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

May 2011 – October 2013

Max
Planck
Institute
for
Software Systems

Contents

1 State of the Institute

This third progress report of the Max Planck Institute for Software Systems (MPI-SWS) covers the period May 2011 – October 2013. 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 14 research groups (10 current, 2 adjunct, and 2 former) that were active during this review period. Finally, Section 16 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. Those who have read the institute's previous progress report can safely skip ahead to Section 1.2.

1.1.1 Challenges in software systems

Software systems is the part of computer science that lays the foundation for the practical use of computers. We interpret the term broadly to include all areas of computer science 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; databases and information retrieval; programming languages and programming systems; software engineering and verification; security and privacy; and human-computer interaction.

With pervasive use of information technology, the field of software systems continues to see new challenges and opportunities:

- The rise of the Internet, personal mobile devices, and wireless data services have enabled the capture, transmission, aggregation, and search over vast amounts of digital information, and has placed this information at the fingertips of anyone who can afford Internet access. This trend continues to pose new challenges in scalable and fault-tolerant distributed systems and networking, data center computing and storage, and security.

- 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 many different languages at many different levels of abstraction. Developing programming methodologies and verification technologies that enable cost-effective design and verification remains a key challenge for software systems research.

- The Internet and web services like online social networks have enabled new forms of expression, communication and interaction for hundreds of millions of people, but pose important challenges for users' privacy, security, and accountability in the online world. The use of software systems to support and enhance a wide range of human activity leads to challenges that must be viewed within their non-technical (i.e., sociological, cultural, economic, and legal) context, for example, by understanding usability and human-computer interfaces.

As a growing research institute in software systems, we seek to build 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 82 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 (to London, Berlin, Munich, and Hamburg). 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, graphics and geometric modeling, 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.

The MPI-SWS has a total budget of just over EUR 10M per year and 17 faculty positions. Additional growth is expected through external funding. The institute buildings at the two sites jointly offer space for over 200 researchers and staff.

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. We expect that some of our graduates and postdocs will take faculty positions at German and European universities, thereby contributing to the strength and breadth of Software Systems research and education. 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 tenure-track model, similar to that of many internationally renowned computer science departments. It is the institute's policy that 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 offers junior faculty a long-term career path at the institute and thus a vital 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 German universities (although TU Munich has recently converted its computer science department—one of the top-ranked in Germany—to a tenure-track model). 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.

So far, the model has worked wonderfully. While the hirings of senior researchers Peter Druschel, Paul Francis, and Rupak Majumdar as scientific directors were notable successes, the reputation of MPI-SWS has grown significantly through the strength and depth of its junior faculty. Each of them has the potential to grow into a star within their respective research communities. In fact, our tenured faculty (Dreyer and Gummadi) are already established in their respective fields and are successfully leading ambitious research programs. The appointment of Michael Backes (UdS, Security) as a Max Planck Fellow and Robert Harper (CMU, Programming Languages) as an external scientific member have provided additional visibility.

Moreover, our demonstrated ability to attract and retain outstanding junior faculty has increased the institute's attractiveness in the eyes of senior candidates. MPI-SWS is no longer an unknown entity—it is now seen as an institute with a strong foundation, solid momentum, and a faculty that offers a vibrant intellectual environment. This was an important consideration, for instance, for both Paul Francis and Rupak Majumdar, and we have received similar signals from other candidates we approached.

A key factor in attracting top talent (both junior and senior) 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 students, travel and equipment. This funding has proved quite attractive to junior faculty, as it gets them started early, reduces the pressure to solicit third-party funding, and enables them to pursue collaborations through mutual visits with colleagues at other institutions. The opportunity to teach, combined with the freedom to choose how much and what they want to teach, is also attractive to some 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 invites candidates for individual interview visits. We interview faculty candidates in the February through April timeframe. The faculty then nominates their selected candidates to an external appointment committee led by the responsible MPS vice-president, 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 since the founding of the institute has been dealing with the fact that our institute is split between two locations: Kaiserslautern and Saarbrücken. Rather than allow the institute to be effectively split into two separate sub-institutes, we have instead made every effort to build a single, unified institute, with as much cross-site collaboration and shared institute culture as possible.

Toward this end, we have made several deliberate strategic decisions. First of all, 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 several 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. In the period that Kaiserslautern has been open, we have used a combination of face-to-face visits and telecommunication. During the first year of its existence (2009), the Kaiserslautern-site researchers visited Saarbrücken twice a month, while Saarbrücken visited Kaiserslautern once a month (in both directions using a chartered bus). Starting mid-year 2010, when Kaiserslautern approached the size of Saarbrücken, we balanced the visits, with each site visiting the other every other week, resulting cumulatively in one visit per week.

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. 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. 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 5 new tenure-track faculty: Björn Brandenburg (real-time systems), Allen Clement (robust systems), Cristian Danescu-Niculescu-Mizil (social information systems), Deepak Garg (foundations of computer security), and Ruzica Piskac (synthesis, analysis and automated reasoning). The institute has also recruited 9 new postdoctoral researchers and 26 new doctoral students.

Several junior faculty have also left the institute: Rodrigo Rodrigues accepted an associate professor position at NOVA University of Lisbon, Umut Acar accepted an assistant professor position at CMU, and Ruzica Piskac accepted an assistant professor position at Yale.

Publications and talks. MPI-SWS has produced 177 peer-reviewed publications during the reporting period. We have published in many of the top conferences and journals in software systems, including POPL, PLDI, CCS, S&P, OSDI, SIGCOMM, NSDI, JACM, TOPLAS, JFP, CSF, OOPSLA, ICFP, CAV, RTSS, LICS, EMSOFT, WWW, ICWSM, ACL, SIGIR, CIKM, IMC, Usenix Security, CoNEXT, and EuroSys. Of these publications, 20 are co-authored by members of two or more groups. Our publications are listed in the per-group sections (2–8).

MPI-SWS faculty gave 9 invited talks at conferences. We have collectively served as program chairs or co-chairs for 4 conferences and 5 workshops. MPI-SWS faculty and postdocs have served on the PCs of over 100 conferences and workshops. The details are given in the individual group sections.

Awards and honors. Björn Brandenburg's dissertation was recognized with three dissertation awards: UNC Chapel Hill's 2012 *Dean Linda Dyk-*

stra Distinguished Dissertation Award in the category "Mathematics, Physical Sciences and Engineering" (April 2012), the 2012 *Council of Graduate Schools / ProQuest Distinguished Dissertation Award* in the category "Mathematics, Physical Sciences and Engineering" (December 2012), and the 2012 *Distinguished Dissertation Award* of the European Design and Automation Association in the area "New directions in embedded systems design and embedded software" (March 2013). (The dissertation was completed prior to the reporting period; the awards were granted during the reporting period.)

Peter Druschel won the 2011 ACM/IFIP/Usenix Middleware 10-year Best Paper Award, together with Ant Rowstron (Microsoft Research, Cambridge). Paul Francis was awarded the 2011 SIGCOMM Test of Time award, together with his co-authors Mark Handley, Richard Karp, Sylvia Ratnasamy, and Scott Shenker. Rupak Majumdar was awarded the 2011 SIGPLAN PLDI Test of Time award, together with co-authors Tom Ball, Todd Millstein, and Sriram Rajamani.

Ruzica Piskac's paper on complete functional synthesis [143] was invited and published as a "Research Highlight" in the *Communications of the ACM*. She was also awarded the Patrick Denantes Prize in 2012 for her doctoral dissertation at EPFL.

Michael Backes was named the leading German scientist under the age of 40 by the Financial Times Germany, and one of the 100 most important IT people in Germany in 2011 by the Computerwoche Newspaper.

Institute members have also won a total of seven conference "Best Paper" awards:

- Rupak Majumdar won the EAPLS Best Paper award at the 2012 ETAPS conference (for his FASE 2012 paper [154],) and the 2011 ACM TODAES best paper award (for the best paper in all volumes of ACM Transactions on Design Automation of Embedded Systems in the year 2011) [76].

- Björn Brandenburg's ECRTS'13 paper [112] was awarded an "outstanding paper award" and his SIES'13 [206] and EMSOFT'11 [53] papers both won the "best paper" award at their respective conferences.

- Cristian Danescu-Niculescu-Mizil's work won the best paper award at WWW 2013 [81] and was nominated for the best paper award at ACL 2013 [176].

- Krishna Gummadi received the Best Paper Award at the AAAI's ICWSM 2012 conference [133].

External grant funding. Although the institute provides its faculty members with generous base funding to run their research groups, we have nevertheless actively encouraged all faculty to seek external funding, particularly in the form of prestigious ERC grants. 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.

Both Umut Acar and Rodrigo Rodrigues won ERC Starting Grants in 2012. Unfortunately, both of them left the institute (for positions at CMU and NOVA University of Lisbon, respectively) soon after obtaining their grants, but we are nonetheless very proud of their accomplishments.

A number of other faculty also secured external funding from Google, Microsoft Research, Yahoo!, Toyota, the European Commission, and the DFG. (For further details, see the individual sections as well as Section 16.)

Teaching. As with external funding, teaching is not a 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 (and especially the junior faculty) have taught 16 courses, including 6 core and 10 advanced seminar courses. (For further details, see Section 16.)

First two tenure cases. During the review period, the institute passed a major milestone, namely the handling of its first two tenure cases: Krishna Gummadi and Derek Dreyer. The tenure process has three stages: First, the institute requests letters from experts familiar with the candidate's research, discusses the case among the tenured faculty, and sends a dossier with the letters, the institute's evaluation of the case and a recommendation to a committee chaired by the Vice President of the CPT section of the MPS. This tenure committee includes scientific members of other Max Planck Institutes, as well as several external experts familiar with the candidate's area of research. Second, candidates present their work and future research agenda to the committee. Third, the committee may request additional letters and finally makes a recommendation to the President of the MPS. A successful tenure case is expected to show outstanding contributions in the candidate's area of research, as well as the potential and ambition to rise to the level of a scientific director in the MPS, or a senior position at another institution.

We are very pleased to report that both Gummadi and Dreyer were granted tenure!

New institute buildings. During the review period, construction of the institute buildings in both Saarbrücken and Kaiserslautern was completed, and we have now comfortably moved into these new buildings. This is a marked improvement over our previous situation, especially in Saarbrücken, where we were previously located off-campus (in the Wartburg building). We are very happy to return to campus, where we can engage fully and interactively with our colleagues in the Saarbrücken Graduate School.

