group has published 4 papers in CONCUR [145, 115, 116, 87], 1 each in ICALP [66], CAV [82], and FSTTCS [64].

In addition, we had one AAAI paper on synthesis for geometry problems [13].

The papers [80] and [65] were nominated for best paper awards.

Software, tools, data and technology transfer. Vinayak Prabhu's work on conformance testing using Skorokhod distances is being used by Toyota Engineering North America and being evaluated at Bosch Research.

Teaching. Rupak Majumdar taught the (required) graduate course on Reactive Systems Verification at University of Kaiserslautern in Summer 2014. He co-taught a course on Foundations of Embedded Systems with Björn Brandenburg. He co-taught a seminar on Counting Complexity with Dmitry Chistikov and Rayna Dimitrova (in preparation for a subsequent TACAS paper).

External funding. Our research is supported in part by an ERC Synergy Award "ImPACT: Privacy, Accountability, Compliance, and Trust in Tomorrow's Internet," with co-PIs Michael Backes, Peter Druschel, and Gerhard Weikum. The project started in February 2015 and is funded for six years.

Additionally, our research is supported in part by industrial grants from Toyota ($75K annually for 2013, 2014, and 2015).

Invited talks and awards. Rupak gave invited talks at UC Berkeley, Microsoft Research, and RWTH University at Aachen. The paper "Abstractions from Proof" from POPL 2004 was awarded the ACM SIGPLAN POPL "Test of time" award in 2014.

Service. At MPI-SWS, Rupak chairs the graduate program.

Rupak served on the program committees of several conferences in the last two years, including LICS, EMSOFT, PLDI, and RTSS. Rupak co-organized the first Verification Mentoring Workshop at CAV 2015. Rupak is chairing RV 2015 and POPL 2016.

8.2 Research agenda

The Rigorous Software Engineering group studies both foundational principles and practical tools for the design and analysis of computer systems. Currently, the research in the group has focused on three different aspects: methodologies and tools for embedded controller design, foundations of infinite-state verification, and software verification.

8.2.1 Embedded Controller Design

In the area of embedded controller design, the research focus of the group is on automated co-design of controllers and their implementations. The following are some highlights from the last two years:

Metrics for hybrid systems Metrics on hybrid systems quantify the notion of similarity between behaviors and generalize notions of bisimilarity and trace equivalence from discrete systems to hybrid systems. While metrics on hybrid state spaces have been used to give semantics to hybrid systems, so far, the algorithmic computation of metrics as well as the use of metrics in conformance testing had not been studied. We have developed algorithms to compute metrics on timed and hybrid systems.

In [62], we designed algorithms to compute the edit distance between timed words and between timed systems modeled as timed automata.

In [143, 79], we considered the Skorokhod distance on hybrid traces. The Skorokhod distance computes a metric on traces that takes into account timing distortions in addition to differences in the continuous state. While it has been used before to give semantics to hybrid and probabilistic systems, and to define continuous bisimulation functions, algorithms to compute the distance were not known. In [143], we describe a polynomial-time algorithm to compute the Skorokhod distance between time-sampled traces completed by linear interpolation (called "polyhedral traces"). Computing the Skorokhod distance is non-trivial because the definition of the distance minimizes over an infinite family of continuous retiming functions. Our polynomial-time algorithm uses geometric characterizations of the space of solutions that were discovered in the study of Fréchet distances in computational geometry. In

[79], we implemented our algorithm in a tool for conformance testing of Simulink models. In collaboration with Jyo Deshmukh at Toyota, we evaluated our algorithm on a set of industrial control system benchmarks. Our implementation shows that the distance can be computed fast and captures the engineering intuition about "close" behaviors. We also characterize the distance using a timed linear time logic with signals and freeze quantifiers.

Reactive synthesis for continuous dynamics The second direction of work has been on automated synthesis of continuous controllers for temporal logic specifications. There are two main ways reactive synthesis is performed for continuous systems: by computing a "close" discrete abstraction and then using discrete synthesis techniques and by directly computing certificates.

We have extended the abstraction-based synthesis technique to stochastic continuous systems by introducing notions of ϵ-bisimulations for continuous stochastic dynamical systems. An ϵ-bisimulation of a continuous dynamical system is a discrete (finite-state) system such that every moment of the continuous stochastic system is matched in expectation by the discrete system to within a distance of ϵ. For continuous systems satisfying a certain well-formedness condition (called incremental input-to-state stability), finite ϵ-bisimulations exist for all $\epsilon > 0$. In [190], we introduce stochastic ϵ-bisimulations, give Lyapunov function characterizations of incremental stochastic stability, and show how a finite discrete abstraction can be computed.

In the second direction, in [80], we have extended certificate-based synthesis to ATL* properties by generalizing the notion of certificates from barriers or Lyapunov functions to certificates for general linear properties and extending a deductive proof system for discrete systems to the continuous case.

Currently, we are working on end-to-end systems that provide language support and automatic synthesis for systems that involve co-ordination among multiple robots. Our intent is to build high-level programming languages that allow end-users to program co-ordination strategies and have a scalable compiler that designs reactive controllers to achieve co-ordination. The reactive controller takes "lower level" robot motion plans and stitches them together to satisfy global temporal specifications. Filip Niksic is working on language design for co-ordinating state machines. Anne-Kathrin Schmuck is working on compositional and hierarchical reactive synthesis. Rayna Dimitrova is working on deductive approaches to controller synthesis.

8.2.2 Infinite-State Verification

Our second focus is in algorithmic foundations and practical tools for infinite state verification. Some highlights in the last two years are:

Analysis of asynchronous programs Asynchronous programming is an idiom to manage concurrent interactions through co-operative scheduling of tasks. They are used in many settings: in systems programs, in Javascript programming on the web, in low-latency smartphone applications, and in embedded systems. We have worked on theoretical foundations of asynchronous programming (decidability of expressive fragments) as well as on software analysis tools. In [85], we showed decidability of safety verification for an expressive model for asynchronous programs by encoding the model into Petri data nets. In [126], we developed a refinement type system to reason about Ocaml programs with asynchronous computation as well as mutable state. The refinement type system combines liquid types with concurrent separation logic. In [96], we present rely-guarantee reasoning principles for modular reasoning about asynchronous programs. We have also developed analysis algorithms for Android applications by defining a happens-before relation for the Android concurrency model [142]. In the Android model, multiple threads of execution run concurrently and can communicate via asynchronous tasks that are executed in FIFO order. We have used our happens-before relation to find race conditions in Android applications.

Model checking parameterized non-atomic networks We have characterized the complexity of verification parameterized systems with a designated "leader" process and an arbitrary number of "follower" processes that communicate with a shared, finite-valued register that does not have an atomic test-and-set operation. We show that safety and liveness verification in this setting are both NP-complete when the leader and followers are implemented by finite-state machines [88, 82]. From a verification perspective, this is surprising: the corresponding problem is PSPACE-hard in a non-parameterized setting and EXPSPACE-hard in the parameterized setting with atomic updates. From an expressiveness perspective, this shows the limited expressive power of parameterized models without atomic operations. We show that the verification problems get harder (PSPACE for safety, NEXPTIME for liveness) when processes are allowed to have unbounded stacks.

Language-theoretic results Finally, we have worked on the complexity of the equivalence problem for unary deterministic pushdown automata (DPDA) [66]. A unary DPDA is a DPDA over a unary alphabet. They accept exactly the unary regular languages; however, a unary DPDA can be exponentially more succinct than an NFA for the same language. Our results make a connection between unary DPDAs and compressed words to show that the equivalence problem for unary DPDAs is PTIME-complete. Our techniques also show better complexity bounds on several related problems, including better complexity bounds for some fragments of Presburger arithmetic.

8.2.3 Merging Logic and Probabilities

A third direction of our research is in combining reasoning with logic and probabilities. We have recently developed randomized algorithms that count the number of satisfying solutions of a formula modulo theories (the "#SMT problem"). For Boolean constraints, the problem of counting the number of satisfying assignments (#SAT) has been studied before, and recently, randomized approaches have been proposed. Our work extends the work on #SAT to work with SMT solvers. For theories such as linear real arithmetic, "counting" involves computing volumes of polytopes. Our approach uses classic results from complexity theory (in particular, that the complexity class #P is contained in $\mathrm{BPP}^{\mathrm{NP}}$) to reduce counting to a randomized algorithm that makes calls to an unmodified SMT solver.

We are exploring factored representations such as Bayesian networks in probabilistic verification. We have shown how independence assumptions on the noise in stochastic dynamical systems can naturally lead to Bayesian network representations and subsequently, better bounds on the error and a much faster analysis algorithm (compared to an exploded Markov chain representation).

We are currently using #SMT as a basis to design Markov Chain Monte Carlo procedures with a mix of hard and soft constraints. Our longer term research agenda is to introduce theory reasoning into machine learning models. We believe that such expressive models that incorporate logical and probabilistic reasoning can lead to precise reasoning about privacy and trust.

9 The Networks and Machine Learning Group

9.1 Overview

The report covers the period from November 2014 – July 2015. The group's research interests are in developing machine learning and large-scale data mining methods for the analysis, modeling and control of large real-world networks and processes that take place over them. The group is particularly interested in problems at the intersection of networks, information and society, with an emphasis on phenomena arising in the Web and social media.

Personnel. The group is led by Manuel Gomez Rodriguez. It currently has one graduate student, Utkarsh Upadhyay (advised since January 2015), and one postdoc, Isabel Valera (since March 2015). Over the reporting period, the group had three interns (Arman Sepehr, Erfan Tavakoli and Mohammad Reza Karimi) and two visiting graduate students (Charalampos Mavroforakis and Mehrdad Farajtabar), each typically staying for 3 months.

Collaborations. The group has collaborated with the social computing group (led by Krishna Gummadi) and the distributed systems group (led by Peter Druschel). Externally, the group has collaborated with researchers at Georgia Institute of Technology (Le Song, Hongyuan Zha), Max Planck Institute for Intelligent Systems (Bernhard Schölkopf), IIT Kharagpur (Niloy Ganguly, Sourangshu Bhattacharya), New York University (Martin Jankowiak) and Stanford University (Jure Leskovec).

Publications. During the reporting period (November 2014 – July 2015), group members have co-authored one NIPS [90], one AISTATS [91], one ICWSM [15], and several workshop [187, 92] publications. During this period, one JMLR journal paper [110] and one TOIS journal paper [111] have been conditionally accepted for publication.

Invited Talks. Manuel was an invited speaker at Princeton University, Ecole Polytechnique, IIT Kharagpur, University of British Columbia, Telefonica Research in Barcelona, Microsoft Research in New York, and Bell Labs in Murray Hill. He was invited to give a guest lecture at the University of California San Diego. He also taught two tutorials, one at WWW 2015 and one at KDD 2015.

Service. Internally, Manuel served on the faculty recruiting committee in 2015. Externally, Manuel served on the following PCs: SDM 2015, WWW 2015, ICWSM 2015, ICML 2015, KDD 2015, IJCAI 2015, WSDM 2015, CIKM 2015, and NIPS 2015. He has been a reviewer for the following journals: JMLR, Machine Learning, ACM TKDD, ACM TWEB, PLOS One, IEEE TKDE, IEEE TSP, IEEE TNSE and The Journal of Web Science. He was co-chair of the "Workshop on Networks – Processes and Causality" at DALI 2015 and the "Workshop on Diffusion, Activity and Events in Networks: Models, Methods and Applications" at WWW 2015. He is co-organizing a Machine Learning Summer School (MLSS) that will take place in May 2016, co-located with AISTATS 2016.

External funding. Isabel Valera has been awarded a Humboldt postdoctoral fellowship (June 2015 – May 2017).

9.2 Research agenda

The Networks and Machine Learning group develops machine learning and large-scale data mining methods for the analysis, modeling, and control of large real-world networks and complex processes that take place over them. The group is particularly interested in problems at the intersection of networks, information and society, with an emphasis on phenomena arising in the Web and social media.

Research at the Networks and Machine Learning group spans several dimensions: (1) developing realistic models of networks and processes that take place over them, as well as assessing their theoretical properties and limitations; (2) developing machine learning algorithms to fit such models and computational methods to influence processes over networks; and (3) validating these models and methods in massive real-world datasets of gigabyte and terabyte scale. In the long term, the group aims to provide computational tools with direct applications in a wide range of domains such as social and information sciences, business intelligence, marketing, epidemiology and national security, among many others.

In the reporting period, the group's research activities spanned the following broad areas: information acquisition, information reliability, opinion dynamics, network control and fairness on learning systems. For the work on fairness on learning systems, please see Krishna Gummadi's section of the report.

9.2.1 Information acquisition

The advent of social media and social networking sites is changing dramatically the way in which people acquire, consume and share information. Social networking sites such as Twitter, Tumblr and Pinterest have become global platforms for public self-expression and conversation, where users play the role of information curators by deciding which information to *post* and which other users in the network to *follow*. In the Networks and Machine Learning group, we are working on understanding different aspects of the information acquisition process in social media and social networking sites. In the reporting period, we have worked on three different projects:

1. In collaboration with the Networked Systems group led by Krishna Gummadi, we introduced a computational framework that allows us to quantify the efficiency of a user as an information curator within a social networking site [15]. Our framework is general and applicable to any social networking site with an underlying information network, in which every user *follows* others to receive the information they produce. We find that social media users are sub-optimally efficient at acquiring information and this lack of efficiency is a consequence of the triadic closure mechanism by which users typically follow other users in social networking sites.

2. In a recent collaboration with researchers at Georgia Institute of Technology, we proposed a probabilistic generative model for the joint dynamics of information diffusion and network evolution. The model successfully captures the following phenomena, observed in social networking sites: users often forward to their *followers* information they are exposed to via the users they follow, triggering the emergence of information *cascades* that travel through the network, and constantly create new links to information sources, triggering changes in the network itself over time. Our model provides a good fit to real data gathered from Twitter and provides more accurate predictions than alternatives. This work is under submission in one of the top machine learning conferences.

3. In a recently started project, led by graduate student Utkarsh, we are developing a probabilistic models of learners (i.e., users learning about a topic) in Q&A websites. The key idea is to leverage the users' interactions within the sites as a proxy for their level of expertise. As users increase their level of expertise, the type of questions they ask, answer,

upvote and downvote as well as the number of upvotes and downvotes their questions and answers receive will change. Moreover, the tags associated with these questions will help us determine whether a person becomes an expert in a narrow topic (i.e., posts share similar tags) or a generalist (i.e., posts have a great variety of tags). Preliminary experiments using data from StackExchange are promising.

9.2.2 Information reliability

In social media and social networking sites, information is often not professionally curated. As a consequence, its high-quality, relevance and credibility are at stake and need to be assessed to avoid the spread of misinformation and false rumors. The Networks and Machine Learning group is developing methods and models that help to assess information reliability. In the reporting period, we have worked on two different projects:

1. Manuel, in collaboration with researchers at Georgia Institute of Technology, developed a method to identify the *source* that originally published a rumor or piece of malicious information in the Web and social media [91]. Being able to do so is critical for curtailing its spread and reducing the potential losses incurred. However, this is a very challenging problem since typically only incomplete information traces are observed and we need to unroll the incomplete traces into the past in order to pinpoint the source. The method first first learns a probabilistic diffusion network model based on historical diffusion traces and then identifies the source of an incomplete diffusion trace by maximizing the likelihood of the trace under the learned model. Experiments on both large synthetic and real-world data show that the framework can effectively "go back to the past", and pinpoint the source node and its initiation time significantly more accurately than other state of the art methods.

2. In a recently started project, executed in collaboration with researchers at the Max Planck Institute for Intelligent Systems and Georgia Institute of Technology, we are developing a generative point process model of information reliability and source (site) trustworthiness. The model is based on the following simple, key idea: trustworthy sites publish reliable information that is not refuted over time and, in contrast, non-trustworthy sites publish information that is eventually refuted. Thus, every time a site publishes a piece of information, our model instantiates a survival process, modulated by the trustworthiness of the

site, which models the time this information goes non-refuted. Our model will allow us to assess the reliability of particular pieces of information and the trustworthiness of sites at any given time as well as understand their temporal evolution.

9.2.3 Opinion dynamics

Social media and social networking sites are increasingly used by people to express their opinions, i.e., give their "hot takes", on the latest breaking news, political issues, sports events, and new products. As a consequence, there has been increasing interest in leveraging social media and social networking sites to measure and predict *opinions* as well as understand *opinion dynamics*. In the Networks and Machine Learning group, we are working on developing realistic models of opinion dynamics in social media and social networking sites that do not only fit fine-grained opinion data, but also provide accurate predictions of the individual users' opinions over time.

In a recent collaboration with researchers at IIT Karagphur, we proposed a continuous time generative model that captures two intuitive key processes driving opinion dynamics: *informational influence* and *social influence*. The former process accounts for the idea that the users may form or update their opinion about a particular topic by learning from the information and opinions that their friends share. The latter process, i.e., social influence, accounts for the impact that different users have on the activity level in the network. In other words, some users may either express their opinions more frequently than others, or be more influential and thus trigger a greater number of replies every time they express their opinion. Remarkably, for several instances of our modeling framework, we identify the conditions under which opinions converge to a steady state of consensus or polarization. In such cases, we derive a closed-form expression for the relationship between the users' steady state opinions and the initial opinions they start with. Experiments on real data gathered from Twitter, Reddit and Amazon show that our model provides a good fit to the data and more accurate predictions than several state of the art models of opinion dynamics. This work is under submission in one of the top machine learning conferences.