Overall, the transition to our new buildings has been viewed very positively by the members of the institute. The new buildings are light and spacious, encouraging transparent and open communication between researchers within each location and between the two locations. A number of lecture rooms and meeting rooms in both locations are equipped with state-of-the-art videoconferencing technology, which we use regularly to enhance inter-site communication and collaboration, as well as to teach courses that can be cross-listed between the two sites.

One point of concern in the Saarbrücken location is that, aside from the kitchen areas on each floor, the new building does not currently have any appropriately furnished public meeting spaces where researchers from different backgrounds can comfortably congregate and engage in informal interactions. There are in fact several large areas in the building that could usefully serve this purpose, but they are currently underutilized. To realize the potential of these areas, we have contracted out with the Fraunhofer-Institut für Arbeitswirtschaft und Organisation to help us transform these areas into "communication hotspots", and we are currently in the process of implementing some of their design proposals.

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: Dmitry Chistikov (Moscow State University), Rayna Dimitrova (Saarland University), Shahram Esmaeilsabzali (University of Waterloo), Dario Fiore (University of Catania/NYU/École Normale Supérieure), Neel Krishnaswami (CMU/Microsoft Research), Stevens Le Blond (INRIA Sophia Antipolis), Roland Perera (University of Birmingham), and Aaron Turon (Northeastern University).

As this list suggests, we have generally been more successful recruiting postdocs in the area of programming languages and verification than in the area of systems/networking, where it is less common for top graduates (particularly those trained in the US) to do postdocs, and where there are fewer top graduates from Europe.

Placing our postdocs and students. Several of our postdocs have moved on to highly sought-after positions. Neel Krishnaswami moved to the University of Birmingham, where he is currently a Birmingham Fellow. Chung-Kil Hur moved onto a postdoc at Microsoft Research Cambridge, and subsequently accepted a tenure-track assistant professor position at Seoul National University. Arthur Charguéraud accepted a researcher position at INRIA Saclay. Allen Clement took a faculty position here at MPI-SWS, and Aniket Kate is now an independent group leader at the MMCI here at Saarland University.

We have had only a few students graduate during this review period. Matthew Hammer moved onto a postdoctoral position with Mike Hicks at the University of Maryland. Max Marcon joined SAP. Nuno Santos will be joining IST Portugal as an assistant professor.

Part I
Current Research Groups

2 The Real-Time Systems Group

2.1 Overview

The report covers the period from October 2011 – October 2013. The Real-Time Systems Group's efforts are focused on the theoretical foundations and practical challenges of multiprocessor real-time systems. The two major focus areas in the report period have been the design and analysis of predictable locking algorithms (i.e., locking algorithms that facilitate the derivation of *a priori* guarantees on maximum blocking) and practical real-time scheduling support in multiprocessor real-time operating systems.

Personnel. The group was newly formed in October 2011. It is led by Björn Brandenburg and currently has four graduate students (Alexander Wieder, Arpan Gujarati, Manohar Vanga, and Felipe Cerqueira).

Collaborations. The group maintains close ties with the Robust Systems Group (led by Allen Clement) and the Dependable Systems group (led by Rodrigo Rodrigues). Two members of the latter group, Pedro Fonseca and Pramod Bhatotia, are co-advised by Björn Brandenburg.

Allen Clement and Björn Brandenburg have jointly offered a seminar on fault-tolerant distributed real-time systems to expose students to topics in the intersection of the two research areas, and to foster new research directions. While no publications have resulted from this early-stage collaboration yet, several interesting research questions have been identified.

Externally, the group has collaborated with Andrea Bastoni of SYSGO AG,[1] which has resulted in a workshop paper, and with Sanjoy Baruah of UNC Chapel Hill, which has resulted in a forthcoming paper.

Publications. The group has published at all major real-time conferences. In the reporting period (2011–2013), group members have published papers at the following venues: RTSS [129], RTAS [52], ECRTS [51, 112], EMSOFT [53], SIES [206], at two workshops [55, 59], and in the journal *Design Automation for Embedded Systems* [54].

Two forthcoming papers [207, 39] co-authored by members of the group will be presented in December 2013 at RTSS'13.

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

Awards and honors. Arpan Gujarati's ECRTS'13 paper [112] was awarded an "outstanding paper award"; Alexander Wieder's SIES'13 paper [206] was recognized with the conference's "best paper" award; and Björn Brandenburg's EMSOFT'11 paper [53] won a "best paper" award.

Björn Brandenburg's dissertation was recognized with three dissertation awards: UNC Chapel Hill's 2012 *Dean Linda Dykstra Distinguished Dissertation Award* in the category "Mathematics, Physical Sciences and Engineering" (April 2012), the 2012 *Council of Graduate Schools / ProQuest Distinguished Dissertation Award* in the category "Mathematics, Physical Sciences and Engineering" (December 2012), and the 2012 *Distinguished Dissertation Award* of the European Design and Automation Association in the area "New directions in embedded systems design and embedded software" (March 2013). (The dissertation was completed prior to the reporting period; the awards were granted during the reporting period.)

Systems and tools. The group's most significant open-source project is LITMUS$^{\mathrm{RT}}$, a real-time extension of the Linux kernel, which has been continuously maintained since 2006 (as of October 2013, nine releases spanning 26 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 25 publications and three dissertations. Another recent software release is SchedCAT, a toolkit for schedulability and blocking analysis, which makes the group's analytical contributions available in runnable form.

Teaching. Björn Brandenburg offered two seminars on, resp., "Real-Time Scheduling and Synchronization" (Fall 2012) and "Fault-Tolerant Distributed Real-Time Systems" (Summer 2013, with Allen Clement), a special topics seminar (Privatissimum) on "Operating System Design and Implementation" (Summer 2013, with Allen Clement), and co-taught the operating systems core lecture (Summer 2013, with Peter Druschel).

Service. Internally to MPI-SWS, Björn Brandenburg has served on the graduate admissions committee in 2012, and chaired said committee in 2013.

Externally, he has served on the program committees of EuroSys'13, RTSS'12 and RTSS'13, ECRTS'13, RTAS'13, RTCSA'12, RTNS'12 and RTNS'13, OSPERT'12 and OSPERT'13, and as a reviewer for the journals *Real-Time Systems*, *Information Processing Letters*, *ACM Operating Systems Review*, *Leibniz Transactions on Embedded Systems*, *IEEE Transactions on Computers*, and for NSF and DFG.

2.2 Research agenda

The real-time systems group studies the algorithmic foundations, design, and implementation of real-time operating systems for multiprocessor platforms.

Motivation. Embedded *real-time systems* are pervasive in everyday technologies such as cars (e.g., engine control, anti lock brakes, etc.), airplanes (e.g., fly-by-wire), and robotics (e.g., signal processing, actuator control). Software in such safety-critical systems must be *predictable* in the sense that outputs should be produced *always at the right time.* For example, a car's airbag is of little use if triggered only fractions of a second late. In general, real-time constraints arise whenever computers interact with the physical environment, or when users perceive delays as unacceptable (e.g., in GUIs, during video playback, audio processing, video gaming, etc.).

Real-time operating systems (RTOSs) provide the foundation for predictable computing by enabling strong *a priori* guarantees on worst-case response times (i.e., on how long the system takes to react to external stimuli in the worst case), which fundamentally requires predictable resource arbitration policies: an RTOS must resolve contention for the processor(s) and other shared resources (e.g., physical memory, I/O ports, message buffers, shared data structures, etc.) such that all timing constraints of all tasks are always met. To this end, *real-time scheduling policies* and *real-time synchronization protocols* form the algorithmic foundation of predictable RTOSs.

RTOSs for *uniprocessors* are a well established, mature technology—in contrast to *multiprocessor* RTOSs, which still face considerable challenges. While uniprocessor RTOSs have been studied and used for decades, (shared-memory) multiprocessor real-time systems have gained widespread practical relevance only with the recent shift to multicores. Worse, virtually none of the known uniprocessor real-time scheduling polices and synchronization protocols is applicable to multiprocessors, in the sense that the desired analytical guarantees of most classical approaches (such as the optimality of *earliest-deadline first* (EDF) or *rate-monotonic* (RM) scheduling) hold only on uniprocessors. In the real-time domain, the widespread adoption of multicore platforms thus represents a revolutionary paradigm shift that necessitates the development of fundamentally new resource allocation policies.

Consequently, while RTOSs that run on multiprocessors are readily available, such current-generation multiprocessor RTOSs are not yet fully based on *analytically sound* foundations (in contrast to their uniprocessor predecessors) simply because these foundations have not yet been developed.

The real-time group's research efforts aim to bridge this gap. In the reporting period, the primary focus area was *predictable locking protocols* for sharing serially-reusable non-processor resources (e.g., I/O ports, shared data structures), complemented by two projects on practical *multiprocessor real-time scheduling* in Linux. The following discussion reviews the main results and then briefly provides an outlook on future research directions.

Predictable locking protocols. The purpose of a real-time locking protocol is to provide tasks with mutually exclusive access to shared resources such that the maximum blocking incurred by any task can be bounded *a priori.* Blocking is problematic as it leads to *priority inversions*, which intuitively exist when a high-priority task is delayed by a lower-priority task and may cause deadline violations if left unchecked. Research into real-time locking protocols is thus primarily concerned with **(i)** techniques for avoiding priority inversions in the first place, and **(ii)** establishing upper bounds on the maximum duration of priority inversion that cannot be avoided.

On a multiprocessor, there are two principal ways to realize blocking: in *semaphore* protocols, tasks wait by *suspending* (i.e., waiting tasks yield the processor to other tasks), whereas in *spin lock* protocols, tasks wait by *busy-waiting* (i.e., waiting tasks execute a tight delay loop) until gaining access to the contested resource. While busy-waiting wastes processor cycles, semaphore protocols incur higher OS overheads and risk losing cache affinity. A discussion of the practical and analytical advantages and disadvantages of either approach is beyond the scope of this overview; however, both techniques are used in practice and merit systematic study. We begin with the real-time group's recent results concerning semaphore protocols.

If tasks require exclusive access to shared resources, some blocking (and hence priority inversion) is inherently possible and cannot be avoided in general. This naturally raises the question of optimality: if priority inversion is inevitable when using locks, then what is the *minimal* bound on the worst-case duration of priority inversion that *any* locking protocol can guarantee? In other words, when is a real-time locking protocol optimal?

This question has been previously formalized in prior work by Brandenburg and Anderson [54]. In the reporting period, the real-time group substantially extended the basic theory of blocking optimality along three lines of work. The first contribution [53, 54] is the development and analysis of the first asymptotically optimal locking protocol for *clustered scheduling.* Under clustered scheduling, a multiprocessor real-time platform is divided into pairwise disjoint *clusters* of processors, where tasks are statically assigned to clusters and each cluster is independently scheduled (*partitioned* *scheduling*, the approach most commonly used in practice, falls into this category). The challenge is to ensure predictable *lock-holder progress* when sharing resources across cluster boundaries (i.e., lock holders should not be preempted while high-priority tasks are waiting in other clusters). However, since the schedulers operate independently, notions of "priority" are incomparable across cluster boundaries, which renders classical approaches such as *priority inheritance* ineffective. As a solution, a new progress mechanism called *priority donation* was introduced [53, 54], which enabled the design of several asymptotically optimal semaphore protocols (notably, the first multiprocessor real-time reader-writer and k-exclusion semaphore protocols).

An undesirable side effect of priority donation (and also all prior progress mechanisms for clustered and partitioned scheduling) is that it induces delays in unrelated tasks when it forces lock-holder progress, which can be problematic in the presence of *latency-sensitive* tasks (i.e., tasks that must be permitted to preempt lower-priority tasks immediately). This was solved in a recent paper [51], which introduced the first *independence-preserving* and asymptotically optimal semaphore protocol for clustered scheduling based on *migratory priority inheritance* [55, 51], wherein preempted lock holders migrate to the cluster of waiting tasks.

The third contribution (currently under submission to the *Leibniz Transactions on Embedded Systems*), motivated by the emergence of manycore platforms, extends the theory of blocking optimality from semaphore protocols for shared-memory multiprocessors to *distributed semaphore protocols*, where all critical sections must be executed on dedicated synchronization processors, and introduces the first two distributed real-time semaphore protocols that are asymptotically optimal in all cases.

Improved analysis of semaphores and spin locks. In practice, asymptotical bounds are too coarse-grained (and hence pessimistic) to obtain useful worst-case response times. Instead, *fine-grained* blocking bounds that reflect task arrival rates, critical section lengths, and other constant factors are required. A major contribution in the reporting period [52] is the development of a new, systematic analysis method for obtaining such fined-grained bounds based on *linear programming* (LP). To overcome the inherent pessimism in prior ad-hoc analysis methods, the new LP-based method expresses the bound on worst-case blocking as the objective function of an LP that, when maximized, yields a bound on the worst-case blocking across the set of all possible schedules. Protocol invariants and task parameters are then expressed as constraints that rule out impossible schedules. The resulting analysis is more concise, flexible, robust, and extensible than prior ad-hoc approaches, and most importantly, much less pessimistic [52].

The LP approach's versatility enabled another major contribution [207]. Motivated by the recent (and underspecified) inclusion of spin locks into the AUTOSAR RTOS standard (deployed in millions of cars), eight types of spin locks were analyzed using the LP-based approach: **(i)** FIFO-ordered spin locks, **(ii)** unordered spin locks, **(iii)** priority-ordered spin locks with unordered tie-breaking, and **(iv)** priority-ordered spin locks with FIFO-ordered tie-breaking, each analyzed assuming both preemptable and non-preemptable spinning. Seven of these lock types had not been analyzed in prior work. Concerning the sole exception (non-preemptable FIFO spin locks), the new analysis is asymptotically less pessimistic since no critical section is accounted for more than once. Finally, four concrete proposals for an improved AUTOSAR spin lock API were derived from the analysis [207].

Another work motivated by AUTOSAR investigated mapping tasks to processors while minimizing worst-case spin delays using integer LPs [206].

Real-time scheduling. Connecting theory and practice, the real-time systems group was the first to point out [112, 39] the theoretical implications of the support for *arbitrary processor affinities* in Linux and other RTOSs.

Ongoing and future work The real-time systems group currently pursues two major research directions. First, w.r.t. real-time synchronization, *nested critical sections* (i.e., when tasks acquire two or more locks) cause considerable analytical problems and cannot yet be handled accurately using existing analysis methods. An extension of the LP-based analysis approach to handle nesting is under way, targeting spin locks at first. Further, several major questions surrounding *independence-preserving real-time locking protocols* remain open (e.g., how to combine independence preservation and reader-writer synchronization). A project analyzing time-space tradeoffs in systems using *read-copy update* (RCU) synchronization is underway as well.

Second, w.r.t. systems issues, two projects exploring the use of real-time resource allocation policies to improve *predictability in data centers* are underway. Finally, while LITMUS$^{\mathrm{RT}}$ has served its intended purpose well, the limitations and tradeoffs inherent in a Linux-based environment have become more noticeable in recent projects; it may thus be time to invest into *a new multiprocessor RTOS based on sound analytical foundations*, both to serve as a testbed for future research and to reassess whether the algorithmic foundation available now is sufficient to support a full RTOS.

3 The Robust Systems Group

3.1 Overview

The report covers the period from July 2012 – October 2013. 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 is 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).

During the reporting period, the group had four graduate student interns (Natacha Crooks, 2 months; Manos Kapritos, 3 months; Priyanka Singla, 3 months; and Chunzhi Su, 3 months).

Collaborations. The group is 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 is engaged in collaborations with researchers at Cambridge University, Grenoble INP, Indian Institute of Science (Bangalore), Microsoft Research (Cambridge), La Sapienza (Universita di Roma), 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 have co-authored two OSDI [151, 128], one PODC [75], one Oakland [10], and one CoNEXT [168] papers.

Teaching. Allen Clement taught the graduate Distributed Systems course with Peter Druschel during Winter 2012/2013, a graduate seminar on Fault-tolerant distributed real-time systems with Björn Brandenburg during Summer 2013, and a graduate seminar on Operating System Design and Implementation with Björn Brandenburg during Winter 2013/2014, all at Saarland University.

Service. Internal to the institute Allen has served on the student recuriting committee for 2012/2013 and the faculty recruiting committee for 2013/2014.

In external service Allen has served (or is currently serving) on the program committees for HotDep 2013, OPODIS 2013, TRiOS 2013, LaDiS 2013, and EuroSys 2014. He was additionally the poster session co-chair for SOSP 2013.

3.2 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.

3.2.1 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 [75]. 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 [128]. 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.

3.2.2 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 [151]. 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.

3.2.3 Sybil attacks

In large open systems, the ability of a user to create multiple accounts is a significant challenge to robustness as ill-mannered users can unbalance the system. The last several years have seen a preponderence of papers developing new techniques for identifying sybil accounts based on the structure of an underlying social graph.

In conjunction with colleagues at Google, la Sapienza, and UT-Austin, the group has demonstrated a deep connection between the sybil detection literature and the theory of random walks [10].

3.2.4 Future and ongoing directions

As mentioned above, the group is actively pursuing further work on Byzantine fault tolerance and concurrency & geo-replication; further work on sybil attacks is, for the moment, not planned.

Additional research directions which the group is likely to explore in the next couple of years include:

1. **Replication: trading liveness for safety.** The key tenant of fault tolerant replication is that additional replicas can be used to ensure that the output of a system is correct. In many environments, a result that is late is no different from an incorrect result. In conjunction with Björn Brandenburg and the Real-time Systems Group, we plan to explore this tradeoff between safety and liveness.

2. **Workload robustness.** A next challenge to robust performance is changes in workload characteristics, specifically rapid and/or unexpected spikes in activity. In settings where server resources are fixed, the group is exploring techniques for deterministically slicing server resources and performing request triage to drop requests that are destined to fail their SLAs early before they consume valuable resources. In settings where server resources are flexbile, the group is exploring ways to adapt eventual consistency to be *elastic consistency*.

4 The Social Information Systems Group

4.1 Overview

The report covers the period from September 2012 – October 2013. (Cristian Danescu-Niculescu-Mizil joined MPI-SWS on October 1st 2013, but was affiliated with MPI-SWS throughout his year as a postdoctoral researcher at Stanford University.) The group's research interests are in social computing and natural language processing. At a high level, the group's research direction 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. Subhabrata Mukherjee will join the group as a doctoral student in Fall 2013.

Collaborations.

The group has an ongoing collaboration with the networked systems group (led by Krishna Gummadi). Externally, the group is collaborating with researchers at Stanford University, Cornell University, University of Bucharest, Facebook and Google.

Publications. The following papers were co-authored under joint MPI-SWS/Stanford affiliation and published at prime natural language processing and social computing venues:

- "No country for old members: user lifecycle and linguistic change in online communities", WWW 2013. Best paper award. [81]

- "A computational approach to politeness with application to social factors", ACL 2013. Nominated for the best paper award. [80],

- "Linguistic models for analyzing and detecting biased language", ACL 2013. [176]

- "Characterizing and Curating Conversation Threads: Expansion, Focus, Volume, Re-entry", WSDM 2013. [37]

Data. In order to encourage further research in these research areas, we released three datasets:

1. The largest corpus with politeness annotations to date (available at `http://www.mpi-sws.org/~cristian/Politeness.html`)

2. Complete longitudinal data for two online review communities (available at `http://www.mpi-sws.org/~cristian/Linguistic_change.html`)

3. Large scale conversational data from Wikipedia talk-pages (available at `http://www.mpi-sws.org/~cristian/Echoes_of_power.html`)

Press. Cristian's research on language and social computing has been featured in popular-media outlets such as the New Scientist, NBC's The Today Show, NPR and the New York Times. Cristian's opinion on external social computing research was solicited and quoted by ABC News.

Awards and invited talks. Cristian's work won the best paper award at WWW 2013 and was nominated for the best paper award at ACL 2013.

Cristian was invited to talk at the Information Science Departament Colloquium at Cornell University, at the Computational Social Science Conference at Stanford University, at the Chair of Systems Design at ETH Zürich, at the Department of Informatics Colloquium at the University of Zürich and at the Department of Computer Science Colloquium at the University of Mannheim.

Service.