9.2.4 Network control

Over the last few years, Manuel and his collaborators have developed realistic models of information diffusion, social activity and opinion dynam-

ics [106, 108, 109]. Can we leverage these models to evaluate the extent to which external interventions can achieve a desired goal? For example, social media and social networking sites may like to incentivize a few *influential* individuals to use their sites hoping that they trigger further use by other people. Large corporations and governments may design orchestrated campaigns in which a number of people are incentivized to swamp the online debate about an ongoing event, influencing people's opinion about it. Social media users would often like to post messages at times when they get the highest level of exposure among the users that *follow* them. In the Networks and Machine Learning group, we are developing a family of methods and algorithms that allow us to investigate the extent to which the above mentioned goals can be achieved. In the reporting period, we have worked on three different projects:

1. In collaboration with researchers at Georgia Institute of Technology, we developed a computational framework for activity shaping in social media and social networking sites [90]. In the activity shaping problem, one aims to drive the overall usage of a site (or service) to a certain level (e.g., at least twice per day per user) by incentivizing a small number of *influential* users to increase their usage. We then exploit a key property of our framework to design a very efficient family of algorithms, based on projected gradient descent, that can easily scale up to networks with tens of thousands of nodes. Experiments on event data gathered from Twitter show that our method can steer the activity of the network more accurately than alternatives.

2. In a recently started project, executed in collaboration with researchers at IIT Karagphur, we are developing computational framework that leverages the model of opinion dynamics described above to quantify to which extent a small set of (incentivized) users can steer users' opinions in a social network to a given state. The motivation to do so is the increasing concerns on the emergence of orchestrated campaigns, often attributed to governments and large corporations, in which a number of people are paid to swamp the online debate about an on-going event, influencing people's opinion about it. Preliminary experiments on synthetic and real data gathered from Twitter already indicate that our framework can accurately determine the quality of a set of incentivized users at a given opinion shaping task.

3. In a recently started project, led by postdoctoral fellow Isabel Valera and visiting graduate student Mehrdad Farajtabar, we are develop-

ing methods for *smart* broadcasting of information in social media and social networking sites. Given a user with a set of followers, our methods will find the times when this user should post information to maximize exposure across her followers, i.e., maximize the time this information remains at the top of their followers' walls. This project could have a tremendous practical impact on fields such as marketing, social awareness and emergency response.

10 The Software Analysis and Verification Group

10.1 Overview

The report covers the period from November 2013 to July 2015. The software analysis and verification group works on the verification of complex software systems and their components, such as compilers and concurrent algorithms. It does so by developing theories and tools for rigorously applying formal reasoning principles to build correct software systems.

Personnel. The group consists of Viktor Vafeiadis (faculty), Ori Lahav (postdoc, since September 2014) and PhD students Marko Doko and Soham Chakraborty. Over the reporting period, the group had six interns (Dan-Cristian Andronic, Nick Giannarakis, Mengda He, Xiao Jia, Wei Li, and Craig McLaughlin) each typically staying for 3 months.

Collaborations. The group collaborates with Derek Dreyer's and Rupak Majumdar's groups. We also collaborate with researchers at Cambridge (Sewell, Sezgin), IMDEA (Gotsman), INRIA (Zappa Nardelli, Balabonski), Leuven (Jacobs), SNU (Hur), TAU (Rinetzky), and UPenn (Zdancewic).

Publications. The group publishes regularly in the top conferences and journals of its field. During the reporting period, group members have co-authored one LMCS [61] journal article, one POPL [177], two PLDI [172, 122] one OOPSLA [176], one ICALP [133] one ECOOP [126], one CPP [120], and one USENIX ATC [139] conference papers. Many of these publications come with machine-checked proof developments in Coq, which are available online.

Teaching. In the 2014 summer semester, Viktor Vafeiadis cotaught the *Concurrency Theory* core course with Roland Meyer at TUKL.

External funding. The group's research has been partially funded by the European Commission's FP7 FET young explorers grant ADVENT (April 2013 – April 2016). The grant currently funds Ori Lahav's postdoc.

Service. In the reporting period, Viktor has served on the program committees of ICFEM 2015, TASE 2015, PDP/4PAD 2015, and ICFEM 2014, as co-chair of the Coq 2014 workshop, and as publicity chair for POPL 2014 and POPL 2015.

10.2 Research agenda

Our research concerns the development of mathematical theories and tools for formally reasoning about software. It aims at improving software quality by making it possible to build provably correct software components. This involves coming up with rigorous mathematical specifications of programming languages and of software components, such as data structure libraries and compilers, developing custom proof techniques for proving adherence to those specifications, as well as improving the underlying general-purpose verification infrastructure.

Most of our work has focused on reasoning about concurrent programs running under the C/C++ weak memory model. We have attacked this problem from two angles: (1) by developing sound program logics that are suitable for reasoning about program correctness, and (2) by reasoning about the correctness of program transformations, such as compiler optimizations. As a result of this ongoing endeavor, we have identified a number of problems in the C/C++ memory model definition and have proposed fixes for them.

10.2.1 Correcting the C memory model

Concurrency semantics [177] The 2011 revisions of the ISO C and C++ standards introduced a fairly complex weak memory model that stipulates the semantics of concurrent memory accesses. One of the main design goals of the C/C++11 memory model was to support the many program transformations and optimizations that compilers do.

In joint work with Francesco Zappa Nardelli's group at INRIA, we showed that the C/C++11 memory model does not meet this objective. The formal memory model definition has serious deficencies that rule out many common source-to-source program transformations (such as expression linearisation and "roach motel" reordering) that modern compilers perform and that are deemed to be correct. In response, we proposed a number of possible local fixes, some strengthening and some weakening the model, and evaluated them by determining which program transformations are valid with respect to each of the patched models. (We provided formal Coq proofs of their correctness or counterexamples as appropriate.)

Integer-pointer casts [122] Another limitation of the ISO C standard is that it does not specify the semantics of many valid programs that use non-portable idioms such as integer-pointer casts. Recent efforts at formal definitions and verified implementation of the C language have inherited

this feature. By adopting high-level abstract memory models, they validate common optimizations but do not support reasoning about low-level code relying on the behavior of common implementations.

In joint work with Jeehoon Kang, Chung-Kil Hur, William Mansky, Dmitri Garbuzov, and Steve Zdancewic, we developed the first formal memory model that allows many common optimizations and fully supports operations on the representation of pointers. All arithmetic operations are well-defined for pointers that have been cast to integers. Crucially, our model, which is formalized in Coq, is also simple to understand and program with.

10.2.2 Program logics for weak memory consistency

Weak memory models, such as the C/C++11 model, formalize the behaviors that may be observed in multithreaded programs running on modern hardware. These behaviors include inconsistent outcomes that cannot be observed simply by interleaving the memory accesses of the various threads of a program. As a result, not only do they complicate the already-difficult task of reasoning about correctness of concurrent code, but they also render impotent the sophisticated formal methods that have been developed to tame concurrency, which almost universally assume a strong (i.e., sequentially consistent) memory model.

To enable formal reasoning about weak memory consistency, we have focused on the release-acquire fragment of the C/C++11 model and have developed two program logics for it.

OGRA [133] Our first logic, called OGRA, is based on the well-known Owicki-Gries method. We showed that even in the absence of auxiliary variables, the Owicki-Gries method for verifying concurrent programs is unsound under even the strongest weak memory models. By strengthening its non-interference check, however, we obtain a useful logic that is sound under release-acquire consistency. We have applied our logic to a number of small challenging examples, and have carried out a preliminary investigation for automating the construction of proofs in the logic.

GPS [176, 172] Our second logic, called GPS, is a more sophisticated program logic providing a full-fledged suite of modern verification techniques (namely, ghost state, protocols, and separation logic) for high-level, structured reasoning about release-acquire consistency. As a result, the soundness proof of GPS is rather complicated and was carried out in Coq.

We have applied GPS to verify several challenging examples drawn from the Linux kernel as well as lock-free data structures. The most challenging among those was an implementation of a concurrent singly-linked list protected by the *read-copy-update* (RCU) synchronization mechanism. Our verified RCU implementation used release-acquire synchronization, which although stronger than what is required by real RCU implementations, it is nonetheless significantly weaker than the assumption of sequential consistency made in prior work on RCU verification. Ours was the first formal correctness proof for an RCU implementation under a weak memory model.

10.2.3 Other verification techniques

Verifying lock-freedom [120] Lock-freedom is a liveness property satisfied by most non-blocking concurrent algorithms. It ensures that at any point at least one thread is making progress towards termination; so the system as a whole makes progress. As a global property, lock-freedom is typically shown by global proofs or complex iterated arguments. We showed that this complexity is not needed in practice. By introducing simple loop depth counters into the programs, we can reduce proving lock-freedom to checking simple local properties on those counters. We have implemented the approach in Cave and have formalized its metatheory in Coq.

Asynchronous liquid separation types [126] In joint work with Johannes Kloos and Rupak Majumdar, we developed a type system for reasoning about asynchronous programs manipulating shared mutable state. Our type system combines two ideas: *refinement types*, such as $\{x : \mathsf{int} \mid x > 5\}$, and *ownership of memory cells* as in concurrent separation logic. It is expressive enough to allow simple data structure invariants to be specified by the user, and guarantees the absence of races and the preservation of any user-specified invariants. In more detail, our types are indexed by sets of resource names and the type system tracks the effect of program execution on individual heap locations and task handles. This, in particular, allows making strong updates to the types of heap locations. Moreover, our types track ownership of shared state across concurrently posted tasks and allow reasoning about ownership transfer between tasks using permissions. We have implemented type inference for our type system and used it to prove complex invariants of asynchronous OCaml programs.

Aspect-oriented linearizability proofs [61] Linearizability of concurrent data structures is usually proved by monolithic simulation arguments

relying on the identification of the so-called linearization points. Regrettably, such proofs, whether manual or automatic, are often complicated and scale poorly to advanced non-blocking concurrency patterns, such as helping and optimistic updates. In joint work with Thomas Henzinger and Ali Sezgin, we proposed a more modular way of checking linearizability of concurrent queue algorithms that does not involve identifying linearization points. We reduced the task of proving linearizability with respect to the queue specification to establishing four basic properties, each of which can be proved independently by simpler arguments. As a demonstration of our approach, we verified the Herlihy and Wing queue, an algorithm that is challenging to verify by a simulation proof.

Detecting consistency requirements [139] In joint work with Cheng Li and others, we applied standard verification technology to assist programming of (weakly consistent) replicated online services. The question we addressed was how to decide which consistency level is needed for each transaction. We developed a tool, called SIEVE, that relieves Java programmers from this error-prone decision process, allowing applications to automatically extract good performance when possible, while resorting to strong consistency whenever required by the target semantics. Our tool takes as input a set of application-specific invariants and a few annotations about merge semantics, and performs a combination of static and dynamic analysis to determine when it is necessary to use strong consistency to preserve these invariants and when it is safe to use causally consistent commutative replicated data types (CRDTs). We evaluated SIEVE on two web applications and showed that the automatic classification overhead is low.

Part II
Adjunct Research Groups

11 The Information Security and Cryptography Group

11.1 Overview

The report covers the period November 2013 – July 2015. The research of this group focuses on the theoretical foundations and applied aspects of information security, privacy, and cryptography. Major research topics have included: the design and analysis of security protocols, privacy and anonymity, linking formal methods and cryptography, and novel approaches for OS and software security.

Personnel. The group is led by Max Planck Fellow Michael Backes. Michael Backes additionally has the chair for information security and cryptography (IS&C) at Saarland University, and he is the director of the Center for IT Security, Privacy, and Accountability (CISPA). The group until recently had one postdoc (Dario Fiore), who is now faculty at IMDEA, Spain. The Max Planck fellowship is currently being used for funding internships and visiting researchers (currently Kangjie Lie from Georgia Tech), who complement the IS&C university group. This group currently consists of nine graduate students (Fabian Bendun, Pascal Berrang, Erik Derr, Praveen Manoharan, Sebastian Meiser, Ivan Pryvalov, Milivoj Simeonovski, Malte Skoruppa, and Philipp von Styp-Rekowsky).

Collaborations. During the review period, the group has engaged in a number of successful collaborations. Internally, the group is collaborating with the groups of Peter Druschel, Deepak Garg, and Krishna Gummadi. Moreover, the group has just started a collaboration with the group of Rupak Majumdar, as part of the research agenda of the ERC Synergy Grant. Externally, we are working with colleagues at the University of Maryland (Jonathan Katz, Michelle Mazurek), Cornell University (Johannes Gehrke), Penn State University (Patrick McDaniel), University of Waterloo (Ian Goldberg), IMDEA (Dario Fiore), Technical University of Darmstadt (Marc Fischlin), University of Bonn (Matthew Smith), and ETH Zurich (Raphael Reischuk).

Publications. The group has published regularly in the top conferences and journals of its field. During the reporting period (21 months), group members have co-authored 22 conference publications: one S&P [16], three

CCS [25, 33, 30], two Usenix Security [39, 22] two CSF [36, 19], two ESORICS [24, 38], one PODC [17], two ACSAC [21, 20], two POST [37, 18], three ACNS [23, 27, 34], one WPES [26], one SAC [28], one HotPETS [32], and one STM [29] publication. Moreover, we have published two journal papers: one at JCS [31] and one at MPCS [35].

Software, tools, and data. The group strives to make its software, tools, and data sets publicly available to the extent possible. A particular highlight in this respect has been AppGuard – a novel tool developed in 2013 for privacy protection of end users on Android devices. More than 1,5 million end users have downloaded the freely available version of AppGuard in the last two years to protect themselves against malicious apps on their Android mobile devices. This gave rise to a commercial version of AppGuard with extended functionality that is marketed by the spinoff SRT, see below. A follow-up approach called Boxify has just been accepted at Usenix Security [22]. Moreover, the source codes of ASF [21] and Scippa [20] were openly released on the group's project website.

Patents and technology transfer. We have just filed a patent related to Boxify within the spinoff company SRT that markets the commercial version of AppGuard. We expect to bring Boxify to market in the spinoff soon.

Press. Our research on providing anonymity guarantees for Tor (MATor [33]) has been exhaustively covered in the public press, in particular by ZDnet, heise, and Deutschlandfunk. Similarly, our research on Appguard and its successor Boxify have been and are stil widely featured by the press (Frankfurter Allgemeine Zeitung, Sueddeutsche Zeitung, c't-Magazin, Technology Review, Hannoversche Allgemeine, and various TV progams such as ORF, NDR Logo, etc.)

Teaching. Michael Backes has initiated the Cybersecurity Bachelor program at Saarland University in 2014. He taught the beginner lecture on cyber security in 2014, the core security lecture in 2013, as well as advanced courses on smartphone security in 2013 and 2014. Moreover, he held two graduate and two undergraduate student seminars in the last two years.

External funding. The research of the group has been partially funded by an ERC Synergy Grant *ImPACT: Privacy, Accountability, Compliance*

and Trust in Tomorrow's Internet, by the Excellence Cluster on Multimodal Computing and Interaction (MMCI), and by the Center for IT Security, Privacy, and Accountability (CISPA).

Invited talks, awards, and honors Michael Backes gave invited talks at the SnT in Luxembourg, at the Leopoldina Symposium on Science Freedom and Responsibilities, and at the Max Planck Forum in Saarbrücken and in Nürnberg. Jointly with Peter Druschel, Rupak Majumdar, and Gerhard Weikum, he received the ERC Synergy Grant in 2013 – one of Europe's most distinguished research awards. He was selected as a member of the German Academy of Science and Engineering (AcaTech) in 2015, and is currently the youngest of its members. In 2014, he received an IEEE Outstanding Community Service Award, the CS Teaching Award of Saarland University for the best CS lecture in 2014, and the Teaching Award of the State of Saarland. Moreover, he was named one of Germany's digital minds by Germany's federal minister of science and education Johanna Wanka in 2014.

Service. Michael Backes initiated the new conference series *IEEE European Symposium on Security & Privacy (EuroS&P)*. EuroS&P is set up to become the European sister conference of IEEE S&P – a task that the security community has been pursuing for more than 10 years, now finally successfully. Michael is the chair of the Steering Committee of EuroS&P, and he will serve as the PC Chair of the conference's first edition in 2016. Michael Backes was the program co-chair of the IEEE Symposium on Security & Privacy (S&P) in 2013 and 2014. He currently serves on the steering committee of IEEE S&P, IEEE EuroS&P, ACM CCS, IEEE CSF and ESORICS. In the reporting period, Michael in particular served on the following major program committees: ESORICS 2013, IEEE S&P 2014, IEEE CSF 2014, ESORICS 2014, ESORICS 2015, and ACM CCS 2015. He is furthermore on the Editorial Board of the Journal on Foundations and Trends in Security and Privacy.

11.2 Research agenda

The group's research interests are in theoretical foundations and applied aspects of information security, privacy, and cryptography. Contributions have been made in the following areas: (1) the design and analysis of security protocols; (2) privacy and anonymity; (3) linking formal methods and cryptography; (4) and novel approaches for OS and software security. In the

last two years, the group's research interest have focused increasingly on the area of privacy assessment and privacy-preserving computation. We expect to further concentrate our research efforts on this area, which corresponds to the research agenda of the ERC Synergy Grant. In the following sections, we highlight our contributions in the four aforementioned research areas.

11.2.1 Design and Analysis of Security Protocols

Designing and analyzing security protocols is known to be difficult, and work aiming at rigorous security guarantees by design or by trustworthy analysis started soon after the first protocols were developed. In our recent research, we in particular designed novel protocols for securely and authentically outsourcing computation, for more efficient multi-party computations under additional assumptions, and for ensuring data authenticity in highly dynamic web scenarios. Moreover, we devised a novel type systems for analyzing security protocols.