During the last year Cristian co-organized the Workshop an Language Analysis in Social Media at NAACL 2013, served on the program commitee of ACL 2013 and reviewed for the Machine Learning Journal.

4.2 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.

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.

Linguistic style coordination in online communities My approach is rooted in the psycholinguistic theory of *linguistic style coordination* [50, 171, 178, 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 [77] 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 [79], 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.

Future directions. There are numerous additional directions suggested by the current results. One set of questions is to further understand the types of social distinctions that are manifested by differences in coordination; while we have seen that multiple forms of power can be observed this way, there could well be other properties that can be exposed as well. More broadly, my work makes clear that language use contains subtle features that reveal latent social information, and identifying new classes of such features promises to extend our understanding of how the social relations between individuals is embedded in their conversational behavior and our ability to predict these relations.

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 mechanisms underlying helpfulness evaluation of opinions. In joint work with G. Kossinets, J. Kleinberg and L. Lee [78], 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 [172], 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 [117] 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.

Future directions. There are a number of interesting directions for further research. First, although our results are generally robust, preliminary experiments also show some small but intriguing variations in the observed effects depending on the considered sub-population. An anecdotal example is the fact that when it comes to science books, reviewers from the state of New Jersey are in general considered the most helpful. Another example is that helpfulness evaluators from Japan tend to be more biased towards appreciating negative reviews. Such variations could be used to form hypotheses about differences in the collective behaviors of the underlying populations. Another idea that we are starting to investigate is to model the effect which polar social relations, such as those encoded in the trust and distrust networks available on Epinions.com, have on "usefulness" perception. From a practical point of view, it would also be interesting to consider mechanisms that can ameliorate some of the biases we observed, so that we can deploy such mechanisms in systems that enable the expression and dissemination of opinions.

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.

5 The Foundations of Programming Group

5.1 Overview

This report covers the period May 2011 – October 2013. The research of this group (formerly known as the Type Systems and Functional Programming Group) focuses on the design, semantics, verification and implementation of modern programming languages and systems. Major topics of study have included: advanced type systems for modular programming and verification; Kripke models and separation logics for reasoning about higher-order, imperative and concurrent programs; and compositional compiler certification.

Personnel. The group is led by **Derek Dreyer**, who joined the institute in January 2008 and received tenure in June 2013. It currently includes one postdoc (**Aaron Turon**) and four PhD students: **Georg Neis**, **Beta Ziliani**, **Scott Kilpatrick**, and **David Swasey** (who is co-advised by Deepak Garg). Aaron Turon received his PhD in 2013 from Northeastern University and joined us in January 2013. During the review period, the group included two other postdocs: **Chung-Kil (Gil) Hur** (who left in September 2012 for a postdoc at Microsoft Research Cambridge and is now an assistant professor at Seoul National University) and **Neel Krishnaswami** (who joined the group in September 2011 after a postdoc at Microsoft Research Cambridge and left in September 2013 for a position at the University of Birmingham). After Umut Acar left the institute, his postdoc **Joshua Dunfield** also joined our group for the period of Sept. 2012 – Sept. 2013.

Collaborations. During the review period, the group has engaged in successful collaborations with a number of leading researchers in Europe and the U.S., including: Lars Birkedal and Jacob Thamsborg (Aarhus University and ITU-Copenhagen), Amal Ahmed (Northeastern), Andreas Rossberg (Google), Claudio Russo (Microsoft Research Cambridge), Georges Gonthier (Microsoft Research Cambridge), Aleks Nanevski (IMDEA Software Institute), Lindsey Kuper and Ryan Newton (Indiana), Robert Harper (CMU), Simon Peyton Jones (Microsoft Research Cambridge), and Simon Marlow (Facebook). The group has collaborated (and continues to collaborate) actively with fellow MPI-SWS faculty Viktor Vafeiadis and Deepak Garg, as well as several (former) members of Umut Acar's group.

Publications. The group has published regularly in top conferences and journals like POPL, ICFP, and JFP. During the review period, group members have co-authored 17 conference papers (four POPL [123, 138, 125, 195], eight ICFP [109, 135, 140, 89, 213, 193, 136, 90], one LICS [124], one OOPSLA [116], one ESOP [137], one ITP [74], and one CSL [139]), as well as 7 journal articles (three JFP [170, 86, 110], two LMCS [85, 94], one TOPLAS [182], and one JAR [41]). Group members also have two papers accepted to the upcoming POPL 2014 conference [130, 145], as well as one journal article that has been conditionally accepted to JFP [183].

Software. In joint work with Viktor Vafeiadis, the group has released two software developments for the Coq proof assistant: Paco (a library for proofs by parameterized coinduction, described in our POPL'13 paper [125]), and Mtac (an extension of Coq for typed tactic programming, described in our ICFP'13 paper [213]). In March 2012, Krishnaswami also released the AdjS programming language for higher-order functional reactive programming. These are available open-source from the MPI-SWS PLV (programming languages/verification) group web page (`http://plv.mpi-sws.org`).

Teaching. In Summer 2011, Dreyer and Vafeiadis co-taught a graduate course on "Concurrent Program Logics". In Winter 2012-13, Dreyer taught a graduate course on "Parametricity and Modular Reasoning".

External funding. The research of the group has been 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 will be used to fund David Swasey.

Invited talks. Dreyer was an invited speaker at LFMTP 2011 and the 2012 Parametricity Workshop. Krishnaswami was an invited speaker at TLDI 2012 and DICE 2012.

Service. Internally to MPI-SWS, Dreyer is serving as coordinator for the preparation of this institute report, as well as a member of the Programming Staff Approval Committee. Since Feb. 2010, he has also served as the elected staff representative for MPI-SWS in the Chemistry, Physics & Technology (CPT) Section of the Max Planck Society, which entails participation in CPT Section meetings three times a year.

Externally, Dreyer founded the highly successful HOPE workshop (co-located with ICFP) and co-chaired the first two instances of that workshop (in 2012 and 2013). He will also serve as co-chair of the 2014 Coq Workshop (co-located with ITP). Beyond that, he has served as a member of the PC for CC 2012, ESOP 2013, CPP 2013, and ML 2013, as well as on the external review committees for POPL 2012 and POPL 2013. Last but not least, he was elected to the position of member-at-large on the ACM SIGPLAN Executive Committee (since July 2012), where he serves as "awards chair" (a significant responsibility, involving several weeks of work per year).

Hur served as PC member for the 2012 Coq Workshop and APLAS 2012. Krishnaswami served as PC member for POPL 2014. Turon was PC co-chair for LaME 2013, as well as PC member for POPL 2014 and PADL 2014.

5.2 Research agenda

During the review period, the focus of the group has broadened considerably to comprise a wide range of different projects, while remaining clearly guided by the group's overall theme of modular programming and verification. These projects concern such varied topics as concurrency, compiler verification, and interactive proof automation. Some of the highlights are:

- **CaReSL** [195, 193]: a modal separation logic for modular verification of *higher-order fine-grained concurrent data structures*

- **Parametric bisimulations** [123]: a groundbreaking new proof technique for *compositional compiler certification*

- **Backpack** [130]: *a new package system for Haskell*, retrofitting the language with support for separate modular development

- **Mtac** [213] and **Paco** [125]: *two different extensions to Coq* supporting more robust and compositional automation of interactive proofs

In the remainder of this section, we discuss the first three of these, as well as several directions for ongoing and future work. The work on Mtac and Paco is in collaboration with Viktor Vafeiadis. For space reasons, please see his section of the report (Section 11.2.2) for further details about it.

CaReSL [195, 193] In the interest of exploiting parallelism, modern concurrency libraries like `java.util.concurrent` support a variety of different *fine-grained concurrent data structures (FCDs)*, which rely on single-instruction synchronization mechanisms like compare-and-swap (CAS) rather than coarse-grained locking. The use of fine-grained synchronization comes at a cost, however: FCDs are very subtle and tricky to reason about, due to the complex interference between concurrently operating threads.

The canonical notion of correctness for FCDs is *linearizability* [119], which essentially says that there is a single point during the execution of any concurrent operation on the data structure—the "linearization point"—at which the operation can be viewed as taking effect atomically. However, there are several problems with linearizability as a correctness criterion. For one, it is not clear how to generalize linearizability (which is typically formalized as a trace property) to the setting of higher-order languages and libraries like `java.util.concurrent`. Moreover, it is not clear what linearizability offers in terms of useful guarantees for the *clients* of FCDs.

An arguably more general and useful criterion is *contextual refinement*, which holds if an FCD is observationally indistinguishable from a simple coarse-grained (lock-based) implementation of the same data structure. The notion of contextual refinement is readily applicable to the higher-order setting. What's more, if one can establish that an FCD refines its coarse-grained counterpart, then clients of the FCD can pretend for verification purposes that they are working with the coarse-grained version (which is comparatively easy to reason about), while at the same time reaping the efficiency benefits of the FCD.

Toward this end, we have developed a new logic called CaReSL [195, 193]. CaReSL is the first logic to support direct proofs of contextual refinement for sophisticated FCDs like the Michael-Scott queue, elimination stacks, conditional compare-and-swap, flat combining, and more, all in a higher-order, Java-like setting. Furthermore, it supports proofs that are thread-local, spatially-local, and temporally-local, even though some of the more advanced FCDs seem almost purposefully designed to stymie such reasoning.

CaReSL is a modal separation logic, which builds closely on the Kripke logical relations (KLR) models and logics for ML-like languages that we developed in previous work [87, 86]. However, CaReSL extends our previous work in nontrivial ways to account for concurrency and fine-grained synchronization. One key innovation in CaReSL is the idea of establishing *local protocols*, which describe the "life story" of each individual piece of an FCD (*e.g.,* each node of a linked list): how it came into being, how its role in the FCD evolves over time, and how it eventually "dies" by being disconnected from the FCD. In addition to offering intuition about how an FCD works, these local protocols provide a unified account of resource invariants and rely-guarantee reasoning. Another key feature of CaReSL is that contextual refinement—a *relational* property—is, perhaps surprisingly, expressible as just a particular mode of use of traditional Hoare triples. As a result, *both* refinement proofs for FCDs *and* Hoare-style proofs about client code may be expressed and combined in a single logical framework.

Parametric bisimulations [123] One of the “grand challenges” in program verification is the problem of compiler certification. Although there has been great progress on this problem (most notably in Xavier Leroy’s CompCert project [149]), existing certified compilers rely on *simulation* proofs, which are closely tied to the implementation of the particular compiler in question, and thus only support the scenario where all modules in a program are compiled using the *same* compiler. They leave open the problem of how to verify the correctness of assembly code that is hand-optimized or linked together from the output of different certified compilers. Solving this problem is crucial in order to ensure the scalability and compositionality of certified compilation, since real programming systems are frequently linked together from modular components, each of which may be built using different compilers and different languages.

Supporting such *compositional compiler certification*, however, is quite challenging. In particular, it seems to demand both *horizontal* compositionality (the ability to link together the results of separately certified compilations) and *vertical* compositionality (certifying a multi-phase compiler by certifying each of its constituent phases independently and then composing those certifications transitively).

In our initial work on this topic from the prior review period [122], we developed a Kripke logical relations (KLR) model supporting horizontally compositional certification of a one-pass compiler from ML to assembly. However, due to limitations of KLRs (namely, the lack of transitivity), that approach could not scale to support vertical compositionality. In contrast, traditional simulation methods like those used in CompCert naturally support vertical compositionality but not horizontal compositionality.

In joint work with Viktor Vafeiadis [123], we have since developed a new and groundbreaking proof technique, called *parametric bisimulations (PBs)* (aka “relation transition systems”), which brings together the advantages of KLRs and simulations. PBs are much like traditional bisimulations, except that they are parameterized explicitly over a relation representing the equivalences that hold in “the rest of the program”, and this extra parameterization facilitates both horizontal and vertical compositionality at once.

Our initial paper on PBs only studies them in the setting of a single language. For his thesis project, funded by a Google doctoral fellowship, Georg Neis is currently exploring the generalization of PBs to the interlanguage setting and the application to compositional compiler certification.

Backpack [130] In our previous work on module system design, we explored a variety of extensions to the ML module system. However, it must be said that even vanilla ML modules support a relatively strong form of modularity—what we call *separate modular development (SMD)*—in which explicit interfaces express assumptions about dependencies, and each module in a program can be typechecked and reasoned about independently. In contrast, module systems like that of Haskell permit only a weak form of modularity—what we call *incremental modular development (IMD)*—in which module implementations depend directly on other implementations and must be processed in dependency order. Tools like the Cabal package management system pick up the slack, enabling Haskell programmers to (ab)use version-range dependencies in order to work around the lack of interfaces. But the Haskell community recognizes that this is a makeshift solution and is actively seeking ways to support SMD properly.

In this project, which is joint work with Simon Peyton Jones and Simon Marlow (lead developers of GHC, the leading Haskell compiler), we develop Backpack [130], 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 [182] that we developed in earlier work, but differs significantly in detail. Like MixML, Backpack supports explicit interfaces and recursive linking. Unlike MixML, Backpack supports a more flexible *applicative* semantics of instantiation. Moreover, its design is motivated less by foundational concerns and more by the *practical* concern of integration into Haskell, which has led us to advocate simplicity—in both the syntax and semantics of Backpack—over raw expressive power. In particular, whereas the semantics of MixML was defined by elaboration to a rather specialized linear type system, the semantics of Backpack packages is defined by elaboration to ordinary Haskell, thus showing how Backpack maintains interoperability with Haskell while extending it with separate typechecking.

Our initial paper on the design of Backpack will serve as the foundation of Scott Kilpatrick's PhD thesis. However, it is only the first step. We are currently investigating how to incorporate support for Haskell's type class mechanism into Backpack, and in future work, we plan to explore how to integrate Backpack into the existing Cabal package management system.

Ongoing and future work

Aside from the future research directions mentioned above, the primary focus of our ongoing and future research is on adapting our CaReSL logic to a variety of different language settings and applications.

Verifying a library for scalable joins Fournet and Gonthier's *join calculus* [96] provides a powerful combination of message passing and *declarative atomicity*, by which processes can receive messages from multiple channels simultaneously. Prior work has demonstrated the expressive power of joins as a high-level concurrency paradigm [40], and recently, Turon and Russo [194] also showed that they can be *implemented* efficiently through fine-grained data structures like "lock-free bags".

In this project, funded by a grant from Microsoft Research, we plan to verify a realistic software stack for scalable joins-based programming. Specifically, we will develop a unifying framework in which to verify both Turon and Russo's scalable joins library and a representative suite of joins-based client code that depends on it. Our hypothesis is that a feasible and effective way to realize such a verification effort is via the method of contextual refinement supported by CaReSL.

Separation logic for security protocols There has been a great deal of work on compositional verification of security protocols, in the form of logics, tools, and type systems like PCL [84], ProVerif [49], and F7 [44]. However, existing work has focused primarily on proving *intrinsic* properties about what happens *during* security protocols (*e.g.,* correspondence assertions), while ignoring *extrinsic* properties describing how a protocol is *useful* to its clients. Such extrinsic properties are essential if one wants to verify the functional correctness of client programs that depend on these protocols.

In ongoing joint work with Deepak Garg and Bob Harper, we are developing a logic, based on CaReSL, for a message-passing language with cryptographic primitives. The "key" idea, so to speak, is to treat secret keys as a resource that can be owned (or shared) by different protocol participants, and to use the Kripke model of CaReSL to express invariants on how the messages signed by different participants may evolve over time. This is, to our knowledge, the first application of separation logic to security protocols.

GPS: a logic for navigating weak memory CaReSL, like essentially all existing concurrent separation logics, assumes a sequentially consistent memory model. However, modern architectures implement weaker memory models, which allow threads to see writes in different orders, and real programs cope with the semantic complexity of weak memory in order to reap gains in performance. In joint work with Viktor Vafeiadis, we are developing GPS, the first logic for direct, local, rely-guarantee reasoning about concurrent programs under weak memory (C11) semantics. GPS is inspired by Viktor's recent work on RSL [199], but goes beyond it to support ghost state and CaReSL-style protocols, as needed for more complex programs.

6 The Distributed Systems Group

6.1 Overview

The report covers the period from May 2011 – October 2013. The group's research during this period has focused mainly on two areas: 1) Enforcing data confidentiality, integrity and accounting policies in distributed data processing environments (in collaboration with Deepak Garg's group); 2) private communication for mobile social applications (in collaboration with Bobby Bhattacharjee and Elaine Shi at the University of Maryland). The group has also collaborated with Krishna Gummadi's group on privacy exposure in social networks, Paul Francis' group on traffic-analysis resistant anonymity networks, and has continued its collaboration with the University of Pennsylvania, Duke University and Akamai on hybrid CDNs. Recently, we've started a collaboration with Bryan Ford and his group at Yale, on online identity management.

Personnel. The group is led by Peter Druschel and currently has six graduate students (Paarijaat Aditya, Eslam Elnikety, Viktor Erdelyi, Aastha Mehta, Nick Merrit, Anjo Vahldiek). Peter is co-advising a post-doc with Paul Francis (Stevens Le Blond). Prof. Lorenzo Alvisi (University of Texas at Austin) was visiting the group (and the institute) in the Fall of 2012 and in the summer of 2013; he was supported by a Humboldt Research Award. William Caldwell works for the group as a research support engineer on a freelance basis.

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, and Akamai.

Publications. Group members have co-authored papers that appeared in ACM SIGCOMM, Usenix NSDI, ACM/Usenix IMC, ACM CoNEXT and COMSNETS [212, 146, 168, 8, 202], as well as a technical report [115]. Druschel co-authored two white papers for the European Network and Information Security Agency (ENISA) [58, 88]. Three papers are under submission [201, 148, 7].

Teaching. Peter Druschel taught a core course on Operating Systems in 2011 (jointly with Rodrigo Rodrigues) and 2013 (jointly with Bjoern Brandenburg) at Saarland University, and taught a core course on Distributed Systems in 2012 (jointly with Allen Clement) 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 proposal, which has made it to the final round of competition, with a reverse site visit taking place in Brussels on Nov 5, 2013.

Awards and invited talks. Druschel won the 2011 ACM/IFIP/Usenix Middleware 10-year Best Paper Award, together with Ant Rowstron (Microsoft Research, Cambridge). He gave invited/keynote talks at the ACM SIGOPS Asia-Pacific Workshop on Systems (APSys) and the 12th ACM/-IFIP/Usenix Middleware Conference in 2011, and the 12th IEEE International Conference on Peer-to-Peer Computing (P2P) and the 28th IEEE International Conference on Data Enginneering (ICDE) in 2012.

Service. Within the MPS, Peter Druschel has served on the strategy committee (Perspektivenkommission) of the Chemical, Physical, and Technology section, the Committee on Information Technology (BAR), and the selection panel of the joint Fraunhofer/Max Planck research program.

Peter Druschel has served on the Technical Advisory Boards of Microsoft Research, Cambridge since 2011 and Microsoft Research, India, since 2013. He has served on the scientific committee of the Laboratory on Information, Networking and Communication Sciences (LINCS), Paris, and on the steering committee of the ACM SIGOPS Asia-Pacific Workshop on Systems (APSys) since 2009, and on the steering committee of the EuroSys/INRIA Winter School on Hot Topics in Distributed Systems (HTDC) since 2008. He served on the SIGOPS EuroSys steering committee through 2012.

Peter serves on the editorial board of the ACM Communications of the ACM (CACM), he chaired the programm committee of the ACM Symposium on Operating Systems Principles (SOSP) in 2011, and he co-chairs the

PC of EuroSys 2014. He also chaired the selection committee for the ACM SIGOPS Hall of Fame Award in 2011 and 2012. Peter also served on the program committees of IEEE ICNP 2012 and IEEE/Usenix HotOS 2013.

6.2 Research agenda

During the reporting period, the group's work has focused on providing security, privacy and accountability in systems at scale. The group's main current focus is on two projects decribed below.

Additionally, we have collaborated with Paul Francis' group, as well as researchers at UW and MSR, on traffic-analysis resistant anonymous networking [146], and with Krishna Gummadi's group and researchers at Northeastern on defending against large-scale crawling in online social networks [168] and on Sybil-resistant system design [202]. We have also wrapped up our joint work with researchers at UPenn, Duke, and Akamai on hybrid content distribution networks [212, 8].

6.2.1 Enforcing declarative data policies

In this project, we study methods to enforce declarative data integrity and confidentiality policies in distributed data processing environments. This work is done in collaboration with Deepak Garg's and Rodrigo Rodrigues' group, as well as researchers at Cornell and Google.