Verifiable outsourced computation. Our main goal here is the development of cryptographic schemes for secure and authentic delegatable computation. More concretely, we address the problem in which a client incrementally stores a large amount of data with an untrusted server in such a way that, at any moment, the client can ask the server to spontaneously compute a function on some portion of its outsourced data. In this scenario, the client must be able to efficiently verify the correctness of the result despite no longer knowing the inputs of the delegated computation. We have proposed a scheme that achieves these goals for computations of quadratic polynomials over multiple variables and which achieves constant-time verification [25]. We have extended this work by developing a highly efficient succinct zero-knowledge proof scheme (ZK-SNARK), which enables privacy-preserving proofs for outsourced computation over authenticated data [16]. Our benchmarks show that our approach is almost as efficient as state-of-the-art approaches without such authentication guarantees.

Multi-party computation. We have developed an asynchronous multi-party computation scheme that significantly improves the communication complexity compared to previous schemes by additionally relying on a small piece of trusted hardware (a trusted incrementer called TrInc) [17]. In this paper, we moreover use non-equivocation to construct the first asynchronous verifiable secret sharing (AVSS) scheme with $t < n/2$, which is of independent interest to threshold cryptography.

Authenticating data in dynamic web scenarios. We have developed

the WebTrust system [27] which provides data authentication properties for interactive content of web pages. WebTrust constitutes the first comprehensive authenticity and integrity framework that allows on-the-fly verification of static, dynamic, and real-time streamed Web content from untrusted servers, in particular for content that is remotely loaded from other servers. We conducted extensive benchmarks that demonstrate the efficiency of our approach.

Type system for protocol analysis. We have developed a novel type system that combines prior work on refinement types, with union, intersection, and polymorphic types, and with the novel ability to reason statically about the disjointness of types [31]. As a result, we can statically characterize novel properties (such as authenticity and integrity achieved by showing knowledge of data) and reason about protocols based on asymmetric cryptography and on zero-knowledge proofs; statically reasoning about these protocols is out of scope of previous approaches.

Miscellaneous. We have developed a generic data lineage framework called LIME for reasoning about data flow across multiple entities in scenarios that guarantee privacy on a contractual level [29]. We have developed a novel approach for strengthening the security of e-voting protocols using out-of-band communication on mobile devices [26].

11.2.2 Privacy and Anonymity

The advent of social networks and digital capture, and the resulting monetizing of personal information by means of targeted advertisement has created unprecedented threats to user privacy.

Anonymous communication. Our main goal in this area is to rigorously assess the anonymity guarantees that the currently most-widely deployed anonymous communication protocol Tor offers. We have developed a light-weight live-monitor called MATor that computes provably accurate anonymity bounds for individual users under structural corruption of Tor nodes [33]. This work is grounded in the framework AnoA that we previously developed for the analysis of anonymous communication protocols. For capturing timing-related attacks (in particular against anonymous communication protocols), we have moreover extended the time-agnostic cryptographic model underlying AnoA to adversaries that can observe and reason about timing behavior. The resulting model, called TUC, contains comprehensive composability results that enable a modular protocol analysis [36]. Moreover, we have shown how to complement the onion construction used in

Tor by an accountability mechanism [23]. This mechanism offers Tor nodes technical means to prove that their routed traffic originated at a different source, hence exonerating them from false accusations.

Novel privacy-enhancing protocols. We have developed two privacy enhancing systems. The first system constitutes a generic privacy-preserving accountable computation system that protects sensitive data while ensuring that any computation that deviates from the specification can be irrefutably linked to the malicious node that performed this computation [24]. Moreover, our approach provably allows honest nodes to disprove false accusations. We have proposed an efficient instantiation of this system for comprehensive classes of computations, based on the recently proposed concept of ZK-SNARKs. The second system, called X-pire2 [28], protects personal data in open environments by realizing a digital expiration date for personal images under realistic assumptions. While prior approaches for realizing a digital expiration date were vulnerable to the data duplication problem (i.e., copying data before the expiration date was reached invalidated the expiration date), X-pire2 utilizes state-of-the-art trusted computing technologies to provide robust protection even against attackers creating such digital copies of data. The system can be used for existing platforms such as Facebook, Google+ and Flickr.

Formalization of privacy case law. We have introduced PriCL, the first framework for formalizing privacy case law that is amenable to automated reasoning [18]. Considering case laws in formal frameworks for privacy law essentially strives to serve two purposes: first, assisting judges in deriving logically consistent court decisions; and second, determining the legal categorization of a given term (e.g., whether the billing address is public or nonpublic personal information in sense stated in the Gramm-Leach-Bliley Act (GLBA). PriCl is parametric in the underlying logic (e.g., LTL) and is sufficiently expressive for court decisions, their justification, the circumstances in which the justification applies, and the court hierarchies. We have identified an efficiently decidable logic subclass and have provided efficient algorithms for major reasoning tasks for those cases, e.g., for deducing legal permissions or extracting norms.

Miscellaneous. We have developed the universal framework Oblivion [167] to support the automation of the *Right to be Forgotten*, as recently introduced by the European Court of Justice, in a scalable, provable, and privacy-preserving manner. We have developed a cryptographic notion for quantifying the security guarantees of systems that use imperfect randomness [34], and thereby demonstrated that even imperfect randomness can

suffice to obtain security guarantees that are useful for many privacy enhancing technologies.

11.2.3 Linking Formal Methods and Cryptography

A successful line of research explores the automation of security proofs while abstracting cryptographic operations into simple equations on terms. To ensure that these abstractions misses no attacks, so-called computational soundness results have been established. This line of research has recently focused on interactive cryptographic primitives and on stronger security properties.

Computational soundness for interactive primitives. Our main goal here was to establish, for interactive primitives, a strong connection between the existing cryptographic notions of interactive primitives expressed in the UC framework and computationally sound symbolic abstractions. We showed that for a large class of interactive primitives, the existence of a securely realized ideal functionality (in the sense of the UC framework) canonically leads to a symbolic abstraction that is computationally sound [38]. Our result is parametric in a set of non-interactive computationally sound symbolic abstractions and comprises an arbitrary number of interactive primitives. Our result holds for arbitrary equivalence properties for single-threaded programs.

Computational soundness of equivalence properties for non-interactive primitives. Since computational soundness results typically pertain to trace properties only, we aimed at extending these results to the set of equivalence properties, both for individual cryptographic primitives and for generic settings. To this end, we developed a proof technique that identifies sufficient conditions so that computational soundness of trace properties implies computational soundness of equivalence properties for pairs of protocols that have the same control flow, so-called uniformity of bi-processes [37]. Using this result, we managed to prove computational soundness of malleable ZK proofs for uniformity of bi-processes [19].

Quantifying information flow in cryptographic systems. Our objective here was to relate quantitative information flow in abstractions of cryptographic systems with their cryptographic instantiations. To this end, we developed a novel notion called transmissible information that is suitable for reasoning about information-theoretically secure (or non-cryptographic) systems, as well as about cryptographic systems (with their polynomially bounded adversaries, error probabilities, etc.) [35]. We showed that trans-

missible information carries over from abstract protocol specifications to concrete instantiations if these instantiations securely realize the abstractions in the sense of the UC framework.

11.2.4 Operating System and Software Security

Operating systems are the foundation for building secure and trusted (end-user) systems. However, recent research has demonstrated that contemporary operating systems do not fulfill the users' expectations in two particular domains: protection of data store on smart devices such as tablets and smartphones, and protection against state-of-the-art software exploitation techniques. In this line of research we successfully contributed different solutions to increase the dependability of operating systems in these two domains.

Security and privacy protection at different layers of the Android software stack. The security architecture of modern mobile devices has been shown to be insufficient to effectively protect the plethora of users' private information stored on those devices. Moreover, it specifically fails in providing any means to adequately address the higher security requirements when those devices are used for business or government purposes—increasingly common use-cases today. To improve on this situation—with strong focus on the popular, open-source Android OS—we first introduced a novel security extension to the Android OS kernel [20], which provides extensive provenance information on inter-process communication. Such provenance information is the cornerstone for any access control enforcement on Android, including the default permission system and various related work. Failing in providing sufficient provenance information has in the past opened an attack surface for different privilege escalation attacks. Second, we introduced a security framework for Android's middleware and kernel layer [21], which allows developers and users to programmatically instantiate various security models. Thus, it accommodates for making Android more suitable for deployment in security contexts with varying security requirements, such as business contexts.

However, Android security extensions are notorious for low deployment rates on common end-user devices due to the high technical expertise required from end-users for installing them. Thus, in our most recent work [22] we introduced a novel line of research for deploying Android security extensions that is based on application virtualization. Our technique allows securely isolating apps on Android in a dedicated, secure runtime environment and integrating security extensions into the virtualization layer without

the need to modify the Android software stack, making our solution highly portable and easily deployable.

New memory management techniques to efficiently and effectively defeat code reuse attacks. Modern operating systems use sophisticated memory management techniques to prevent malicious code from exploiting software security vulnerabilities, such as buffer overflows. Among the most important countermeasures today is memory address layout randomization (short ASLR). However, highly efficient, fine-grained ASLR has two important disadvantages. First, it prevents code sharing among different processes, a well-established mechanism to optimize memory usage. Second, it still can be circumvented through a code reuse attack by an attacker that gains insight into the memory layout, e.g., through a memory disclosure attack. We presented two new memory management techniques that address those problems. First, we introduced a fine-grained memory randomization technique on a per-process level [39] that does not interfere with code sharing. Executables and libraries built with our system feature "memory-layout-agnostic code." Furthermore, it is the first solution that offers comprehensive protection against state-of-the-art code reuse attacks (so called JIT-ROP) and to demonstrate that fine-grained memory randomization is feasible without forfeiting the enormous memory savings of shared code. Second, we introduced a new primitive we call Execute-no-Read (XnR) [30], which ensures that code can still be executed by the processor, but at the same time code cannot be read as data. This primitive ultimately prevents the self-disassembly (i.e., memory disclosure) that is necessary for state-of-the-art code reuse attacks to work. To the best of our knowledge, XnR is the first approach to prevent memory disclosure attacks of executable code and JIT-ROP attacks in general.

12 The Dependable Systems Group

12.1 Overview

This section of the report describes the activities and achievements of the Dependable Systems Group, led by Rodrigo Rodrigues, and whose research focuses on building reliable software systems. Despite the fact that Rodrigo Rodrigues left the institute before the beginning of this reporting period, the group's work is still ongoing since there are a number of MPI-SWS students who are advised by him. As such, this section of the report focuses only on the subset of the work conducted by Rodrigo Rodrigues where one or more MPI-SWS students are involved.

Personnel. During this period, the group comprised four doctoral students, namely Pramod Bhatotia, Nancy Estrada, Pedro Fonseca, and Cheng Li.

Collaborations. The members of the group had joint publications and collaborations with the following groups: Distributed Systems (led by Peter Druschel), Real-Time Systems (led by Björn Brandenburg), Software Analysis and Verification (led by Viktor Vafeiadis) and Robust Systems (led by Allen Clement).

Externally, group members have collaborated with researchers at Microsoft Research, CMU, Nova University of Lisbon, and the MMCI.

Publications. Group members have co-authored papers in the following conferences: EuroSys [156], ASPLOS [43], OSDI [94], Middleware [42], Usenix ATC [139], and the PaPoC workshop [138]. In addition, a chapter in a book edited by CRC Press was also co-authored by members of the group [44].

Awards and honors The Middleware 2014 publication [42] won the best student paper award.

Systems and tools We have made the SKI [94] tool available, and it was used by Linux kernel developers for testing the kernel code. Part of the iThreads [43] package (namely a tracing library) has been made available.

Teaching. Pramod Bhatotia taught a course entitled Big Data Systems in February 2015. A total of 48 students signed up for this course.

Invited talks. Pedro Fonseca gave talks at Northeastern University, Yale University, Columbia University, MIT, IST-Lisbon, and EPFL.

Cheng Li gave talks at IMDEA Software Institute, Facebook, Huawei European Research Center, Microsoft Research Asia, IBM Research China.

Pramod Bhatotia gave talks at TU Dresden, Universitaet Paderborn, Johannes Gutenberg-Universitaet Mainz, Bell Labs Germany, NEC Research Heidelberg, EPFL, Telefonica Research Barcelona, IMDEA Networks Madrid, Microsoft Research Cambridge, and NOVA University of Lisbon.

In addition Pramod Bhatotia's work on the Slider system [42] was presented at the Hadoop Summit'15, in Brussels, April 2015.

Service. Pramod Bhatotia was a student representative for MPI-SWS in the Max Planck Society (2011-2013). Nancy Estrada organized a reading group and group meetings for the systems and networking groups (2013-2014).

Degrees. Pramod Bhatotia concluded his doctoral studies in April 2015, and Pedro Fonseca defended his doctoral thesis in June 2015.

12.2 Research agenda

The research of the members of the group focused on two broad subjects.

We are researching better ways to achieve a balance between **consistency and performance**, by coming up with principles and designing systems that make use of fast, weakly consistent operations whenever that does not jeopardize the required semantics that users should perceive, and only resort to slower, strongly consistent operations when that is necessary for enforcing those semantics.

We are also researching new ways to make **parallel and concurrent systems** perform better, namely by taking advantage of incremental computations in order to update the output efficiently and incrementally as the input evolves. In addition to improving the performance of these systems, we are also looking at ways to make them more reliable by improving the methods for developing these systems, e.g., by building new testing tools or new abstractions that programmers can leverage in the construction of parallel and concurrent applications.

Part III
Former Research Groups

13 The Robust Systems Group

13.1 Overview

The report covers the period from October 2013 – September 2014. The group's research interests broadly span the foundations of distributed systems with special attention paid to fault tolerance, reliability, and consistency.

Personnel. The group was led by Allen Clement and currently has three graduate students (Natacha Crooks, joined July 2013; Nancy Estrada, switched from Rupak Majumdar's group August 2013; and Reinhard Munz, joined when Umut Acar left August 2012).

The group disbanded when Allen Clement left for Google in September 2014. At this time the students moved on:

Reinhard Munz transferred to the group of Paul Francis in May 2014, Nancy Estrada moved to Lisbon with Rodrigo Rodrigues in September 2014, Natacha Crooks moved to the University of Texas in September 2014.

Collaborations. The group was collaborating with the Real-Time Systems Group (led by Björn Brandenburg) and the Software Analysis and Verification Group (led by Viktor Vafeiadis).

Externally the group was engaged in collaborations with researchers at Cambridge University, Grenoble INP, Microsoft Research (Cambridge), Universidade Nova de Lisboa, and the University of Texas at Austin.

Publications. During the reporting period the group has published at top conferences and journals in its field. Group members published papers at Eurosys [156, 102], USENIX [139], and PaPoC [138].

Teaching. No teaching took place during this period.

Service. Allen was on the program committee for OSDI 2014.

13.2 Research agenda

Computer systems play an important and pervasive role in modern life. It is important that these systems work properly, e.g., that they are *robust* to a range of environmental and adversarial factors. The Robust Systems Group

studies the theoretical foundations of, and practical design and implementation issues for, robust systems.

The work of the group during the reporting period has addressed three distinct factors that may lead to systems not working as intended: Byzantine failures, concurrency and geo-replication, and Sybil attacks.

13.2.1 Byzantine fault tolerance

One obvious challenge to building robust distributed systems is the fact that individual components can fail in unexpected ways. The *Byzantine* fault model allows for arbitrary failures and is attractive, in principle, due to this generality. Conventional, and technically correct, wisdom states that Byzantine fault tolerant systems require at least $3f + 1$ processes to tolerate up to f arbitrary faults.

The group's work in this area strives to identify and understand theoretical conditions for when fewer processes suffice and to provide practical Byzantine fault tolerant implementations.

Theory. A fundamental problem with the Byzantine failure model is that it is too general; it makes no assumptions on how individual processes may fail or how multiple faulty processes may coordinate. While it is true that any system designed to be robust to Byzantine faults will be robust under any stronger fault model, the generality of the Byzantine failure model does come at a cost. The group's work on the theoretical side of Byzantine fault tolerance focuses on understanding the extent to which restrictions on the capabilities of faulty processes reduces the required replica requirements.

We have shown that, contrary to conventional wisdom, enforcing *non-equivocation*—i.e., eliminating the ability for a faulty process to tell multiple different stories—is not sufficient to reduce the required number of processes, though the combination of non-equivocation and digital signatures does suffice. In ongoing work with Aniket Kate at MMCI we are exploring refinements to non-equivocation that (a) effectively reduce the required processes below $3f + 1$ and (b) are implementable with modern hardware.

The group is also working with Rodrigo Rodrigues and the group formerly known as the Dependable Systems Group to identify a failure model that accurately reflects the realities of modern data center hardware. Two key observations in the data center context are (i) while individual machines can and do fail in unexpected ways, the failures are generally not coordinated, and (ii) while the network is technically asynchronous, the vast majority of the time it behaves synchronously. Preliminary results indicate that

using a fault model based off of these observations (*Visigoth fault tolerance*) it is possible to reduce the required number of processes to $f+1$. Additional information on this line of research can be found in the Dependable Systems Group section.

...and Practice. In addition to the foundational work discussed above, the group is engaged in three distinct design and implementation efforts related to (Byzantine) fault tolerant state machine replication.

1. In collaboration with colleagues at UT-Austin and Grenoble-INP we have shown that it is possible to use state machine replication techniques with multi-threaded servers that process requests in parallel [124]. The key to this result lies in reversing the normal order of operation: traditional state machine replication first agrees on an order of requests and then executes in the specified order; our system instead executes the requests and then agrees that the resulting state and outputs are the same across all replicas. The key to making this system work is efficient fine-grained rollback and state transfer.

2. The substantial body of literature on state machine replication demonstrates that it is possible to use replication to make a single service robust to failures. When state machine replication is applied to modern systems that compose multiple services to process a single user request, unanticipated anomalies occur. Consider, for example, a standard MapReduce cluster: the user of a MapReduce cluster interacts with the job coordinator (service 1) which in turn interacts with the HDFS NameNode (service 2). When both services are replicated, a single user request to service 1 can induce multiple physical requests to service 2. In collaboration with colleagues at UT-Austin, the group is working to design and implement efficient techniques for chaining replicated state machines together.