Guardat: A foundation for policy-protected persistent data In [201], we present Guardat, an architecture that enforces rich data access policies at the storage layer. Users, application developers and system administrators can provide per-object policies to Guardat. Guardat enforces these policies and provides attestations about the state of stored objects. With Guardat, the data integrity, confidentiality and access accounting rules for a collection of objects can be stated as a single declarative policy. Policy enforcement relies only on the integrity of the Guardat controller and any external policy dependencies; it does not depend on correct software, configuration and operator actions in other parts of a system. Guardat allows developers, system administrators and third-party hosting platform providers to enforce concise, system-wide data protection policies based on a small trusted computing base, and to demonstrate their compliance to any party that trusts the Guardat layer. We present a design and prototype implementation of Guardat, show experimentally that the overhead of making policy checks and storing additional metadata are low, and discuss applications and policies.

Current work: In current work, we are desiging *Thoth*, a system that extends Guardat by mediating not only storage accesses, but all input and output of tasks in a distributed, parallel computation. Therefore, Thoth can enforce data confidentiality and integrity policies in a complex data retrieval system like a search engine. For this purpose, we are extending the policy language to support provenance and declassification policies, which control the upstream and downstream data flows, respectivelty. The policy interpreter and enforcement logic will be implemented as part of the Xen virtual machine monitor. A poster on Thoth appeared at SOSP 2013 [91].

6.2.2 Private, secure communication for mobile social apps

In this project, we study private and secure communication support for mobile social applications. 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 radios to discover and share content with nearby users. The discovery protocol to locate other user's devices used by these applications must be extremely power efficient since they are run continuously in the background. A good discovery protocol must also preserve user privacy (users cannot be tracked by third-parties), while providing linkability (users can recognize friends when strangers cannot) and efficient silent revocation (users can unfriend without permission and without rekeying their entire friend set).

In [148], we introduce 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. 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 [7], we introduce 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 communication and sharing, we evaluated EnCore's utility using a live testbed deployment with 35 users.

Current work: In ongoing work, we are developing versions of EnCore and *Context* for both iOS and Android that don't require rooting the phone. We are also working on a fully functional data sharing backend service independent of Facebook, and on giving the *Context* user interface more professional polish. Once this is done, we plan to make the system and application available for download and hope to attract a larger user base, which would enable us to perform an extensive evaluation of EnCore's capabilities.

We are also working on group communication support that exploits the graph of secure encounters. Using such support, it is possible to communicate securely and anonymously with participants of a large event (like a concert or sports match), the people who might have contracted a disease, directly or indirectly, from a particular person, or the set of people who might have found a lost item.

Finally, we are working on an audio challenge/response protocol that can verify that two devices are in the same room, and can provide more accurate distance and relative positioning information than what is possible using Bluetooth signal strength. Also, we hope to engage with Google to try and get support for SDDR/EnCore included in the mainline Android distribution.

7 The Large Scale Internet Systems Group

7.1 Overview

The report covers the period from May 2011 – October 2013. 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 user data gathering, analytics and advertising, while protecting user privacy. The group is also working on network anonymity systems. We have launched a spinoff called Aircloak (www.aircloak.com), which works closely with the group.

Personnel. The group is led by Paul Francis. It currently has three graduate students (Alexey Reznichenko, Istemi Ekin Akkus, and Muhammad Imran Khan) and one postdoc (Stevens Le Blond). The group is supported by three research support staff who are also members of the spinoff (Sebastian Probst Eide, Felix Bauer, Matthias Kretschmer), and by a fourth developer working on a freelance basis, Jeff Hoye.

Two postdocs left the group during this reporting period: Ruichuan Chen, who is now with Bell Labs, and Bin Cheng, now with NEC Research.

Collaborations. The group has joint publications with the distributed systems group (led by Peter Druschel). Externally, the group has collaborated with researchers at Cornell (one of whom moved to Twitter during the collaboration), Microsoft Research, the University of Washington, AT&T Research, and the Technical University Kaiserslautern.

Publications. During the reporting paper, the group has published papers in SIGCOMM [66, 146], CCS [9, 179], NSDI [67], and CoNEXT [196, 68].

Invited Talks. Prof. Francis gave the keynote at ACM CoNEXT 2011, and gave invited talks at Technicolor Research, the WiTAP Workshop (Stanford), MSR India, UC Berkeley, the INTIMATE Workshop in Paris, the Dagstuhl Workshop on Decentralizing Systems for Privacy, ENISA, MSR Redmond, and the University of Washington.

Patents and Technology Transfer. During this time period, the group has submitted seven patent applications, written a proposal for privacy standards in behavioral advertising, written a white paper describing our private analytics technology, and submitted reports to two privacy certification organizations, TViT (Germany) and EuroPriSe (EU). In addition, Prof. Francis visited a number of privacy policy and advocacy groups, including ENISA, Privacy International, FTC, and EFF.

External funding. The group was awarded an EXIST startup grant from the German Government (Bundesministerium fr Wirtschaft und Technologie) for roughly EUR500K over 18 months.

Service Prof. Francis served on the following PCs:

- ACM Hotnets 2011 Program Committee
- USENIX NSDI 2011 Program Committee
- WiTap 2011 (Workshop on Internet Tracking, Advertising, and Privacy), Program Co-chair
- ACM CONEXT SWID 2011 (Special Workshop on the Internet and Disasters), Program Co-chair
- CCS 2011 Program Committee
- APF 2012 (Annual Privacy Forum) Program Committee

Awards and media Prof. Francis was awarded the SIGCOMM Test of Time award for 2011 for the paper "A Scalable Content-Addressable Network," along with Mark Handley, Richard Karp, Sylvia Ratnasamy, and Scott Shenker.

Stevens Le Blond's work on security flaws in Skype and other peer-to-peer applications received global media attention, for instance in the WSJ, Le Monde (French), die Zeit (German), Daily Mail, Slashdot, Wired, and New Scientist.

7.2 Research agenda

For the past 4.5 years, my group has been focused on the problem of online privacy. Although people generally think of online privacy as being a problem for individuals, in fact it is also a problem for business. For individuals, the problem is the risk of harm inherent in user data being collected, stored, and often, exchanged. As companies learn more and more about users, and get better and better at profiling, users increasingly find themselves in a situation where the companies they interact with know a great deal about them. In some cases, this may benefit the user, for instance as with recommendation services. In other cases, however, this knowledge may be exploited to the user's detriment.

As we explore the issue of online privacy, we increasingly find that it is a problem for industry as well. Companies can benefit greatly from aggregate user data analysis, including user data held by other companies, but are reluctant to share that data because of concerns over their users' privacy.

From the beginning, we've taken a practical approach to the research. We are focused on solutions that can succeed in practice. For us, this ideally means designing solutions that do not hurt industry, We initially focused on privacy-preserving behavioral advertising, on the premise that advertising motivates a large amount of user data collection. Towards this end, we designed, built, and are currently deploying a privacy-reserving behavioral advertising system called Privad (published before this evaluation period [111]).

Privacy-preserving Aggregate Analytics The last two years, however, we have focused on privacy-preserving aggregate (statistical) analytics. This is in part because aggregate analytics is needed by Privad, but mostly because analytics is stand-alone an important application. In the last two years we have designed four privacy-preserving analytics systems, each improving on the last. The first three of these were published in NSDI (PDDP [67]), CCS (Akkus [9]), and SIGCOMM (SplitX [66]) respectively. Each of these is based on a similar architecture: user data is held at user devices (clients). Honest-but-curious analysts transmit queries to clients, and receive answers via honest-but-curious proxies. The proxies blindly add differentially-private noise to the answers.

A core design component in each case is blindly adding noise at the proxies. The answers must be encrypted so that the proxies cannot deduce information about clients. The proxies must add noise without knowing how much noise they added. This is because if the noise answer is publicly published, the proxy would otherwise be able to reverse the pre-noisy answers by subtracting the noise it added.

The first system, PDDP led by postdoc Ruichuan Chen, accomplished this by exploiting the XOR homomorphism of the Goldwasser-Micali bit-crypto system. Clients provide a pool of random coins to the proxy, which "re-flips" them by XOR-ing them with their own random values. These coins are then mixed in with the real client answers before forwarding them to the analyst. Because the proxy does not know the original value of each coin, it does not know how much noise was added to the answers.

PDDP was designed in support of Privad, and indeed has been deployed along with Privad. As an adjunct to Privad, PDDP piggy-backed query distribution on top of Privad's ad distribution. PDDP, however, does not scale well as a stand-alone analytics system. The primary problem is that PDDP requires all queries from all analysts to be transmitted to all clients. In, for instance, a web analytics application, every website can be a potential analyst. The second scaling problem with PDDP is that it scales linearly with the number of possible answers to a query. For the query "count the number of clients that visit each website", this can requires millions of possible answers.

The second system, led by student Ekin Akkus, addresses the shortcomings of PDDP specifically for the web context. In particular, Akkus' system exploits normal client access to web servers to distribute queries: the query is retrieved from the website when the client otherwise accesses the website. Akkus' system also eliminates the need for a separate administrative entity to run the proxies. Instead, it uses the websites themselves for this purpose. Instead of blind XOR homomorphism, Akkus uses double noise addition in series. The web server adds noisy answers (non-blind) to encrypted user answers. The analytics service provider decrypts and tallies the real and noisy answers, adds its own noise, and gives the double-noisy answer to the web server. The web server then subtracts its noise to get the single-noisy answer. Only the double-noisy answer may be made public. Akkus uses standard RSA encryption.

Unlike the first two systems, the third system, SplitX, is a fully general, stand-alone privacy-preserving analytics system. It also scales three order of magnitude better in both message size and computation, than the previous systems. The breakthrough here is the use of simple XOR (not homomorphic) to encrypt user answers and noisy answers. SplitX requires two administratively distinct honest-but-curious proxies. User answers are XOR'd with a random string. The random string is sent to one proxy, and the XOR'd string to the other. The proxies cooperate to generate blind noisy answers, with each proxy providing half of each answer. The real and noisy answers are mixed and forwarded to the analyst, which can decrypt them simply by XOR'ing the strings. Queries are also distributed using simple XOR: clients subscribe to queries from specific analysts via the pair of proxies, and the queries are published via the proxies.