3. In collaboration with the Dependable Systems Group and Rodrigo Rodrigues we are implementing a Visigoth fault tolerant replication protocol.

13.2.2 Concurrency and geo-replication

Geo-replicated systems that require immediate responses to user requests must frequently respond without waiting for cross-site coordination to occur. The resulting concurrent execution of contending requests at independent sites can, if not handled carefully, lead to unexpected application behaviors. Systems that are robust to this usage pattern are said to be *eventually consistent*.

The group, in collaboration with the Dependable Systems Group, has shown that for applications that require state convergence (i.e., any pair of replicas that have processed the same set of requests are in the same state), eventually consistent operation as described above is only possible for operations that (a) commute and (b) are incapable of causing a system invariant to be violated. We successfully implemented a prototype system that leverages the distinction between the *original execution* and the *shadow execution* (or induced side effects) of an operation in order to increase the number of commutative operations in a system. The team has expanded to include members of Viktor Vafeiadis and the Software Analysis and Verification group in a new effort to automatically produce shadow operations and identify at run-time which of the produced shadow operations require coordination.

14 The Social Information Systems Group

14.1 Overview

The report covers the period October 2013 – December 2014. (In January 2015 Cristian Danescu-Niculescu-Mizil joined Cornell University as an assistent professor in the department of information science.) The group's research interests are in social computing and natural language processing. At a high level, the group aims at developing computational frameworks that can lead to a better understanding of human social behavior and shape the future of social information systems.

Personnel.

The group is led by Cristian Danescu-Niculescu-Mizil. Vlad Niculae started as a PhD student in January 2014. Justine Zhang was a summer intern from June 2014 to August 2014, visiting from Stanford University.

Collaborations.

The group has an ongoing collaboration with the networked systems group (led by Krishna Gummadi). Externally, the group was collaborating with researchers at Cornell University, Google, Harvard University, Max Planck Institute for Informatics, Stanford University.

Publications.

The following papers were co-authored under the MPI-SWS affiliation and published at prime natural language processing and social computing venues:

- "How to Ask for a Favor: A Case Study on the Success of Altruistic Requests", ICWSM 2014 [12]
- "How Community Feedback Shapes User Behavior", ICWSM 2014 [63]
- "People on Drugs: Credibility of User Statements in Health Communities.", KDD 2014 [170]
- "Brighter than Gold: Figurative Language in User Generated Comparisons", EMNLP 2014 [152]

Data and code.

In order to encourage further research in these research areas, we released three datasets and one natural language processing API:

1. A collection of comparisons annotated for figurativeness (available at `http://vene.ro/figurative-comparisons/`)
2. A dataset of textual requests together with their outcome and meta-data (available at `http://cs.stanford.edu/~althoff/raop-dataset/`)
3. Large scale conversational data from Wikipedia talk-pages (available at `http://www.mpi-sws.org/~cristian/Echoes_of_power.html`)
4. Released the Stanford Politeness API and the Politeness Web App (`politeness.mpi-sws.org`)

Press.

During the period covered in this report, Cristian's research on language and social computing has been featured in popular-media outlets such as the New York Times, The Guardian, Huffington Post, Lifehacker, Gizmodo, Business Insider and Sddeutsche Zeitung.

Awards and invited talks.

Cristian gave invited talks at Google, Facebook, Cornell University, ETH Zürich, University of Zürich, at the Computational Linguistics in Political Science Conference, gave an invited keynote at the Interaction and Exchange in Social Media Workshop and was an invited panelist at the ACL Workshop on Language and Social Dynamics

Service.

During the last year Cristian co-organized the Workshop on Language Technologies and Computational Social Science at ACL 2014, served in the program committee of WSDM 2015, WWW 2014, EACL 2014 and reviewed for the Transactions of the Association for Computational Linguistics.

14.2 Research agenda

More and more of life is now manifested online, and many of the digital traces that are left by human activity are in natural-language format. This is a time when the exploitation of these resources under a computational framework can bring a phase transition in our understanding of human social

behavior. Our group's research takes advantage of this opportunity and aims at discovering, understanding and modeling complex human social behavior starting from very large-scale textual data.

Within this paradigm we are currently pursuing two specific goals. In one of them, we are investigating how key aspects of social relations between individuals are embedded in (and can be inferred from) their conversational behavior. In the other, we are exploring and leveraging social and textual factors that affect how users perceive the "usefulness" of online content and of their fellow community members.

A. Conversational Behavior and Social Relations

With the arrival of detailed data on the social interactions within online communities, an active line of research has attempted to uncover the rules that govern these interactions. To date, these analyses have mainly used structural features of the interactions, including who talks to whom, how frequently, and how these patterns of interaction form larger network structures. But the interactions themselves are generally taking place in natural language — both spoken and written — and the language content of these interactions has been a long-acknowledged missing ingredient in this style of investigation. The reason for this is clear: while it is reasonable to suppose that signals within the language could provide insight into the social structure of the community, it has been challenging to extract useful language-level signals that are domain independent.

Success of altruistic requests Requests are at the core of many social media systems such as question and answer sites and online philanthropy communities. While the success of such requests is critical to the success of the community, the factors that lead community members to satisfy a request are largely unknown. Success of a request depends on factors like who is asking, how they are asking, when are they asking, and most critically what is being requested, ranging from small favors to substantial monetary donations. Together with Tim Althoff and Dan Jurafsky [12], we present a case study of altruistic requests in an online community where all requests ask for the very same contribution and do not offer anything tangible in return, allowing us to disentangle what is requested from textual and social factors. Drawing from social psychology literature, we extract high-level social features from text that operationalize social relations between recipient and donor and demonstrate that these extracted relations are predictive of success. More specifically, we find that clearly communicating need through

the narrative is essential and that that linguistic indications of gratitude, evidentiality, and generalized reciprocity, as well as high status of the asker further increase the likelihood of success. Building on this understanding, we develop a model that can predict the success of unseen requests, significantly improving over several baselines. We link these findings to research in psychology on helping behavior, providing a basis for further analysis of success in social media systems.

Linguistic style coordination in online communities My approach is rooted in the psycholinguistic theory of *linguistic style coordination* [51, 153, 159, inter alia], which accounts for the general observation that participants in conversations tend to immediately and unconsciously adapt to each other's language styles – to the extent that a speaker will even adjust the number of articles and other function words in their next utterance in response to the number in their partner's immediately preceding utterance.

This fascinating phenomenon was previously observed and studied almost exclusively in small-scale or controlled laboratory studies. A priori, it was not all clear whether linguistic coordination would occur under the constraints imposed by the online setting, where most conversations are not face-to-face, do not happened in real-time and are subject to various formatting limitations. Furthermore, there was no formalism that allowed the quantification of this phenomenon at large. In collaboration with with Michael Gamon and Susan Dumais [71] I changed the status quo by proposing a probabilistic framework that can model the coordination phenomenon and measure its effects in a large scale, "in the wild" setting. By applying this framework to a large Twitter conversational dataset specifically developed for this task (comprising 210,000 conversations, this was arguably the largest complete conversational dataset to date), we showed for the first time that linguistic style coordination is prevalent in online conversations. Moreover, the experiment provided new insights into the phenomenon which suggest that linguistic coordination could be used as a signal for uncovering key factors of the social structure of communities.

Echoes of power And indeed, in recent work with Lillian Lee, Bo Pang and Jon Kleinberg [73], I show that in group discussions power differentials between participants are revealed by how much one individual immediately echoes the linguistic style of the person they are responding to. Starting from this observation, we propose an analysis framework based on linguistic coordination that can be used to shed light on power relationships and

that works consistently across multiple types of power — including a more "static" form of power based on status differences, and a more "situational" form of power in which one individual experiences a type of dependence on another. Using this framework, we show how conversational behavior can be successfully employed to reveal power relationships in two very different settings: discussions among Wikipedians and arguments before the U.S. Supreme Court.

Computational politeness Politeness is another conversational phenomenon intimately related to the power dynamics of social interactions [57, 134]. In recent work with Moritz Sudhof, Dan Jurafsky, Jure Leskovec and Christopher Potts ([74]), I developed the first computational framework for characterizing and identifying politeness strategies in requests. By applying this framework to large-scale conversational data (comprising 400,000 requests from Wikipedia editors and from users of question-answering forums) we reveal new interactions between politeness and social power: in what seems like a "power-corrupts" type of situation, we find that polite users are more likely to achieve high status in their community, but, once elevated, they become less polite.

B. Analyzing the Perception of "Usefulness"

As the web evolves towards being more and more user-centric, instances in which users evaluate the "usefulness" of online content become increasingly common. There are two basic facets of online "usefulness" evaluation. In one case it occurs *explicitly*, namely, when users are asked to express whether they find some content "useful" or not: on review sites like Amazon.com, each review is accompanied by a question like "Was this review helpful to you?"; on community Q&A sites like Askville and Yahoo! Answers users have the option to rate the available answers; and social news sites like Digg.com and Reddit.com are based on this kind of feedback mechanism. The other facet of online "usefulness" evaluation occurs when users *implicitly* express their "usefulness" judgments through their actions: for recommender systems like Netflix, ordering a product or visiting its page is considered to be an indication that the user might find that and similar products interesting; in the case of search engines, clicks are often taken as implicit relevance feedback indicating whether the user thinks that the offered search result or ad is more "useful" relative to other results in the context of their information or commercial need.

My intention is to explore the social, contextual and textual factors that influence the way in which users perceive the "usefulness" of online content. The implications of such a study are two-fold: from a practical perspective it would allow us to build systems that can better accommodate the individuality of each user or that can optimize collective satisfaction; from a social psychological point of view it would give us insight into the mechanisms behind the perception of "usefulness".

Social and linguitic signals of credibility. Online health communities are a valuable source of information for patients and physicians. However, such user-generated resources are often plagued by inaccuracies and misinformation. Together with Subhabrata Mukherjee, Gerhard Weikum [170] we propose a method for automatically establishing the credibility of user-generated medical statements and the trustworthiness of their authors by exploiting linguistic cues and distant supervision from expert sources. To this end we introduce a probabilistic graphical model that jointly learns user trustworthiness, statement credibility, and language objectivity. We apply this methodology to the task of extracting rare or unknown side-effects of medical drugs - this being one of the problems where large scale non-expert data has the potential to complement expert medical knowledge. We show that our method can reliably extract side-effects and filter out false statements, while identifying trustworthy users that are likely to contribute valuable medical information.

Social mechanisms underlying helpfulness evaluation of opinions. In joint work with G. Kossinets, J. Kleinberg and L. Lee [72], I am the first to develop a framework for understanding and modeling how opinions are evaluated with respect to helpfulness within on-line communities. The problem is related to the lines of computer-science research on opinion, sentiment, and subjective content [155], but with a crucial twist in its formulation that makes it fundamentally distinct from that body of work. Rather than asking questions of the form "What did Y think of X?", we are asking, "What did Z think of Y's opinion of X?" Crucially, there are now three entities in the process rather than two. Such three-level concerns are widespread in everyday life, and integral to any study of opinion dynamics in a community. For example, political polls will more typically ask, "How do you feel about Barack Obama's position on taxes?" than "How do you feel about taxes?" or "What is Barack Obama's position on taxes?" (though all of these are useful questions in different contexts). Also, Heider's theory of

structural balance in social psychology [114] seeks to understand subjective relationships by considering sets of three entities at a time as the basic unit of analysis. But there has been relatively little investigation of how these three-way effects shape the dynamics of on-line interaction, and this is the topic we considered in our work.

We formulated and assessed a set of theories that govern the evaluation of opinions, and applied these to a dataset consisting of over four million reviews (arguably the most comprehensive review dataset studied to date) of roughly 675,000 books on Amazon.com's site. The resulting analysis provided a way to distinguish among competing hypotheses for the social feedback mechanisms at work in the evaluation of Amazon reviews: we offered evidence against certain of these mechanisms, and showed how a simple model can directly account for a relatively complex dependence of helpfulness on review and group characteristics. We also used a novel experimental methodology that takes advantage of the phenomenon of review "plagiarism" to control for the text content of the reviews, enabling us to focus exclusively on factors outside the text that affect helpfulness evaluation.

Antisocial behavior in online communities. When users evaluate content contributed by fellow users (e.g., by liking a post or voting on a comment), these evaluations create complex social feedback effects. In work with Justin Cheng and Jure Leskovec [63] we investigate how ratings on a piece of content affect its author's future behavior. By studying four large comment-based news communities, we find that negative feedback leads to significant behavioral changes that are detrimental to the community. Not only do authors of negatively-evaluated content contribute more, but also their future posts are of lower quality, and are perceived by the community as such. Moreover, these authors are more likely to subsequently evaluate their fellow users negatively, percolating these effects through the community. In contrast, positive feedback does not carry similar effects, and neither encourages rewarded authors to write more, nor improves the quality of their posts. Interestingly, the authors that receive no feedback are most likely to leave a community. Furthermore, a structural analysis of the voter network reveals that evaluations polarize the community the most when positive and negative votes are equally split.

Conclusion

Our groups's research aims at developing computational frameworks that can transform our understanding of human social behavior by unlocking the unprecedented potential of the large amounts of natural language data generated online. So far this enterprise provided key insights into diverse aspects of human conduct — such as the social patterns governing conversational behavior and the mechanisms behind the evaluation of "usefulness"— and the tools necessary to convert these insights into practical applications that can enhance our online experience.

Part IV
Details

15 Details

In this section, we provide detailed information about the institute, following the outline required by the Max Planck Society's rules for scientific advisory board status reports.

15.1 Structure and organization

Faculty As discussed in Section 1.1, the institute has a flat organization, with currently ten independent research groups, each led by a faculty member (tenure-track, tenured, or director). Max Planck Fellow Michael Backes leads an additional research group (fellow groups are limited to a maximum of two terms of five years each; Backes is in his second term). In addition, Robert Harper (CMU) has an appointment as an external scientific member, and Rodrigo Rodrigues is an adjunct faculty member.

The faculty appointment dates, tenure status and (for tenured faculty) retirement dates are shown in Figure 1.

Leadership As stated in the institute bylaws, institute policy is decided jointly by the faculty. The faculty typically meets weekly, with the location alternating between the two sites. The day-to-day operation of the institute is in the hands of the Managing Director (currently Peter Druschel), assisted by the head of the administrative department, Volker Maria Geiss. The position of Managing Director rotates among the directors (normally every two years).

Administrative support The MPI-SWS and the MPI for Informatics (MPI-INF) in Saarbrücken are supported by a shared administrative department headed by Volker Maria Geiss. The department provides personnel, finance, and purchasing services. Volker also handles much of the public relations, relations with local governments, and relations with other research institutions in Kaiserslautern and Saarbrücken. We share the core IT support group with MPI for Informatics. Separately MPI-SWS has its own user-facing IT support group. Both the core IT and user-facing IT groups report to Volker as well.

MPI-SWS shares a library jointly with MPI-INF, DFKI (German Research Center for Artificial Intelligence), and the Mathematics and Computer Science departments of Saarland University. The joint library reports to Volker Maria Geiss.

Group Name	Group Leader	Start	Status	Retire
Real-Time Systems	Brandenburg	2011	tenure-track	—
Robust Systems	Clement	2012	left in 2014, tenure-track	—
Social Information Systems	Danescu-Niculescu-Mizil	2013	left in 2015, tenure-track	—
Foundations of Programming	Dreyer	2008	tenured in 2013	2047
Distributed Systems	Druschel	2005	director	2025
Large Scale Internet Systems	Francis	2009	director	2023
Foundations of Computer Security	Garg	2011	tenure-track	—
Networked Systems	Gummadi	2005	tenured in 2012	2046
Rigorous Software Engineering	Majumdar	2010	director	2042
Software Analysis and Verification	Vafeiadis	2010	tenure-track	—
Information Security and Cryptography	Backes	2008	fellow	—
Dependable Systems	Rodrigues	2008	left in 2012, now adjunct	—
Networks and Machine Learning	Gomez Rodriguez	2014	tenure-track	—
Automated Verification and Approximation	Darulova	Fall 2015	tenure-track	—

Figure 1: MPI-SWS research groups

Administrative assistance for faculty, staff, postdocs and students is provided by an administrative team consisting of four members—Brigitta Hansen and Claudia Richter in Saarbrücken, Vera Schreiber and Susanne Girard (currently on maternity leave, temporarily replaced by Roslyn Stricker) in Kaiserslautern. In addition, Maria-Louise Albrecht serves as coordinator for the MPI-SWS graduate program.

IT services Support for core information technology services (network and core network services, telephony, and storage/email/web services) is provided by a team headed by Jörg Herrmann. This team (currently 10

members) is also shared with the MPI for Informatics. Working together with the core team is a four-member IT support team (headed by Christian Mickler), which provides dedicated support for the all other IT needs of SWS researchers, such as audio/video conferencing, hardware, and software issues. Locating this dedicated team alongside the offices of SWS researchers (both in Kaiserslautern and Saarbrücken) has made it much easier for them to respond effectively to researchers' often-spontaneous requests for assistance.

English language support It is critically important that young researchers develop their communication skills. Moreover, we feel that English language support is particularly important for non-native English speakers. Therefore, the institute has a strict policy of using English as the working language. We feel this is necessary, not only to accommodate our highly international staff, but also to help the non-native English speakers develop their language skills.

The institute employs an English support coordinator who provides English language speaking, writing and presentation support for all institute members. Rose Hoberman, who currently occupies the position, has a Ph.D. in computer science from CMU. She offers regular courses on presentation, reading, and writing skills, and additional soft skills courses as needed. She also provides feedback on institute members' presentations, papers, and other documents. We plan to hire additional staff as the institute grows.

Research support team The institute also has several funded positions available for software developers. We have been filling these positions on a temporary per-project basis. In this reporting period, we have used six such developers, William Caldwell (network anonymity), Cedric Gilbert (email attachment malware), Sebastian Probst-Eide, Matthias Kretschmer, Cristian Berneanu, and Sasa Juric (anonymized analytics). We have also used an additional developer, Jeff Hoye, on a consulting basis for both research projects and institute administrative tools such as our admissions system.

15.2 Research program and groups

This information is provided in previous sections.

15.3 Personnel structure

Currently, the institute has 96 members (excluding interns and visitors). Among these, there are 64 researchers and 32 non-research staff. Of these,

25 are administrative staff shared with MPI-INF, and 7 are IT staff. During the reporting period, 34 members joined MPI-SWS, and 33 left.

Permanent faculty	5
Tenure-track faculty	5
Permanent staff	32
Temporary contracts (incl. 1 Fellow and 6 undergrad. assistants)	12
Postdocs	11
PhD students	39

15.4 Structure of the budget

The institute's total yearly budget is EUR 10.51M per year. Of that, the institute's yearly expenditure amounted to EUR 7.53 per year, including EUR 2.7M for material expenses, EUR 189K special financing for basic scientific equipment, EUR 835K for investment in major equipment, EUR 3.19M for personnel expenses (excluding stipends) and EUR 617K for graduate and postdoctoral stipends and contracts. (Personnel funds can be used to fund additional stipends but not vice versa.)

The institute is alloted 5 senior faculty (director, W3) positions, and up to 12 junior and mid-career (tenure-track or tenured, W2) positions. Currently, the institute uses only part of its full budget, since only 10 faculty positions have been filled (including 3 director positions).

15.5 Provision of material, equipment, and working space

Material The nature of the institute's research in software systems is such that it does not require materials beyond normal office supplies.

Equipment The institute has a state-of-the art, reliable and fail-safe computing infrastructure. A redundant network backbone of 10 Gigabit links connects the Kaiserslautern site, the Saarbrücken site, the MPI-INF and Saarland University via a multi-gigabit link to the X-WIN—the German research network. Basic network services, as well as email and web servers, are implemented in a reliable and fail-safe manner. Storage services provide backup and access to about 90TB of storage. All services are monitored by a system that notifies the IT staff via SMS and e-mail in case of trouble. Institute members have personal desktop and notebook computers.

The institute currently maintains three clusters for research. The oldest cluster has ninety Xeon hexa-core systems with each two CPUs, the second

cluster is a Blade-Server and has thirty-two Xeon octa-core systems with each two CPUs and the third cluster with its quad twelve-core CPUs are equipped with 1.5 TB RAM. All clusters are connected to the institute's intranet and have direct access to the storage services. The institute also contributes six nodes to the PlanetLab testbed and 70 nodes to the VICCI testbed.

The computing infrastructure will be expanded as needed to accommodate new research demands and growth. For instance, future faculty hires may require more specialized laboratories.

Space At our current level of staffing, the two buildings in Saarbrücken and Kaiserslautern provide ample office space, lab space, meeting and conference rooms, open space, event space, and machine room space. Each site is able to accommodate all members of the other site on our weekly visit days. The lecture halls and meeting rooms of both sites are periodically used by external organizations. The Kaiserslautern building is also used by a group from the TU-KL math department.

15.6 Junior scientists and guest scientists

Junior scientists Attracting, supporting, mentoring and creating opportunities for outstanding young researchers is a top priority at the institute.

The purpose of the institute's tenure-track systems is to attract the very best young PhDs internationally and provide them with conditions (independence, resources, mentorship, full participation in the institute governance) that will allow them to grow as researchers and future leaders. The institute has a formal faculty mentorship program.

We have an active program to attract and support outstanding postdoctoral researchers from diverse backgrounds. Postdoctoral positions are normally granted for two years, and can be extended to three years. Currently, we have 11 postdocs from nine countries. A list of our current postdocs can be found online at `http://www.mpi-sws.org/index.php?n=people&s=function&c=postdocs`.

A high priority for the institute is to attract the best graduate students and provide them with the training necessary for them to obtain academic and research positions at the world's best universities and research labs. We seek to maintain a highly talented, highly motivated and diverse body of graduate students. Moreover, we provide intensive training in small groups (less than six students per faculty). We emphasize high-risk, high-impact research and publication in top venues.

We currently have 39 doctoral students from 16 countries. A list of our current doctoral students can be found online at `http://www.mpi-sws.org/index.php?n=people&s=function&c=doctoral`.

Guest researchers As part of the institute's strategy to increase visibility, create opportunities for collaborations with other institutions, and contribute to a vibrant intellectual environment, the institute has a very active program for short and longer term visitors at all seniority levels.

We host both undergraduate and graduate interns at the institute. During the reporting period, MPI-SWS hosted 7 undergraduate interns and 40 graduate interns from 15 countries.

Researchers from other institutions frequently come for research visits. There were around 100 such short-term visitors, including: Azalea Raad (Imperial College London), Michael W. Hicks (University of Maryland), KyoungSoo Park (KAIST), Robert West (Stanford), Marcos K. Aguilera (unaffiliated, formerly MSR SV), Oksana Denysyuk (INES-ID and University of Lisbon), Rachid Guerraoui (EPFL), Anne-Marie Kermarrec (Inria), Christian Rossow (MMCI), Ingmar Weber (MPII), Nico Pfeifer (MPII), Robert van Renesse (Cornell University), Marcel Schulz (MMCI), Mike Dodds (University of York), Flavio Chierichetti (Sapienza University of Rome), Andrei Sabelfeld (Chalmers), Hongjin Liang (UST China), Andreas Karrenbauer (MPII), Jilles Vreeken (MMCI), Ras Bodik (UC Berkeley), Kavita Bala (Cornell University), Andreas Wiese (MPII), Lana Hadarean (NYU), Ori Lahav (Tel Aviv University), Michael Roitzsch (TU Dresden), Aaron Carroll (NICTA Sydney), Andrew Myers (Cornell University), Noah Smith (Carnegie Mellon University), Koushik Sen (UC Berkeley), Steve Zdancewic (University of Pennsylvania), Simon Peter (Seattle University), Jim Anderson (UCNC at Chapel Hill), Ankit Singla (UIUC), Mark Crovella (Boston University), Alin Deutsch (UCSD), Amir Houmansadr (University of Texas at Austin), Tim Kraska (Brown University), Andrei Sabelfeld (Chalmers), Gilles Barthe (IMDEA), Limin Jia (CMU), Sanjit Chatterjee (Indian Institute of Science), Sanjiva Prasad (IIT-Delhi), Deian Stefan (Stanford), Devdatta Akhawe (UC Berkeley), Aniket Kate (MMCI), Michael Carl Tschantz (ICSI), Andrew Tomkins (Google), Evimaria Terzi (Boston University), Chengxiang Zhai (UIUC), Mark Crovella (Boston University), Javier Esparza (TU Munich), Pierre Ganty (IMDEA Madrid), Stefan Göller (ENS Cachan), Eva Darulova (EPFL), Aditya Kanade (IISc Bangalore), Anthony Lin (NUS Singapore), Joel Ouaknine (Oxford), Majid Zamani (TU Munich), Koushik Sen (UC Berkeley), Thomas Wies (NYU), Jyo Deshmukh

(Toyota), James Kapinski (Toyota), Jeffrey Fischer (Quaddra Software), Matthias Woehrle (Bosch), Indranil Saha (Postdoc, UC Berkeley), Milos Gligoric (PhD student, UIUC), Liana Hadarean (PhD student, NYU), Jean Yang (PhD student, MIT), Roopsha Samanta (Postdoc at IST), Verena Wolf (MPII), Alexey Gotsman (IMDEA), Cezara Dragoi (IST Austria), Michael Pradel (TU Darmstadt), Bryan Ford (Yale), I-Ting Angelina Lee (CSAIL), Julian McAuley (Stanford), Anja Feldmann (TU Berlin), Hongseok Yang (Oxford), Viktor Kuncak (EPFL), Foteini Baldimtsi (Brown), Mark Silberstein (Israel Institute of Technology), Jürgen Steimle (MMCI), Sanjoy Baruah (UNC at Chapel Hill), Christian Grothoff (TU Munich), Hans Boehm (Google), Alessandro Panconesi (Informatica-Sapienca, Madrid), Martin Jankowiak (NYU), Bobby Bhattacharjee (University of Maryland, College Park), Lorenzo Alvisi (UT Austin), KyoungSoo Park (KAIST), Nicole Megov (TU Berlin), Robbert van Renesse (Cornell), Arvind and Hema Krishnamurthy (UW), Aaron Carroll (NICTA), Michael Roitzsch (TU Dresden), Christian Grothoff (INRIA), and Elaine Shi (University of Maryland, College Park).

We have also had two long-term faculty visitors, Amir Herzberg from Bar-Ilan University who spent his sabbatical here in 2015 and Jonathan Katz from University of Maryland who is spending one year as part of his Humboldt Award.

15.7 Equal opportunity

Ensuring gender diversity is a well-known perennial problem in computer science departments worldwide. At the beginning of the review period, 13.8% of the scientific staff (students, postdocs, and faculty) were women. During the review period, 19% of the joining scientific staff were women (1 of 16 students, 3 of 8 postdocs, and 1 of 2 faculty). The current number is 20.7% (6 of 37 doctoral students, 5 of 11 postdocs, and 1 of 10 faculty).

15.8 Relations with domestic and foreign research institutions

Local We continue to work towards integration with UdS and TU-KL. Professors from the two departments are present in our faculty recruitment committees and our graduate student admission committees. MPI-SWS is a member of the UdS CS graduate school. There are a number of joint research projects with UdS and TU-KL faculty, including Björn Brandenburg with Gerhard Fohler, Deepak Garg with both Christian Hammer and

Aniket Kate, and Peter Druschel with both Matteo Maffei and Aniket Kate. Druschel is a PI in CISPA, the Excellence Cluster at UdS, and the privacy SFB proposal. Druschel, Majumdar, and Backes are co-PIs on the imPACT ERC Synergy Grant.

MPI-SWS faculty have taught courses in their areas of expertise. During this reporting period, faculty at MPI-SWS taught the following courses:

- Operating System Design and Implementation, TU Kaiserslautern, Winter, 2013/2014
- Cyber-Physical Systems, TU Kaiserslautern, 2014/2015
- Categorical Logic, Saarland University, 2014/2015
- Core course on Distributed Systems, Saarland University and TU Kaiserslautern, 2014/2015
- Graduate course on Logics in Security, Saarland University, 2014/2015
- Language-based Security, Saarland University, 2015.
- A graduate seminar on Social computing systems, 2013.
- Graduate course on Reactive Systems Verification, TU Kaiserslautern, 2014.
- Concurrency Theory, TU Kaiserslautern, 2014.

International Institute members maintain numerous collaborations with researchers at international universities and research institutions, including:

Universities Stanford, Georgia Tech., IIT Kharagpur, MPI-IS, Harvard University, Oxford University, University of Cambridge, Northeastern University, University of York, University of Modena, UNC Chapel Hill, Columbia University, Colorado State University, Warwick University, IISc India, ISI Calcutta, CMU, Purdue University, Chalmers, University of Dundee, University of Texas at Austin, ENS Cachan, UCLA, IISc Bangalore, EPFL, LABRI Bordeaux, ETH Zurich, UIUC, Louisiana State University, University of Waterloo, University of Pennsylvania, NYU, University of Leuven, Tel Aviv University, Seoul National University, Teesside University, Cornell University, Maryland University, Nova University of Lisbon, Aarhus University, University of Glasgow, Yale University, Duke University, IST Lisbon, TU

Braunschweig, Humboldt University Berlin, Ludwig-Maximilians-University Munich (Faculty of Medicine, Eye Clinic), TU Munich, Uni Freiburg, University of Bonn, Friedrich-Alexander-University Erlangen-Nürnberg, Universitätsklinikum Heidelberg, and the University of Darmstadt.

Research Institutes INRIA Grenoble, Telefonica Research, MMCI, ICSI, IMDEA Software Institute Madrid, IST Austria, INRIA in Nancy, and SnT Luxembourg.

Industry Google, Cisco Berlin Innovation Center, Microsoft Research, Toyota, Bosch, Praemandatum, SRT, Center for Genomics and Transcriptomics GmbH, Twitter, Linux, Mozilla Research, Facebook, and InnoZ (Berlin Innovation Center for Mobility and Societal Change).

The institute is a partner in the MPS's collaboration agreement with the Indo-German Max Planck Center in Computer Science (IMPECS).

15.9 Activities regarding the transfer of knowledge/relations with industry

Gummadi's group's work in retrieving trustworthy and relevant information is deployed as part of Twitter search. To date, over 400 research groups at universities and research labs worldwide have used the group's data sets as part of measurement studies of online social networks. Over the last two years, more than 400,000 end users worldwide have used the Glasnost software designed by the group to test the traffic management policies of their access ISPs (e.g., cable and DSL providers). His group has a project with Telefonica Research to launch the Data Transparency Lab.

Majumdar's group continues to work closely with Toyota on a Conformance testing tool for Simulink models used in internal testing at Toyota.

Backes' group has been working on the smartphone application Appguard, currently being marketed by SRT.

Pedro Fonseco's code testing tool SKI was used by Linux kernel developers.

As mentioned elsewhere, Francis' group uses the startup Aircloak as a critical part of its research agenda. Three patent applications are currently outstanding.

Druschel's and Garg's groups obtained a Google research award on the Thoth project. They have applied for a patent on Storage Leases (awarded in the US, pending in Europe and China).

15.10 Symposia, conferences, etc.

The institute organized the fourth institute retreat in August 2014 at Marienburg, Zell an der Mosel. The primary purpose of this retreat was to make all the research groups at MPI-SWS aware of one another's ongoing work. During the retreat, faculty and students presented and discussed their current work. Faculty also used the opportunity to gather feedback and present the future goals and vision of the institute. Other activities included work-in-progress presentations, discussions devoted to academic issues and institute life, talks and discussions on the nature and methodologies of CS research, birds-of-a-feather sessions, and discussions to help students make the most of their graduate studies and prepare for future roles as leading researchers and faculty.

The institute has an ongoing distinguished lecture series. The purpose of this series is to bring senior leaders in software systems to the institute (typically, for two days), have them give a talk, showcase the institute, have them meet faculty, postdocs and students, and last but not least, seek feedback on our strategy and advice in identifying potential hires. In this reporting period, we have had 19 distinguished lecturers: Gilles Barthe (IMDEA Madrid), Marcos K. Aguilera (unaffiliated), Rachid Guerraoui (EPFL), Anna-Marie Kermarrec (Inria), Noah Smith (Carnegie Mellon University), Joel Ouaknine (University of Oxford), Robbert van Renesse (Cornell University), Andrei Sabelfeld (Chalmers University), Ras Bodik (UC Berkeley), Kavita Bala (Cornell University), Andrew Myers (Cornell University), Anja Feldmann (Telekom Innovation Laboratories, TU Berlin), Hongseok Yang (University of Oxford), Viktor Kuncak (EPFL), Sanjoy Baruah (University of North Carolina at Chapel Hill), Andrew Tomkins (Google), Hans Boehm (Google), Alessandro Panconesi (Informatica-Sapienza, University of Rome) and Alin Deutsch (UC San Diego). The full list of lecture abstracts and titles is available online at `http://www.mpi-sws.org/index_flash.php?n=lectures/dlseries/program`.

This series has been very effective in raising the institute's visibility and identifying potential hires, and we have received valuable feedback and advice regarding our strategy.

A number of the faculty have given keynote and invited talks at various institutions and conference during the reporting period. These are detailed within the individual sections.

15.11 Committee work of the faculty

MPI-SWS researchers have served on the program committees of over 75 conferences and workshops, and have chaired or co-chaired the PCs of 9 conferences and 7 workshops. The information is provided in detail in the individual research group sections.

Brandenburg and Druschel served as reviewers for the German funding agency DFG.

Dreyer was elected and served as "awards czar" on the ACM SIGPLAN Executive Committee from July 2012 to June 2015. This entailed overseeing the management of all SIGPLAN awards, and chairing the committees for test-of-time awards for the four major SIGPLAN conferences. In 2014 he was also was invited to join the editorial board of the Journal of Functional Programming, as well as the IFIP Working Group 2.8 on Functional Programming.

Druschel is a co-PI in the successful renewal of both Saarland University's MMCI Cluster of Excellence and the Saarbrcken Graduate School in Computer Science, funded by the German National Science Foundation (DFG). He is also a co-PI and assistant director of the Center for Information Security, Privacy and Trust, funded by the German ministry of science (BMBF).

Within the MPS, Druschel continues to serve on the strategy committee (Perspektivenkommission) of the Chemical, Physical, and Technology section, and the selection panel of the joint Fraunhofer/Max Planck research program. He served on the Committee on Information Technology (BAR) through June 2014. Currently, Druschel also serves on two presidential committees of the MPS: The committee on the Support of Junior Scientists, and the committee on IT Security. Lastly, Druschel co-organized a Symposium on Foundations of Cyper-Security and Privacy for the CPT Section of the MPS in July 2015.

In addition, he continues to serve on the Technical Advisory Boards of Microsoft Research, Cambridge and Microsoft Research, India, the scientific committee of the Laboratory on Information, Networking and Communication Sciences (LINCS), Paris, and on the steering committee of the EuroSys/INRIA Winter School on Hot Topics in Distributed Systems (HTDC). He served on the steering committee of the ACM SIGOPS Asia-Pacific Workshop on Systems (APSys) through 2014. He also works on the editorial boards of the ACM Communications of the ACM (CACM) and the Royal Society Open Science Journal.