With the design of SplitX, we felt we had a system that scales adquately well that we could base an analytics startup on SplitX technology. A particulary nice advantage, from a research point of view, of doing a privacy spinoff is that the technology should be transparant in order to build trust. As a result, virtually all of the innovation done by the startup can be openly published. I started looking for a team in early 2012, and by August 2012 the first co-founder joined (Sebastian Probst Eide, a PhD studnet from Cambridge).

Interestingly, within 2 months of beginning work, we abandoned the client-based distributed architecture for a totally new approach. The change stemmed from a concrete examination of what kind of company should deploy the honest-but-curious proxies, and the associated trust and cost implications. We found ourselves exploring ways to make these proxies as transparent as possible, and in the end decided on trusted computing and remote attestation as a basis for establishing trust. Once this decision was made, we found that we could dispense with the proxies altogether, and instead store data in hardened and attested servers.

We call this approach *cloaked computing*. Trusted computing (the TPM) is used in two ways: clients can remotely attest the servers (cloaks) as proof of what software is running, and user data, when stored on disk, is sealed to the software running on the cloak. This software is hardened with SELinux such that nobody, including system operators or designers, have access to data on the system. Answers to queries are anonymized using a combination of differential-privacy like noise adding and other filters that make it difficult to infer user data through repeated queries. We are currently working on a publication, but in the meantime an overview of the approach can be found at [1].

Deploying Privad and PDDP PhD student Alexey Reznichenko is working on the implementation and large-scale deployment of Privad and PDDP. The purpose is to test the research ideas in Privad and PDDP in a setting that is as close to reality as we can get. Alexey has implemented Privad and, along with Ruichuan Chen, PDDP, for Mozilla. It does real targeting based on user search terms on shopping sites, and delivers real ads by pulling them from online shopping APIs. We get users by bundling with existing Firefox addons. One challenge is that we need about 30 or more installations to get one useful user. This is because opt-in rate is about

1/20, and some fraction of users who opt-in use ad-block. Never-the-less, a deployment of many thousands of useful users appears realistic, and we'll be gathering data in the coming months.

Network Anonymity Postdoc Stevens Le Blond has been working independently on network anonymity. He has designed and built an onion-routing system that is resistant against traffic analysis attacks as well as node compromise, and yet exhibits good performance. By contrast, Tor is only resistant against the latter. The challenge is to keep packet latency low while still mixing packets sufficiently to defeat traffic analysis. One innovation in Steven's system is to spread traffic over many onion nodes in parallel so as to maximize mixing. Another innovation is to synchronize the beginning and end of sessions among a large group of users, adding chaff to make all sessions appear to be the same length. Stevens simulated the system for the BitTorrent application using real packet traces, and found only minor degradation in performance due to the network anonymization. The paper is published in SIGCOMM'13 ([146]).

8 The Foundations of Computer Security Group

8.1 Overview

The Foundations of Computer Security Group started in October 2011. This reports covers the group's activities up to October 2013. The group is broadly interested in solving problems in formal logic and programming languages that have application to security and privacy of computer systems. During the reporting period, the group's internal activities have focused on declarative access control, efficient information flow control in Web browsers, enforcement of privacy policies, program tracing, and separation logic for verifying security protocols.

Personnel. The group is led by Deepak Garg and currently has three graduate students — Ezgi Cicek (since October 2012), Vineet Rajani (since June 2012) and Paul (David) Swasey (since August 2012). Before joining this group, Ezgi was a student in the former programming languages and systems group (led by Umut Acar) for one year. Roland Perera was a postdoc in the group from October 2012 to May 2013. Roland moved on to a postdoc position at the University of Edinburgh.

Collaborations. Internally, the group has collaborated with the distributed systems group (led by Peter Druschel), the type systems and functional programming group (led by Derek Dreyer) and the former programming languages and systems group (led by Umut Acar).

Externally, the group collaborated with researchers at the Saarland University, Carnegie Mellon University, University of Illinois at Urbana-Champaign, University of Luxembourg, University of Torino, University of Helsinki and the IMDEA Software Institute during the reporting period.

Publications. During the reporting period, the group's direct activities have resulted in five conference papers at CCS [101], ICISS [83], LICS [100], CSF [103] and ICFP [140], two journal papers in ACM TOPLAS [169] and the Journal of Computer Security (JCS) [102], and one workshop paper [104]. Working with external collaborators, Vineet co-authored two papers, one each at SERVICES [175] and SCC [142]. Similarly, Dave co-authored a paper published at ESORICS [97].

Teaching. Deepak taught a graduate course on introductory Proof Theory in the summer of 2012. He was also invited to teach a course in authorization at the DFG/RS3 summer school in 2013 for graduate students working in security.

External funding. The group's research on security in Web browsers has been 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 the Saarland University. The funding lasts from October 2012 to October 2014.

Invited talks. Deepak was an invited speaker at the TaPP 2012 Workshop (TaPP = Theory and Practice of Provenance) and the TAFC 2013 Workshop (TAFC = Trustworthiness, Accountability and Forensics in the Cloud).

Service. Internally to MPI-SWS, Deepak served on the student recruiting committee in 2011–2012 and the faculty hiring committee in 2011–2012 and 2012–2013. Externally, Deepak served on the program committees of CSF 2012, APLAS 2012, ASIACCS 2013, POST 2013, Oakland 2013, FSTTCS 2013 and several workshops. Deepak has been the publications chair of CSF since its 2012 edition. He will also serve on the program committees of CSF 2014, Oakland 2014, SACMAT 2014 and POPL 2015.

8.2 Research agenda

The foundation of computer security (FCS) group works on logic- and language-based techniques for solving computer security and privacy problems. This covers both theory and practice. The group's research activities in the reporting period span the following broad areas: declarative access control, information flow control in Web browsers, enforcement of privacy policies, program tracing and separation logics for proving protocol correctness. The first four of these are described below. For the work on separation logics, please see Derek Dreyer's section of the report.

Declarative access control. Access policies are often represented in low-level configurations (access control lists, firewall tables) that state what access is allowed. Alternatively, policies can be expressed in declarative, abstract languages that record why an access is allowed. The latter simplifies policy specification, reduces the access control system's attack surface and makes complex access policies easier to debug. As a result, declarative access policies are a relevant (and active) area of research. Within the reporting period, the FCS group has worked on significantly advancing both the theory and the practice of declarative access control.

Theory. Deepak collaborated with researchers at the universities of Helsinki, Luxembourg and Torino to advance the study of access control logics. An access control logic is a domain-specific logic for representing and enforcing access policies. At their core, access control logics are multi-modal logics, with a modality "Alice says X" for each principal Alice. Access control logics are both interesting and complex because interaction between modalities in the logic's proof system must be carefully controlled within a range — with less interaction, we may deny required access (a usability issue) and with excessive interaction we may allow undesirable access (a security issue).

Our first contribution is a decision procedure that uniformly covers a broad range of modal interactions [100]. The decision procedure, based on so-called labeled sequent calculi, is similar to well-known tableaux procedures but is more powerful in its explicit use of symbolic states and general termination conditions. The procedure is also constructive — it provides countermodels when a formula is unprovable. Countermodels are relevant in access control to justify why an access is denied. An extension of this basic framework [103] covers access rights delegation and demonstrates how countermodels generated by the decision procedure can be used to find necessary credentials when access is denied. The latter is the first procedure of its kind for arbitrary access policies (prior work is limited to policies without left-nested implications). A further extension generalizes the decision procedure was to top-down or goal-directed search. This improves efficiency in practice [104].

Practice. Deepak has an ongoing collaboration with the Distributed Systems Group led by Peter Druschel. The goal is to push declarative access policy specifications down to the lowest layers of the software stack, thus providing a safety net against misconfigurations, bugs, vulnerabilities, operator errors and attacks in higher layers of software. A recently built system called *Guardat* enforces high-level, declarative access policies on persistent data, entirely in the storage layer. Guardat policies are rich: they may depend on authentication information, externally certified facts like wall-clock time, and file state (e.g., the contents of a file's first 2 bytes). Some examples of policies expressible in Guardat are: a signed-update only policy for system binary files, an append-only policy for log files, a mandatory log policy for access to medical records (without a special policy primitive for logging), and a time capsule to prevent backup data deletion or modification until a stipulated point of time. Guardat's policy language is expressive but easily interpreted with very moderate overheads. Guardat's design minimally adds high-level primitives to the storage layer to enforce practical policies. For example, our storage layer exposes a flat namespace of objects (files) and authenticates sessions. Policy-irrelevant primitives like disk-space allocation and directory hierarchies are delegated to an untrusted file system. Thus Guardat delicately balances minimality and efficiency with expressiveness.

In recently started work, we are extending Guardat's design to distributed data-retrieval systems, e.g., large-scale search engines. Due to bugs or misconfigurations, a search engine may leak a user's private emails (which are indexed in the same data structure as public content) during another user's search. Our system, *Thoth*, seeks to prevent such illicit information flows. A declarative policy attached to the private email specifies legitimate flows of information and a policy enforcement logic that encapsulates each individual computation task enforces the policy. End-to-end enforcement combines coarse-grain information flow tracking at the level of tasks with policy comparison algorithms, a novel combination that we have recently proved adequate.

Information flow-control in Web browsers. Modern Web applications often combine JavaScript code from several domains in a single trusted page. Some code may be untrusted or semi-trusted (advertisements or third-party libraries respectively) but the browser's standard protection mechanisms offer no defense when such code is malicious. An effective defense against many data confidentiality and integrity attacks is information flow tracking within the browser and, centrally, within JavaScript (JS). This is difficult and expensive because JS is very dynamic. Existing approaches either build taint-tracking interpreters for JS from scratch or employ multiple executions (for different security levels). In either case, overheads are prohibitive.

In a project led by graduate student Vineet Rajani, executed in collaboration with researchers at Saarland University, and funded by the DFG, we are designing and implementing a practical information flow tracking approach for JS: We are working in the *bytecode* interpreter. Specifically, we have instrumented WebKit, the JS engine used by the Safari Web browser. 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. Our overheads are moderate and improve on existing work by an order of magnitude. We incorporate a novel technique to correctly and efficiently track information flows across exceptions.

We have also built a complete formal model of the instrumented interpreter (through a thorough examination of nearly 20,000 lines of interpreter code) and proved that our information flow analysis is sound.