Peter Druschel, Rupak Majumdar, Michael Backes, and Gerhard Weikum

(MPI-INF) are co-PIs on an ERC Synergy Grant on Privacy, Accountability, Compliance, and Trust in the Internet (EUR9.25M, 2015–2021).

Gummadi currently serves as an associate editor for Transactions on the Web, a steering committee member of Measurement Lab and as the chair of the technical advisory board of the ACM conference on online social networks (COSN).

Vaveiadis served as publicity chair for POPL 2014 and POPL 2015.

Garg has served as the Publications Chair of CSF every year since 2012.

Danescu-Niculescu-Mizil co-organized the 2014 Workshop on Language Technologies and Computational Social Science at ACL 2014.

Backes is the chair of the Steering Committee of EuroS&P, and he is serving as the PC Chair of the conference's first edition in 2016. He currently serves on the steering committees of IEEE S&P, IEEE EuroS&P, ACM CCS, IEEE CSF and ESORICS. He serves on the Editorial Board of the Journal on Foundations and Trends in Security and Privacy.

15.12 Publications

All publications are listed in the per-group sections. Here, we provide summary information.

During the reporting period, the institute produced 125 peer-reviewed conference, workshop, and journal publications: [160, 90, 91, 15, 187, 92, 110, 68, 157, 78, 119, 69, 180, 48, 154, 47, 135, 10, 136, 50, 23, 167, 179, 83, 11, 156, 43, 94, 42, 139, 138, 44, 61, 177, 172, 122, 176, 133, 120, 140, 161, 184, 149, 98, 130, 127, 45, 97, 128, 99, 183, 166, 16, 25, 33, 30, 39, 22, 36, 19, 24, 38, 17, 21, 20, 37, 18, 27, 34, 26, 28, 32, 29, 31, 35, 41, 56, 59, 168, 185, 186, 60, 55, 113, 54, 125, 132, 121, 163, 62, 143, 165, 171, 190, 77, 80, 79, 89, 101, 146, 142, 65, 96, 126, 85, 58, 86, 144, 145, 115, 116, 87, 66, 82, 64, 13, 12, 63, 170, 152, 75, 74, 158, 40, 14, 103]. Of these, 19 are collaborative across research groups.

15.13 Long-term archiving of research results

MPI-SWS has a policy of keeping all source data used for published research results archived through our normal system backup procedure. When this data is useful for other researchers' work, the data—and, where appropriate, the tools used to produce the data—are also made available on our website.

15.14 Appointments, scientific awards and memberships

- Majumdar wan the POPL Test of Time award in 2014 for his 2004 paper "Abstractions from Proof".
- Gomez Rodriguez won an Outstanding Paper Award at the Neural Information Processing (NIPS) Conference in December 2013.
- Aaron Turon received the 2014 ACM SIGPLAN John C.Reynolds Doctoral Dissertation Award for his thesis.
- Reznichenko received a Best Student Paper Award at CCS 2014.
- Gummadi and Kooti received the Best Paper Award at the AAAI's ICWSM 2012 conference.
- Bhatotia won the Best Student Paper award at Middleware 2014.
- Soudjani won the DIC Best PhD-Thesis Award.
- Sen won the ACM-India Best Doctoral Dissertation Award 2014.
- Valera and Sen were both awarded two-year Humboldt postdoctoral fellowships in March 2015.
- MPI-SWS and Aircloak jointly won the CISCO IoT Security Grand Challenge competition.
- Garg and Druschel were awarded a Google Faculty Research Award in 2014.
- The Aircloak system received a TÜViT Trusted Process certification.
- Druschel, Majumdar, Backes, and Weikum (MPI-INF) received the ERC Synergy Grant in 2013 – one of Europe's most distinguished research awards.
- Dreyer was awarded a 2013 Microsoft Research PhD Scholarship.
- Backes was selected as a member of the German Academy of Science and Engineering (AcaTech) in 2015. He was named one of Germany's "digital minds" by Germany's federal minister of Science and Education Johanna Wanka in 2014. He received an IEEE Outstanding Community Service Award, the CS Teaching Award of Saarland University for the best CS lecture in 2014, and the Teaching Award of the State of Saarland.

15.15 External funding

- Garg, with Prof. Hammer from UdS, obtained funding from the DFG, program RS3, for a period of 4 years.
- Garg and Druschel won a Google Faculty Research Award.
- Gummadi received a gift from AT&T, two Humboldt Postdoctoral Fellowships, two IMPECS fellowships, and a grant for 'Analysis and Design of Online Social Networks' from IMPECS. He continues to be funded by the DFG in Germany and IST in India (2011-2015)
- Druschel, Majumdar, Backes, and Weikum won the ERC Synergy Award 'ImPact: Privacy, Accountability, Compliance, and Trust in Tomorrow's Internet'.
- Majumdar has been the recipient of a Toyota research contract since 2013.
- Vafeiadis will continue until March 2016 to use his European Commission ADVENT Fund.
- Backes received grants from BMBF and CISPA.
- Dreyer obtained a Microsoft Research PhD Scholorship in 2013.
- Francis continued to use the EXIST Forschungstransfer (Technology Transfer) grant for the spinoff Aircloak. MPI-SWS and Aircloak won the Cisco IoT Security Grand Challenge prize.
- Neis continued to be funded from his 2012 Google European Doctoral Fellowship.

15.16 Public relations work

Articles describing Gummadi's social network research have appeared in numerous popular news media and technology blogs including the New York Times, Harvard Business Review, MIT Technology Review, New Scientist, Wired magazine, Slashdot, Businessweek, Sueddeutsche Zeitung (Germany), Science TV (Korea), and MTV (Brazil).

Francis presented at one of the "Impulse aus der Zukunft" open public lectures in Berlin sponsored by Technologie Stiftung Berlin. The event was on the topic "Rethinking online tracking," and included a lecture by Francis

and a panel discussion with Dr. Christoph Peylo, VP of Duetsche Telekom Innovation Lab in Berlin.

Danescu-Niculescu-Mizil's research on language and social computing has been featured in popular-media outlets such as the New York Times, The Guardian, Huffington Post, Lifehacker, Gizmodo, Business Insider and Sddeutsche Zeitung.

Backes's research on providing anonymity guarantees for Tor has been covered by ZDnet, Heise, and Deutschlandfunk. His research on Appguard and its successor Boxify have been featured by the Frankfurter Allgemeine Zeitung, Sueddeutsche Zeitung, c't-Magazin, Technology Review, Hannoversche Allgemeine, and various TV progams such as ORF and NDR Logo.

In addition, the Institute participated in the following public relations activities:

- In Kaiserslautern, efficient collaboration between all relevant academic institutions and strong innovative industrial partners is achieved through the "Science Alliance" umbrella organization.

- An open public Max Planck Forum was held on the topic "When machines get too smart. How secure is our data?" in Saarbrücken in November 2014. A panel discussion was held with Prof. Dr. Michael Backes, Fellow at the Max Planck Institute for Software Systems in Saarbrücken, Prof. Dr. Paul Francis, Director at the Max Planck Institute for Software Systems in Kaiserslautern, Jürgen Lennartz, State Secretary, Head of the State Chancellery and Representative of Saarland to the federal government in Berlin, and Thomas Schauf, from the German Federal Association of Digital Commerce (Bundesverband Digitale Wirtschaft (BVDW) e.V.).

- Together with the TU Kaiserslautern, the Institute participated in the event "Long Night of Sciences", where the university and its affiliated institutes, presented themselves to a wider public audience. This event serves as a link between research and the region, introducing visitors to the entire spectrum of scientific knowledge from basic research to laboratory samples and prototypes.

- High school students interested in computer science gather together every year at the German Federal Computer Science Competition (Bundeswettbewerb Informatik). In September 2013, the TU Kaiserslautern provided the location for the 31st final round, with MPI-SWS as co-host. Of the roughly 1200 participants, 28 advanced to the final round in Kaiserslautern where the winner was determined.

- As part of the Saar Music Festival, the Philharmonic Orchestra Rzeszow performed a concert on 8 June with the theme "Chopin meets Max Planck". The orchestra was accompanied by two soloists, and more than 400 guests at the "Platz der Informatik" on the university campus enjoyed the works by Chopin and other Polish composers. The enthusiastic audience was then invited to a reception at MPI-SWS and a guided tour through the building.

- Publication of articles and book chapters, with the goal of making the Institute well-known to the broader public:

  - The Saarland University publication "Magazine Forschung" reports on the latest scientific findings from the Saarbrcken campus. In an article called "Software systems, the backbone of the wired world", MPI-SWS was presented.
  - The Center for Productivity and Technology Saar (Zentrale fr Produktivität und Technologie Saar e.V.) supports regional companies and organizations with the brochure "Innovatives Saarland". In the most recent 9th edition, all relevant innovation partners are portrayed, including the MPI-SWS. The publication, printed in both German and English, is used by the Saarland State Chancellery to advise partners on results of research and development and cooperation opportunities.

References

[1] Glasnost: Bringing transparency to the Internet.

[2] FriendList Manager: Simplifying FriendList management.

[3] TrulyFollowing: Discover Twitter accounts with suspicious followers.

[4] TrulyTweeting: Detect tampered tweet promotions.

[5] What is happening: Discover real-time topical news on Twitter.

[6] Measure and manage your information diet.

[7] Whom to follow: Discover topic authorities on Twitter.

[8] P. Aditya, B. Bhattacharjee, P. Druschel, V. Erdélyi, and M. Lentz. Brave new world: privacy risks for mobile users. In *Proceedings of the ACM MobiCom Workshop on Security and Privacy in Mobile Environments, SPME@MobiCom 2014, Maui, Hawaii, USA, September 11, 2014*, pages 7–12, 2014.

[9] P. Aditya, B. Bhattacharjee, P. Druschel, V. Erdélyi, and M. Lentz. Brave new world: Privacy risks for mobile users. *Mobile Computing and Communications Review*, 18(3):49–54, 2014.

[10] P. Aditya, V. Erdelyi, M. Lentz, E. Shi, B. Bhattacharjee, and P. Druschel. Encore: Private, context-based communication for mobile social apps. In *The 12th International Conference on Mobile Systems, Applications, and Services (MobiSys'14)*, 2014.

[11] P. Aditya, R. Sen, P. Druschel, B. Bhattacharjee, M. Fritz, and B. Schiele. ipic: Controlling unwanted photo capture. Submitted for publication.

[12] T. Althoff, C. Danescu-Niculescu-Mizil, and D. Jurafsky. How to ask for a favor: A case study on the success of altruistic requests. In *Proceedings of ICWSM*, 2014.

[13] C. Alvin, S. Gulwani, R. Majumdar, and S. Mukhopadhyay. Synthesis of geometry proof problems. In C. E. Brodley and P. Stone, editors, *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, July 27 -31, 2014, Québec City, Québec, Canada.*, pages 245–252. AAAI Press, 2014.

[14] L. Alvisi, A. Clement, A. Epasto, S. Lattanzi, and A. Panconesi. Communities, random walks, and social sybil defense. *Internet Mathematics*, 10(3-4):360–420, 2014.

[15] M. Babaei, P. Grabowicz, I. Valera, and M. Gomez-Rodriguez. Quantifying Information Overload in Social Media and its Impact on Social Contagions. In *Proceedings of the 9th International AAAI Conference on Weblogs and Social Media*, 2015.

[16] M. Backes, M. Barbosa, D. Fiore, and R. M. Reischuk. Adsnark: Nearly-practical privacy-preserving proofs on authenticated data. In *Proceedings of the 36th IEEE Symposium on Security and Privacy (S&P)*, page to appear. IEEE, 2015.

[17] M. Backes, F. Bendun, A. Choudhury, and A. Kate. Asynchronous mpc with a strict honest majority using non-equivocation. In *Proceedings of the 2014 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 10–19. ACM, 2014.

[18] M. Backes, F. Bendun, J. Hoffmann, and N. Marnau. Pricl: Creating a precedent. a framework for reasoning about privacy case law. In *Proceedings of the 4th Conference on Principles of Security and Trust (POST)*, pages 344–363. Springer, 2015.

[19] M. Backes, F. Bendun, M. Maffei, E. Mohammadi, and K. Pecina. Symbolic malleable zero-knowledge proofs. In *Proceedings of the 28th IEEE Computer Security Foundations Symposium (CSF)*, page to appear. IEEE, 2015.

[20] M. Backes, S. Bugiel, and S. Gerling. Scippa: System-centric ipc provenance on android. In *Proceedings of the 30th Annual Computer Security Applications Conference (ACSAC)*, pages 36–45. ACM, 2014.

[21] M. Backes, S. Bugiel, S. Gerling, and P. von Styp-Rekowsky. Android security framework: Extensible multi-layered access control on android. In *Proceedings of the 30th Annual Computer Security Applications Conference (ACSAC)*, pages 46–55. ACM, 2014.

[22] M. Backes, S. Bugiel, C. Hammer, O. Schranz, and P. von Styp-Rekowsky. Boxify: Full-fledged App Sandboxing for Stock Android. In *Proceedings of the 24th USENIX Security Symposium (USENIX)*, page to appear. USENIX, 2015.

[23] M. Backes, J. Clark, P. Druschel, A. Kate, and M. Simeonovski. BackRef: Accountability in Anonymous Communication Networks. In *Proceedings of the 12th International Conference on Applied Cryptography and Network Security (ACNS 2014)*, volume 8479 of *Lecture Notes in Computer Science*, pages 380–400. Springer, 2014.

[24] M. Backes, D. Fiore, and E. Mohammadi. Privacy-preserving accountable computation. In *Proceedings of the 18th European Symposium on Research in Computer Security (ESORICS)*, pages 38–56. Springer, 2013.

[25] M. Backes, D. Fiore, and R. M. Reischuk. Verifiable delegation of computation on outsourced data. In *Proceedings of the 20th ACM Conference on Computer and Communications Security (CCS)*, pages 863–874. ACM, 2013.

[26] M. Backes, M. Gagne, and M. Skoruppa. Using mobile device communication to strengthen e-voting protocols. In *Proceedings of the 12th ACM workshop on Workshop on Privacy in the Electronic Society (WPES)*, pages 237–242. ACM, 2013.

[27] M. Backes, R. Gerling, S. Gerling, S. Nürnberger, D. Schröder, and M. Simkin. Webtrust - a comprehensive authenticity and integrity framework for http. In *Proceedings of the 12th International Conference on Applied Cryptography and Network Security (ACNS)*, pages 401–418. Springer, 2014.

[28] M. Backes, S. Gerling, S. Lorenz, and S. Lukas. X-pire 2.0 - a user-controlled expiration date and copy protection mechanism. In *Proceedings of the 29th ACM Symposium on Applied Computing (SAC)*, pages 1633–1640. ACM, 2014.

[29] M. Backes, N. Grimm, and A. Kate. Lime: Data lineage in the malicious environment. In *Proceedings of 10th International Workshopon Security and Trust Management (STM)*, pages 183–187. Springer, 2014.

[30] M. Backes, T. Holz, B. Kollenda, P. Koppe, S. Nürnberger, and J. Pewny. You can run but you can't read: Preventing disclosure exploits in executable code. In *Proceedings of the 21st ACM conference on Computer and Communications Security (CCS)*, pages 1342–1353. ACM, 2014.

[31] M. Backes, C. Hriţcu, and M. Maffei. Union, intersection, and refinement types and reasoning about type disjointness for secure protocol implementations. *Journal of Computer Security*, 22(2):301–353, 2013.

[32] M. Backes, A. Kate, P. Manoharan, S. Meiser, and E. Mohammadi. Anoa: A framework for analyzing anonymous communication protocols. In *Proceedings of the 6th Workshop on Hot Topics in Privacy Enhancing Technologies (HotPETs)*, 2013.

[33] M. Backes, A. Kate, S. Meiser, and E. Mohammadi. (Nothing else) MATor(s): Monitoring the anonymity of tor's path selection. In *Proceedings of the 21st ACM Conference on Computer and Communications Security (CCS)*, pages 513–524. ACM, 2014.

[34] M. Backes, A. Kate, S. Meiser, and T. Ruffing. Secrecy without perfect randomness: Cryptography with (bounded) weak sources. In *Proceedings of the 13th International Conference on Applied Cryptography and Network Security (ACNS)*, page to appear. Springer, 2015.

[35] M. Backes and B. Koepf. Quantifying information flow in cryptographic systems. *Mathematical Structures in Computer Science*, 25(02):457–479, 2015.

[36] M. Backes, P. Manoharan, and E. Mohammadi. TUC: Time-sensitive and modular analysis of anonymous communication. In *Proceedings of the of the 27th IEEE Computer Security Foundations Symposium (CSF)*, pages 383–397. IEEE, 2014.

[37] M. Backes, E. Mohammadi, and T. Ruffing. Computational soundness results for proverif - bridging the gap from trace properties to uniformity. In *Proceedings of the 3rd Conference on Principles of Security and Trust (POST)*, pages 42–62. Springer, 2014.

[38] M. Backes, E. Mohammadi, and T. Ruffing. Computational soundness for interactive primitves. In *Proceedings of the 20th European Symposium on Research in Computer Security (ESORICS)*, page to appear. Springer, 2015.

[39] M. Backes and S. Nürnberger. Oxymoron - making fine-grained memory randomization practical by allowing code sharing. In *Proceedings of the 23th USENIX Security Symposium (USENIX)*, pages 433–447. USENIX, 2014.

[40] L. Backstrom, J. Kleinberg, L. Lee, and C. Danescu-Niculescu-Mizil. Characterizing And Curating Conversation Threads: Expansion, Focus, Volume, Re-Entry, year = 2013, pages = 13–22, booktitle = Proceedings of WSDM.

[41] S. Baruah and B. Brandenburg. Multiprocessor feasibility analysis of recurrent task systems with specified processor affinities. In *Proceedings of the 34th IEEE Real-Time Systems Symposium*, RTSS'13, pages 160–169, 2013.

[42] P. Bhatotia, U. A. Acar, F. P. Junqueira, and R. Rodrigues. Slider: Incremental sliding window analytics. In *Proceedings of the 15th International Middleware Conference*, Middleware '14, pages 61–72. ACM, 2014.

[43] P. Bhatotia, P. Fonseca, U. A. Acar, B. B. Brandenburg, and R. Rodrigues. iThreads: A threading library for parallel incremental computation. In *Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS'15, pages 645–659. ACM, 2015.

[44] P. Bhatotia, A. Wieder, U. A. Acar, and R. Rodrigues. Incremental mapreduce computations. In S. Sakr and M. Gaber, editors, *Large Scale and Big Data: Processing and Management*. CRC Press, 2014.

[45] P. Bhattacharya, S. Ghosh, M. B. Zafar, J. Kulshrestha, M. Mondal, N. Ganguly, and K. P. Gummadi. Deep Twitter Diving: Exploring Topical Groups in Microblogs at Scale. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW'14)*, 2014.

[46] P. Bhattcharya, M. B. Zafar, N. Ganguly, S. Ghosh, and K. P. Gummadi. Inferring User Interests in the Twitter Social Network. In *Proceedings of the 8th ACM Conference on Recommender Systems (RecSys'14)*, 2014.

[47] A. Bichhawat, V. Rajani, D. Garg, and C. Hammer. Generalizing permissive-upgrade in dynamic information flow analysis. In *Proc. PLAS'14*, 2014.

[48] A. Bichhawat, V. Rajani, D. Garg, and C. Hammer. Information flow control in WebKit's JavaScript bytecode. In *Proc. POST*, 2014.

[49] A. Biere and R. Bloem, editors. *Computer Aided Verification - 26th International Conference, CAV 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 18-22, 2014. Proceedings*, volume 8559 of *Lecture Notes in Computer Science*. Springer, 2014.

[50] S. L. Blond, A. Uritesc, C. Gilbert, Z. L. Chua, P. Saxena, and E. Kirda. A look at targeted attacks through the lense of an ngo. In *23rd USENIX Security Symposium (USENIX Security 14)*, pages 543–558, San Diego, CA, Aug. 2014. USENIX Association.

[51] J. K. Bock. Syntactic persistence in language production. *Cognitive Psychology*, 18(3):355 – 387, 1986.

[52] B. Brandenburg. *Scheduling and locking in multiprocessor real-time operating systems*. PhD thesis, UNC Chapel Hill, 2011.

[53] B. Brandenburg. A fully preemptive multiprocessor semaphore protocol for latency-sensitive real-time applications. In *Proceedings of the 25th Euromicro Conference on Real-Time Systems*, pages 292–302, 2013.

[54] B. Brandenburg. Blocking optimality in distributed real-time locking protocols. *Leibniz Transactions on Embedded Systems*, 1(2):1–22, 2014.

[55] B. Brandenburg. The FMLP$^+$: An asymptotically optimal real-time locking protocol for suspension-aware analysis. In *Proceedings of the 26th Euromicro Conference on Real-Time Systems*, ECRTS'14, pages 196–206, 2014.

[56] B. Brandenburg. A synchronous IPC protocol for predictable access to shared resources in mixed-criticality systems. In *Proceedings of the 35th IEEE Real-Time Systems Symposium*, RTSS'14, pages 196–206, 2014.

[57] P. Brown and S. C. Levinson. *Politeness: Some universals in language usage*. Cambridge University Press, 1987.

[58] G. Calin, E. Derevenetc, R. Majumdar, and R. Meyer. A theory of partitioned global address spaces. In A. Seth and N. K. Vishnoi, editors, *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2013, December*

12-14, 2013, Guwahati, India, volume 24 of *LIPIcs*, pages 127–139. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2013.

[59] F. Cerqueira, A. Gujarati, and B. Brandenburg. Linux's processor affinity API, refined: *Shifting* real-time tasks towards higher schedulability. In *Proceedings of the 35th IEEE Real-Time Systems Symposium*, RTSS'14, pages 249–259, 2014.

[60] F. Cerqueira, M. Vanga, and B. Brandenburg. Scaling global scheduling with message passing. In *Proceedings of the 20th IEEE Real-Time and Embedded Technology and Applications Symposium*, RTAS'14, pages 263–274, 2014.

[61] S. Chakraborty, T. A. Henzinger, A. Sezgin, and V. Vafeiadis. Aspect-oriented linearizability proofs. *Logical Methods in Computer Science*, 11(1:20)2015.

[62] K. Chatterjee, R. Ibsen-Jensen, and R. Majumdar. Edit distance for timed automata. In M. Fränzle and J. Lygeros, editors, *17th International Conference on Hybrid Systems: Computation and Control (part of CPS Week), HSCC'14, Berlin, Germany, April 15-17, 2014*, pages 303–312. ACM, 2014.

[63] J. Cheng, C. Danescu-Niculescu-Mizil, and J. Leskovec. How community feedback shapes user behavior. In *Proceedings of ICWSM*, 2014.

[64] D. V. Chistikov. Notes on counting with finite machines. In *34th International Conference on Foundation of Software Technology and Theoretical Computer Science, FSTTCS 2014, December 15-17, 2014, New Delhi, India*, pages 339–350, 2014.

[65] D. V. Chistikov, R. Dimitrova, and R. Majumdar. Approximate counting in SMT and value estimation for probabilistic programs. In C. Baier and C. Tinelli, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 21st International Conference, TACAS 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings*, volume 9035 of *Lecture Notes in Computer Science*, pages 320–334. Springer, 2015.

[66] D. V. Chistikov and R. Majumdar. Unary pushdown automata and straight-line programs. In J. Esparza, P. Fraigniaud, T. Husfeldt, and E. Koutsoupias, editors, *Automata, Languages, and Programming -*

41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part II, volume 8573 of *Lecture Notes in Computer Science*, pages 146–157. Springer, 2014.

[67] O. Chowdhury, D. Garg, L. Jia, and A. Datta. Equivalence-based security for querying encrypted databases: Theory and application to privacy policy audits. In *Proc. CCS*, 2015. To appear.

[68] O. Chowdhury, L. Jia, D. Garg, and A. Datta. Temporal mode-checking for runtime monitoring of privacy policies. In *Proc. CAV*, 2014.

[69] E. Cicek, D. Garg, and U. Acar. Refinement types for incremental computational complexity. In *Proc. ESOP*, 2015.

[70] D. Correa, L. A. Silva, M. Mondal, F. Benevenuto, K. P. Gummadi, and S. Ghosh. Characterizing Information Diets of Social Media Users. In *Proceedings of the 9th International AAAI Conference on Weblogs and Social Media (ICWSM'15)*, 2015.

[71] C. Danescu-Niculescu-Mizil, M. Gamon, and S. Dumais. Mark my words! Linguistic style accommodation in social media. In *Proceedings of WWW*, pages 745–754, 2011.

[72] C. Danescu-Niculescu-Mizil, G. Kossinets, J. Kleinberg, and L. Lee. How opinions are received by online communities: A case study on Amazon.com helpfulness votes. In *Proceedings of WWW*, pages 141–150, 2009.

[73] C. Danescu-Niculescu-Mizil, L. Lee, B. Pang, and J. Kleinberg. Echoes of power: Language effects and power differences in social interaction. In *Proceedings of WWW*, 2012.

[74] C. Danescu-Niculescu-Mizil, M. Sudhof, D. Jurafsky, J. Leskovec, and C. Potts. A computational approach to politeness with application to social factors. In *Proceedings of ACL*, 2013.

[75] C. Danescu-Niculescu-Mizil, R. West, D. Jurafsky, J. Leskovec, and C. Potts. No country for old members: User lifecycle and linguistic change in online communities. In *Proceedings of WWW*, 2013.

[76] P. R. D'Argenio and H. C. Melgratti, editors. *CONCUR 2013 - Concurrency Theory - 24th International Conference, CONCUR 2013,*

Buenos Aires, Argentina, August 27-30, 2013. Proceedings, volume 8052 of *Lecture Notes in Computer Science*. Springer, 2013.

[77] E. Darulova, V. Kuncak, R. Majumdar, and I. Saha. Synthesis of fixed-point programs. In *Proceedings of the International Conference on Embedded Software, EMSOFT 2013, Montreal, QC, Canada, September 29 - Oct. 4, 2013*, pages 22:1–22:10. IEEE, 2013.

[78] A. Datta, D. Garg, D. Kaynar, D. Sharma, and A. Sinha. Program actions as actual causes: A building block for accountability. In *Proc. CSF*, 2015. To appear.

[79] J. V. Deshmukh, R. Majumdar, and V. S. Prabhu. Quantifying conformance using the skorokhod metric. In D. Kroening and C. S. Pasareanu, editors, *Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part II*, volume 9207 of *Lecture Notes in Computer Science*, pages 234–250. Springer, 2015.

[80] R. Dimitrova and R. Majumdar. Deductive control synthesis for alternating-time logics. In T. Mitra and J. Reineke, editors, *2014 International Conference on Embedded Software, EMSOFT 2014, New Delhi, India, October 12-17, 2014*, pages 14:1–14:10. ACM, 2014.

[81] D. Dreyer, G. Neis, and L. Birkedal. The impact of higher-order state and control effects on local relational reasoning. *Journal of Functional Programming*, 22(4&5):477–528, Sept. 2012. Special issue devoted to selected papers from ICFP 2010.

[82] A. Durand-Gasselin, J. Esparza, P. Ganty, and R. Majumdar. Model checking parameterized asynchronous shared-memory systems. In Kroening and Pasareanu [129], pages 67–84.

[83] E. Elnikety, A. Mehta, A. Vahldiek-Oberwagner, D. Garg, and P. Druschel. Thoth: Ensuring compliance in data retrieval systems. Submitted for publication.

[84] E. Elnikety, A. Vahldiek, A. Mehta, D. Garg, and P. Druschel. Thoth: Efficiently enforcing data confidentiality and integrity in distributed data processing systems, Oct. 2013. Poster at SOSP 2013.

[85] M. Emmi, P. Ganty, R. Majumdar, and F. Rosa-Velardo. Analysis of asynchronous programs with event-based synchronization. In J. Vitek,

editor, *Programming Languages and Systems - 24th European Symposium on Programming, ESOP 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings*, volume 9032 of *Lecture Notes in Computer Science*, pages 535–559. Springer, 2015.

[86] S. Esmaeilsabzali, R. Majumdar, T. Wies, and D. Zufferey. Dynamic package interfaces. In S. Gnesi and A. Rensink, editors, *Fundamental Approaches to Software Engineering - 17th International Conference, FASE 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014, Proceedings*, volume 8411 of *Lecture Notes in Computer Science*, pages 261–275. Springer, 2014.

[87] J. Esparza, P. Ganty, J. Leroux, and R. Majumdar. Verification of population protocols. In *CONCUR 2015 - Concurrency Theory - 26th International Conference*, 2015.

[88] J. Esparza, P. Ganty, and R. Majumdar. Parameterized verification of asynchronous shared-memory systems. In N. Sharygina and H. Veith, editors, *Computer Aided Verification - 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13-19, 2013. Proceedings*, volume 8044 of *Lecture Notes in Computer Science*, pages 124–140. Springer, 2013.

[89] J. Esparza, R. Ledesma-Garza, R. Majumdar, P. Meyer, and F. Niksic. An smt-based approach to coverability analysis. In Biere and Bloem [49], pages 603–619.

[90] M. Farajtabar, N. Du, M. Gomez-Rodriguez, I. Valera, H. Zha, and L. Song. Shaping Social Activity by Incentivizing Users. In *Advances in Neural Information Processing Systems*, 2014.

[91] M. Farajtabar, M. Gomez-Rodriguez, N. Du, M. Zamani, H. Zha, and L. Song. Back to the Past: Source Identification in Diffusion Networks from Partially Observed Cascades. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, 2015.

[92] M. Farajtabar, M. Gomez-Rodriguez, Y. Wang, S. Li, H. Zha, and L. Song. Co-evolutionary dynamics of information diffusion and network structure. In *Workshop in Diffusion, Activity and Events in*

Networks: Models Methods and Applications at the 24th International World Wide Web Conference, 2015.

[93] F. Figueiredo, J. Almeida, F. Benevenuto, and K. P. Gummadi. Does Content Determine Information Popularity in Social Media? A Case Study of YouTube Videos' Content and their Popularity. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'14)*, 2014.

[94] P. Fonseca, R. Rodrigues, and B. B. Brandenburg. SKI: Exposing kernel concurrency bugs through systematic schedule exploration. In *Proceedings of the 11th USENIX Conference on Operating Systems Design and Implementation*, OSDI'14, pages 415–431, 2014.

[95] C. Freitas, F. Benevenuto, S. Ghosh, and A. Veloso. Reverse Engineering Socialbot Infiltration Strategies in the Twitter Social Network. In *Proceedings of the IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM'15)*, 2015.

[96] I. Gavran, F. Niksic, A. Kanade, R. Majumdar, and V. Vafeiadis. Rely-guarantee reasoning for asynchronous programs. In *CONCUR 2015 - Concurrency Theory - 26th International Conference*, 2015.

[97] S. Ghosh, N. Sharma, F. Benevenuto, N. Ganguly, and K. P. Gummadi. Cognos: Crowdsourcing Search for Topic Experts in Microblogs. In *Proceedings of the 35th Annual SIGIR Conference (SIGIR'12)*, 2012.

[98] S. Ghosh, B. Viswanath, F. Kooti, N. K. Sharma, K. Gautam, F. Benevenuto, N. Ganguly, and K. P. Gummadi. Understanding and Combating Link Farming in the Twitter Social Network. In *Proceedings of the 21st International World Wide Web Conference (WWW'12)*, 2012.

[99] S. Ghosh, M. B. Zafar, P. Bhattacharya, N. Sharma, N. Ganguly, and K. P. Gummadi. On Sampling the Wisdom of Crowds: Random vs. Expert Sampling of the Twitter Stream. In *Proceedings of the 22nd ACM International Conference on Information and Knowledge Management (CIKM'13)*, 2013.

[100] A. Girard and S. Sankaranarayanan, editors. *Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control, HSCC'15, Seattle, WA, USA, April 14-16, 2015*. ACM, 2015.

[101] M. Gligoric, R. Majumdar, R. Sharma, L. Eloussi, and D. Marinov. Regression test selection for distributed software histories. In Biere and Bloem [49], pages 293–309.

[102] I. Gog, M. Schwarzkopf, N. Crooks, M. P. Grosvenor, A. Clement, and S. Hand. Musketeer: All for one, one for all in data processing systems. In *Proceedings of the Tenth European Conference on Computer Systems*, EuroSys '15, pages 2:1–2:16, New York, NY, USA, 2015. ACM.

[103] I. Gog, M. Schwarzkopf, N. Crooks, M. P. Grosvenor, A. Clement, and S. Hand. Musketeer: all for one, one for all in data processing systems. In *Proceedings of the Tenth European Conference on Computer Systems, EuroSys 2015, Bordeaux, France, April 21-24, 2015*, page 2, 2015.

[104] O. Goga, P. Loiseau, R. Summer, R. Teixeira, and K. P. Gummadi. On the Reliability of Profile Matching Across Large Online Social Networks. In *Proceedings of the 21st ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'15)*, 2015.

[105] O. Goga, G. Venkatadri, and K. P. Gummadi. The Doppelgänger Bot Attack: Exploring Identity Impersonation in Online Social Networks. In *Proceedings of the 15th ACM SIGCOMM Conference on Internet Measurement (IMC'15)*, 2015.

[106] M. Gomez-Rodriguez, D. Balduzzi, and B. Schölkopf. Uncovering the temporal dynamics of diffusion networks. In *Proceedings of the 28th International Conference on Machine Learning*, 2011.

[107] M. Gomez-Rodriguez, K. P. Gummadi, and B. Schölkopf. Quantifying Information Overload in Social Media and its Impact on Social Contagions. In *Proceedings of the 8th International AAAI Conference on Weblogs and Social Media (ICWSM'14)*, 2014.

[108] M. Gomez-Rodriguez, J. Leskovec, and A. Krause. Inferring Networks of Diffusion and Influence. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery in Data Mining*, 2010.

[109] M. Gomez-Rodriguez, J. Leskovec, and B. Schölkopf. Modeling information propagation with survival theory. In *Proceedings of the 30th International Conference on Machine Learning*, 2013.

[110] M. Gomez-Rodriguez, L. Song, H. Daneshmand, and B. Schölkopf. Estimating diffusion networks: Recovery conditions, sample complexity & soft-thresholding algorithm. *Journal of Machine Learning Research*, 2015.

[111] M. Gomez-Rodriguez, L. Song, N. Du, H. Zha, and B. Schölkopf. stimating diffusion networks: Recovery conditions, sample complexity & soft-thresholding algorithm. *ACM Transactions on Information Systems*, 2015.

[112] A. Gujarati, F. Cerqueira, and B. Brandenburg. Schedulability analysis of the Linux push and pull scheduler with arbitrary processor affinities. In *Proceedings of the 25th Euromicro Conference on Real-Time Systems*, pages 69–79, 2013.

[113] A. Gujarati, F. Cerqueira, and B. Brandenburg. Multiprocessor real-time scheduling with arbitrary processor affinities: From practice to theory. *Real-Time Systems*, 51(4):440–483, 2015.

[114] F. Heider. Attitudes and Cognitive Organization. *Journal of Psychology*, 21:107–112, 1946.

[115] R. Hüchting, R. Majumdar, and R. Meyer. A theory of name boundedness. In D'Argenio and Melgratti [76], pages 182–196.

[116] R. Hüchting, R. Majumdar, and R. Meyer. Bounds on mobility. In P. Baldan and D. Gorla, editors, *CONCUR 2014 - Concurrency Theory - 25th International Conference, CONCUR 2014, Rome, Italy, September 2-5, 2014. Proceedings*, volume 8704 of *Lecture Notes in Computer Science*, pages 357–371. Springer, 2014.

[117] C.-K. Hur and D. Dreyer. A Kripke logical relation between ML and assembly. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, 2011.

[118] C.-K. Hur, D. Dreyer, G. Neis, and V. Vafeiadis. The marriage of bisimulations and Kripke logical relations. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, 2012.

[119] L. Jia, D. Garg, and A. Datta. A logic of programs with interface-confined code. In *Proc. CSF*, 2015. To appear.

[120] X. Jia, W. Li, and V. Vafeiadis. Proving lock-freedom easily and automatically. In X. Leroy and A. Tiu, editors, *Proceedings of the 2015 Conference on Certified Programs and Proofs, CPP 2015, Mumbai, India, January 15-17, 2015*, pages 119–127. ACM, 2015.

[121] R. Jung, D. Swasey, F. Sieczkowski, K. Svendsen, A. Turon, L. Birkedal, and D. Dreyer. Iris: Monoids and invariants as an orthogonal basis for concurrent reasoning. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, 2015.

[122] J. Kang, C. Hur, W. Mansky, D. Garbuzov, S. Zdancewic, and V. Vafeiadis. A formal C memory model supporting integer-pointer casts. In D. Grove and S. Blackburn, editors, *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation, Portland, OR, USA, June 15-17, 2015*, pages 326–335. ACM, 2015.

[123] J. Kang, Y. Kim, C.-K. Hur, D. Dreyer, and V. Vafeiadis. Lightweight verification of separate compilation. Submitted for publication, July 2015.

[124] M. Kapritsos, Y. Wang, V. Quema, A. Clement, L. Alvisi, and M. Dahlin. All about Eve: execute-verify replication for multi-core servers. In *Proceedings of the 10th USENIX conference on Operating Systems Design and Implementation*, OSDI'12, pages 237–250, Berkeley, CA, USA, 2012. USENIX Association.

[125] S. Kilpatrick, D. Dreyer, S. Peyton Jones, and S. Marlow. Backpack: Retrofitting Haskell with interfaces. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, 2014.

[126] J. Kloos, R. Majumdar, and V. Vafeiadis. Asynchronous liquid separation types. In J. T. Boyland, editor, *29th European Conference on Object-Oriented Programming, ECOOP 2015, July 5-10, 2015, Prague, Czech Republic*, volume 37 of *LIPIcs*, pages 396–420. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.

[127] F. Kooti, M. Cha, K. P. Gummadi, and W. Mason. The Emergence of Conventions in Online Social Networks. In *Proceedings of the 6th International AAAI Conference on Weblogs and Social Media (ICWSM'12)*, 2012.

[128] F. Kooti, W. Mason, K. P. Gummadi, and M. Cha. Predicting Emerging Social Conventions in Online Social Networks. In *Proceedings of the 21st ACM International Conference on Information and Knowledge Management (CIKM'12)*, 2012.

[129] D. Kroening and C. S. Pasareanu, editors. *Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part I*, volume 9206 of *Lecture Notes in Computer Science*. Springer, 2015.

[130] J. Kulshrestha, F. Kooti, A. Nikravesh, and K. P. Gummadi. Geographic Dissection of the Twitter Network. In *Proceedings of the 6th International AAAI Conference on Weblogs and Social Media (ICWSM'12)*, 2012.

[131] J. Kulshrestha, M. B. Zafar, L. E. Noboa, K. P. Gummadi, and S. Ghosh. Characterizing Information Diets of Social Media Users. In *Proceedings of the 9th International AAAI Conference on Weblogs and Social Media (ICWSM'15)*, 2015.

[132] L. Kuper, A. Turon, N. R. Krishnaswami, and R. R. Newton. Freeze after writing: Quasi-deterministic parallel programming with LVars. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, 2014.

[133] O. Lahav and V. Vafeiadis. Owicki-gries reasoning for weak memory models. In M. M. Halldórsson, K. Iwama, N. Kobayashi, and B. Speckmann, editors, *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part II*, volume 9135 of *Lecture Notes in Computer Science*, pages 311–323. Springer, 2015.

[134] R. Lakoff. The Logic of Politeness; Or, Miding Your P's and Q's. In *Proceedings of the 9th Meeting of the Chicago Linguistic Society*, pages 292–305, 1973.

[135] S. Le Blond, D. Choffnes, W. Caldwell, P. Druschel, and N. Merritt. Herd: A Scalable, Traffic Analysis Resistant Anonymity Network for VoIP Systems. In *Proceedings of the ACM SIGCOMM Conference*, SIGCOMM '15, 2015.

[136] M. Lentz, V. Erdélyi, P. Aditya, E. Shi, P. Druschel, and B. Bhattacharjee. Sddr: Light-weight, secure mobile encounters. In *23rd*

USENIX Security Symposium (USENIX Security 14), pages 925–940, San Diego, CA, Aug. 2014. USENIX Association.

[137] X. Leroy. Formal verification of a realistic compiler. *Communications of the ACM*, 52(7), 2009.

[138] C. Li, J. a. Leitão, A. Clement, N. Preguiça, and R. Rodrigues. Minimizing coordination in replicated systems. In *Proceedings of the First Workshop on Principles and Practice of Consistency for Distributed Data*, PaPoC '15. ACM.

[139] C. Li, J. a. Leitão, A. Clement, N. Preguiça, R. Rodrigues, and V. Vafeiadis. Automating the choice of consistency levels in replicated systems. In *Proceedings of the 2014 USENIX Conference on USENIX Annual Technical Conference*, USENIX ATC'14, pages 281–292, 2014.

[140] Y. Liu, K. P. Gummadi, B. Krishnamurthy, and A. Mislove. Analyzing Facebook privacy settings: User expectations vs. reality. In *Proceedings of the 11th ACM SIGCOMM Conference on Internet Measurement (IMC'11)*, 2011.

[141] Y. Liu, B. Viswanath, M. Mondal, K. P. Gummadi, and A. Mislove. Simplifying Friendlist Management. In *Proceedings of the 21st International World Wide Web Conference (WWW'12), Demo Paper*, 2012.

[142] P. Maiya, A. Kanade, and R. Majumdar. Race detection for android applications. In M. F. P. O'Boyle and K. Pingali, editors, *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '14, Edinburgh, United Kingdom - June 09 - 11, 2014*, page 34. ACM, 2014.

[143] R. Majumdar and V. S. Prabhu. Computing the skorokhod distance between polygonal traces. In Girard and Sankaranarayanan [100], pages 199–208.

[144] R. Majumdar, S. D. Tetali, and Z. Wang. Kuai: A model checker for software-defined networks. In *Formal Methods in Computer-Aided Design, FMCAD 2014, Lausanne, Switzerland, October 21-24, 2014*, pages 163–170. IEEE, 2014.

[145] R. Majumdar and Z. Wang. Expand, enlarge, and check for branching vector addition systems. In D'Argenio and Melgratti [76], pages 152–166.

[146] R. Majumdar and Z. Wang. Bbs: A phase-bounded model checker for asynchronous programs. In Kroening and Pasareanu [129], pages 496–503.

[147] M. Mondal, P. Druschel, K. P. Gummadi, and A. Mislove. Beyond Access Control: Managing Online Privacy via Exposure. In *Proceedings of the Workshop on Useable Security (USEC'14)*, San Diego, CA, February 2014.

[148] M. Mondal, P. Druschel, K. P. Gummadi, and A. Mislove. Deep Twitter Diving: Exploring Topical Groups in Microblogs at Scale. In *Proceedings of the Workshop on Usable Security (USEC'14)*, 2014.

[149] M. Mondal, B. Viswanath, A. Clement, P. Druschel, K. P. Gummadi, A. Mislove, and A. Post. Defending against large-scale crawls in online social networks. In *Proceedings of the 8th international conference on Emerging Networking Experiments and Technologies*, CoNEXT '12, New York, NY, USA, 2012.

[150] M. Mondal, B. Viswanath, K. P. Gummadi, and A. Mislove. Deep Twitter Diving: Exploring Topical Groups in Microblogs at Scale. In *Proceedings of the 10th Symposium on Usable Security and Privacy (SOUPS'14)*, 2014.

[151] G. Neis, C.-K. Hur, J.-O. Kaiser, C. McLaughlin, D. Dreyer, and V. Vafeiadis. Pilsner: A compositionally verified compiler for a higher-order imperative language. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*, 2015.

[152] V. Niculae and C. Danescu-Niculescu-Mizil. Brighter than gold: Figurative language in user generated comparisons. In *Proceedings of EMNLP*, October 2014.

[153] K. G. Niederhoffer and J. W. Pennebaker. Linguistic style matching in social interaction, journal = Journal of Language and Social Psychology, year = 2002, volume = 21, number = 4, pages = 337–360,.

[154] S. E. Oh, J. Y. Chun, L. Jia, D. Garg, C. A. Gunter, and A. Datta. Privacy-preserving audit for broker-based health information exchange. In *Proc. CODASPY*, 2014.

[155] B. Pang and L. Lee. Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, 2(1-2):1–135, 2008.

[156] D. Porto, J. a. Leitão, C. Li, A. Clement, A. Kate, F. Junqueira, and R. Rodrigues. Visigoth fault tolerance. In *Proceedings of the Tenth European Conference on Computer Systems*, EuroSys '15. ACM, 2015.

[157] V. Rajani, A. Bichhawat, D. Garg, and C. Hammer. Fine-grained information flow control for event handling and the DOM. In *Proc. CSF*, 2015. To appear.

[158] M. Recasens, C. Danescu-Niculescu-Mizil, and D. Jurafsky. Linguistic Models for Analyzing and Detecting Biased Language. In *Proceedings of ACL*, 2013.

[159] D. Reitter, F. Keller, and J. D. Moore. A computational cognitive model of syntactic priming. *Cogn Sci*, 35(4):587–637, 2011.

[160] A. Reznichenko and P. Francis. Private-by-Design Advertising Meets the Real World. In *ACM Conference on Computer and Communications Security*, 2014.

[161] T. Rodrigues, F. Benevenuto, M. Cha, K. P. Gummadi, and V. Almeida. On Word-of-Mouth Based Discovery of the Web. In *Proceedings of the 11th ACM SIGCOMM Conference on Internet Measurement (IMC'11)*, 2011.

[162] A. Rossberg and D. Dreyer. Mixin' up the ML module system. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 35(1:2), Apr. 2013.

[163] A. Rossberg, C. Russo, and D. Dreyer. F-ing modules. *Journal of Functional Programming*, 24(5):529–607, Sept. 2014.

[164] K. Rudra, S. Ghosh, N. Ganguly, P. Goyal, and S. Ghosh. Extracting Situational Information from Microblogs during Disaster Events: A Classification-Summarization Approach. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management (CIKM'15)*, 2015.

[165] I. Saha, S. Baruah, and R. Majumdar. Dynamic scheduling for networked control systems. In Girard and Sankaranarayanan [100], pages 98–107.

[166] N. Sharma, S. Ghosh, F. Benevenuto, N. Ganguly, and K. P. Gummadi. Inferring Who-is-Who in the Twitter Social Network. In *Pro-*

ceedings of the 4th ACM SIGCOMM Workshop On Social Networks (WOSN'12), 2012.

[167] M. Simeonovski, F. Bendun, M. R. Asghar, M. Backes, N. Marnau, and P. Druschel. Oblivion: Mitigating privacy leaks by controlling the discoverability of online information, 2015.

[168] R. Spliet, M. Vanga, B. Brandenburg, and S. Dziadek. Fast on average, predictable in the worst case: Exploring real-time futexes in LITMUS$^{\mathrm{RT}}$. In *Proceedings of the 35th IEEE Real-Time Systems Symposium*, RTSS'14, pages 96–105, 2014.

[169] G. Stewart, L. Beringer, S. Cuellar, and A. W. Appel. Compositional CompCert. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, 2015.

[170] G. W. Subhabrata Mukherjee and C. Danescu-Niculescu-Mizil. People on drugs: Credibility of user statements in health communities. In *Proceedings of the 20th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'14)*, 2014.

[171] P. Tabuada, S. Y. Caliskan, M. Rungger, and R. Majumdar. Towards robustness for cyber-physical systems. *IEEE Trans. Automat. Contr.*, 59(12):3151–3163, 2014.

[172] J. Tassarotti, D. Dreyer, and V. Vafeiadis. Verifying read-copy-update in a logic for weak memory. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, 2015.

[173] A. Turon. *Understanding and Expressing Scalable Concurrency.* PhD thesis, Northeastern University, Apr. 2013.

[174] A. Turon, D. Dreyer, and L. Birkedal. Unifying refinement and Hoare-style reasoning in a logic for higher-order concurrency. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*, 2013.

[175] A. Turon, J. Thamsborg, A. Ahmed, L. Birkedal, and D. Dreyer. Logical relations for fine-grained concurrency. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, 2013.

[176] A. Turon, V. Vafeiadis, and D. Dreyer. GPS: Navigating weak memory with ghosts, protocols, and separation. In *ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, 2014.

[177] V. Vafeiadis, T. Balabonski, S. Chakraborty, R. Morisset, and F. Z. Nardelli. Common compiler optimisations are invalid in the C11 memory model and what we can do about it. In S. K. Rajamani and D. Walker, editors, *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*, pages 209–220. ACM, 2015.

[178] V. Vafeiadis and C. Narayan. Relaxed separation logic: A program logic for C11 concurrency. In *ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, 2013.

[179] A. Vahldiek-Oberwagner, E. Elnikety, A. Mehta, D. Garg, P. Druschel, R. Rodrigues, J. Gehrke, and A. Post. Guardat: enforcing data policies at the storage layer. In *Proceedings of the Tenth European Conference on Computer Systems, EuroSys 2015, Bordeaux, France, April 21-24, 2015*, page 13, 2015.

[180] A. Vahldiek-Oberwagner, E. Elnikety, A. Mehta, D. Garg, P. Druschel, R. Rodrigues, J. Gehrke, and A. Post. Guardat: Enforcing data policies at the storage layer. In *Proc. EuroSys*, 2015.

[181] B. Viswanath, M. A. Bashir, M. Crovella, S. Guha, K. P. Gummadi, B. Krishnamurthy, and A. Mislove. Towards Detecting Anomalous User Behavior in Online Social Networks. In *Proceedings of the 23rd Usenix Security Symposium (Usenix Security'14)*, 2014.

[182] B. Viswanath, M. A. Bashir, M. B. Zafar, S. Bouget, S. Guha, K. P. Gummadi, A. Kate, and A. Mislove. Strength in Numbers: Robust Tamper Detection in Crowd Computations. In *Proceedings of the 3rd ACM Conference on Online Social Networks (COSN'15)*, 2015.

[183] B. Viswanath, M. Mondal, A. Clement, P. Druschel, K. P. Gummadi, A. Mislove, and A. Post. Exploring the design space of social network-based Sybil defense. In *Proceedings of the Third International Conference on Communication Systems and Networking (COMSNETS'12)*, Bangalore, India, 2012. Invited paper.

[184] B. Viswanath, M. Mondal, K. P. Gummadi, A. Mislove, and A. Post. Canal: Scaling Social Network-based Sybil Tolerance Schemes. In *Proceedings of the 7th European Conference on Computer Systems (EuroSys'12)*, 2012.

[185] A. Wieder and B. Brandenburg. On spin locks in AUTOSAR: blocking analysis of FIFO, unordered, and priority-ordered spin locks. In *Proceedings of the 34th IEEE Real-Time Systems Symposium*, RTSS'13, pages 45–56, 2013.

[186] A. Wieder and B. Brandenburg. On the complexity of worst-case blocking analysis of nested critical sections. In *Proceedings of the 35th IEEE Real-Time Systems Symposium*, RTSS'14, pages 106–117, 2014.

[187] B. Zafar, I. Valera, M. Gomez-Rodriguez, and K. Gummadi. Fairness constraints: A mechanism for fair classification. In *Workshop in Fairness, Accountability and Transparency in Machine Learning at the 32nd International Conference on Machine Learning*, 2015.

[188] M. B. Zafar, P. Bhattacharya, N. Ganguly, K. P. Gummadi, and S. Ghosh. On Sampling the Wisdom of Crowds: Random vs. Expert Sampling of the Twitter Stream. In *Proceedings of the ACM Journal on Transactions on the Web*, volume 9, pages 12:1–12:33, 2013.

[189] M. B. Zafar, I. Valera, and M. Gomez-Rodriguez. Fairness Constraints: A Mechanism for Fair Classification. In *Proceedings of the 2nd Workshop on Fainess, Accountability, and Transparency in Machine Learning (FATML'15)*, 2015.

[190] M. Zamani, P. M. Esfahani, R. Majumdar, A. Abate, and J. Lygeros. Symbolic control of stochastic systems via approximately bisimilar finite abstractions. *IEEE Trans. Automat. Contr.*, 59(12):3135–3150, 2014.

[191] B. Ziliani. *Interactive Typed Tactic Programming in the Coq Proof Assistant*. PhD thesis, Saarland University, Mar. 2015.

[192] B. Ziliani, D. Dreyer, N. R. Krishnaswami, A. Nanevski, and V. Vafeiadis. Mtac: A monad for typed tactic programming in Coq. *Journal of Functional Programming*, Aug. 2015. To appear in special issue devoted to selected papers from ICFP 2013.

[193] B. Ziliani and M. Sozeau. A unification algorithm for Coq featuring universe polymorphism and overloading. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*, 2015.