The overall scope of this project extends to information flow control in all browser components that contain sensitive data, not just the JS interpreter. Currently, we are adding taint-tracking to the document object model (DOM) — the browser state that persists across invocations of the JS interpreter. Our next step will be to consider other persistent state like bookmarks and history. Simultaneously, we are also looking into a suitable language for specifying information flow and declassification policies. Eventually, we envision a browser where users can annotate their data like passwords or history items with policies that are automatically enforced by tracking information flow in all browser components.

Enforcement of legal privacy mandates. As sensitive user data in the hands of companies and governments grows, so do legal mandates to protect its privacy. To maintain compliance with such mandates, there is need for computer assistants that find violations of a given privacy policy by observing day-to-day operations of data handling and transmission within an organization and comparing them to the policy. Prior work has established that legal privacy policies are effectively represented in extensions of linear temporal logic (LTL) and that finding violations amounts to model checking the policies on system logs. However, there are both theoretical and practical challenges. For example, policies require quantification over infinite domains and, in general, model checking with such quantification is undecidable. Moreover system logs may be incomplete, often due to failure to integrate information from different databases in time for audit.

Working with collaborators from Carnegie Mellon University, we have designed two different approaches to address these problems. First, in work published at the start of the reporting period (but completed earlier) [101, 83], an *offline* algorithm for auditing privacy violations was presented. The algorithm iterates over time to compensate for log incompleteness (converging to all violations if the log converges) and handles infinite quantification through a nontrivial application of mode-analysis, a static policy check often used in logic programming. In recent work, an alternate, *online* algorithm has been designed. By monitoring events online, the algorithm can summarize relevant log information in caches and discard the rest immediately, thus saving space and, in some cases, time. A novel, time-sensitive mode analysis that guides what to cache is the work's key contribution. In ongoing work, we are examining the space-time trade-off in moving from batch offline processing to online processing. In collaboration with researchers from University of Illinois at Urbana Champaign, we are also working on graded encryption techniques to protect the event log itself and to make the audit algorithms sensitive to such encryption.

Trace slicing and trace continuity. A broad but somewhat preliminary direction of research within the FCS group is study of program traces and their applications (to security problems and others). During the reporting period, we have pursued two such problems — slicing and continuity.

Slicing. A key problem in program debugging as well assigning blame for security violations is that of slicing — finding a minimal program part that explains a particular output. Motivated by this problem, postdoc Roland Perera developed a framework for slicing any higher-order concurrent computation. The key contribution is a novel characterization of minimal explanatory program slices through Galois connections and an algorithm for their computation through backward propagation of information through program execution traces.

Trace-continuity. Led by graduate student Ezgi Cicek and in collaboration with Umut Acar — now at Carnegie Mellon University but earlier at MPI-SWS — we have been working on type systems to establish what we call *trace-continuity* of programs. Roughly, trace-continuity means that if a program is run twice with slightly different inputs, the traces of the two executions will differ only slightly. This has applications in incremental computation (where programs are re-executed periodically with updated inputs): We have recently built a type system for a list processing language that allows a programmer to estimate quantitative upper-bounds on incremental execution time for her program. In future work, we plan to generalize this idea to a higher-order programming language.

9 The Networked Systems Group

9.1 Overview

This section describes the activities of the Networked Systems group between May 2011 and October 2013. The group's research interests are in the 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 protecting the privacy of users sharing personal data, understanding and leveraging word-of-mouth exchanges to spread information virally, and finding relevant and trustworthy sources of information in crowds.

Personnel. The group is led by Krishna Gummadi. It is currently comprised of five graduate students (Bimal Viswanath since November 2008, Mainack Mondal from October 2010, Juhi Kulshretha from April 2011, Muhammad Bilal Zafar from October 2012, Giridhari Venkatadri from October 2013) and one postdoctoral researcher (Przemyslaw Grabowicz from October 2013).

Massimiliano Marcon graduated with a PhD in November 2011. He took a position as a senior software engineer at SAP. Nuno Santos, who is graduating with a PhD student from the Dependable Systems group, was co-advised by Krishna Gummadi and Rodrigo Rodrigues. Nuno will be joining IST Portugal as an assistant professor starting November 2013.

Collaborations. The group members have close collaborations with the distributed, dependable, and social information systems groups led by Peter Druschel, Rodrigo Rodrigues, and Cristian Danescu-Niculescu-Mizil respectively.

External collaborators include researchers from Saarland University (Aniket Kate), Microsoft Research (Stefan Sariou and Emre Kiciman), UFMG (Fabricio Benevenuto), Northeastern (Alan Mislove), AT&T (Balachander Krishnamurthy), KAIST (Meeyoung Cha), Stevens Institute of Technology (Winter Mason), and IIT Kharagpur (Niloy Ganguly).

Publications. The group members regularly publish their research in the top conferences and workshops in their field. During the reporting period, members have co-authored papers at IMC [152, 181], EuroSys [203], CoNEXT [168], Usenix Security [187], WWW [106], ICWSM [141, 133], CSCW [48] SIGIR [105], CIKM [134, 107], COMSNETS [202]and WOSN [189].

Software, tools, and data. The group strives to make its software, tools, and data sets publicly available 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. Over the last two years, more than 500,000 end users worldwide have used the Glasnost software we designed to test the traffic management policies of their access ISPs (e.g., cable and DSL providers).

We also developed an application on the Facebook social network called FriendList Manager to help users manage their privacy settings better. To date, the app has more than 1250 users. Further, we have publicly deployed search systems that we built over the Twitter social network for finding expertise of users, experts on a given topic, and finding interesting and important content on a given topic. The developers at Twitter have acknowledged (in personal communication) the influence of our system designs on the design of their official Twitter search systems.

Teaching. Krishna Gummadi and Bimal Viswanath taught a graduate seminar on social computing systems in 2013.

External funding. The research of the group has been partially funded by fellowships from Google Research. Juhi Kulshrestha won a Google European Doctoral Fellowship as well as an Anita Borg Scholarship. The group also receives funding under an IMPECS grant (Indo-Max Planck Center for Computer Science) for its research collaborations with IITs in India.

Invited talks and awards. Krishna was an invited speaker at several emerging workshops on online social networks including PSOSM and SIMPEX (cohosted with WWW conference), WOSS (cohosted with VLDB), and PADE (cohosted with SIGMETRICS). Krishna was also invited to give a series of lectures on online social networks at the ESSCASS Summer School in 2012 and was on the expert panel on social search at CIKM 2011.

Farshad Kooti and Krishna Gummadi also received the Best Paper Award at the AAAI's ICWSM 2012 conference for their work on understanding the diffusion of social conventions in the Twitter social network.

Service. Internally to MPI-SWS, Krishna served as the chair of the Faculty Recruiting Committee for the 2013 hiring season, and as a member of the same committee for the 2012 hiring season.

Externally, Krishna has served as the program chair of IMC 2013 and the WWW 2013 Security & Privacy track. He also served on the program committees of IMC 2012, WSDM 2012, SIGCOMM 2012, NSDI 2013, COSN 2013, VLDB 2013, and NDSS 2014. Krishna currently serves on the steering committee member of Measurement Lab and on the technical advisory board of the newly formed ACM conference on online social networks (COSN).

9.2 Research agenda

9.2.1 Towards Dependable Social Computing Systems

Our recent research has focused on understanding and building more dependable social computing systems. Social computing systems refer to online computer systems that enable users to interact, collaborate, and compete with one another. Examples of social computing systems include social networking sites like Facebook and Google+, blogging and microbloging sites like Twitter and LiveJournal, content sharing sites like YouTube and Flickr, social bookmarking sites like Delicious and Reddit, crowd-sourced opinion sites like Yelp and eBay seller ratings, and social peer production sites like Wikipedia and Amazon's Mechanical Turk. Recently these systems have become tremendously popular and they are bringing profound changes in the ways individuals, organizations, and governments exchange information in our society.

Despite their popularity, social computing systems are still in their infancy and they bring with them new threats, challenges, and opportunities: first, users in many social computing systems operate behind weak identities, leaving the systems vulnerable to attacks from Sybil (multiple forged or fake) identities. Second, social computing systems caputre unprecendented amounts of personal information about users' activities and preferences that could be disclosed to others in ways that violate users' privacy. Third, the search and recommendation services offered by many social computing systems can on one hand overload users with irrelavent information, if the services are not personalized to their individual tastes, but on the other hand, lead to "filter bubbles", in which users are rarely exposed to diverse information that could challenge their world view, if the services are tailored to individuals' known interests. Finally, the ability to virally propagate ideas and information through the social networks between users in these systems, offers novel opportunities for raising public awareness on important societal issues and conducting word-of-mouth marketing campaigns, but it could also be misused for manipulating public opinion by spreading misinformation and rumors.

Below we present our recent work aimed at (and future plans for) **building more dependable social computing systems**. Specifically, we discuss our strategies for (i) mitigating the security and privacy threats facing these systems, (ii) tackling the challenge of designing information retrieval systems that provide trustworthy, relevant yet diverse perspectives on topics of interest to users, and (iii) exploiting new opportunities that arise from viral dissemination and propagation of information in societies. We use a two-step methodology to achieve our goals: in the first step, we plan to measure, analyze, and model the social networks and interactions between the hundreds of millions of users in current social computing systems to understand the micro- and macro-level behaviors of users. In the second step, we plan to leverage the insights from our computational social science studies to design more dependable social computing systems.

9.2.2 Determining the trustworthiness of weak social identities

Past Work: In the past, we focused on methods to leverage social networks to defend against Sybil (multiple identity) attacks. We studied Sybil identity detection schemes that rely on analyzing social networks' structure to identify Sybil nodes and evaluated their performance over real-world social networks. Our analysis revealed that the effectiveness of Sybil identity detection schemes is limited over real-world social networks that exhibit well-defined community structures.

Recent Work – during the review period: In light of the challenges with detecting Sybil identities in social networks, we proposed a fundamentally different approach to leveraging social networks against Sybil attacks: Sybil tolerance [202]. Rather than focus on identifying nodes as Sybils, Sybil tolerance focuses on designing systems that strictly bound the impact of Sybil nodes. We proposed a general methodology for designing Sybil tolerant systems using credit networks, which provide a way to model trust between nodes in a social network and support payments between arbitrary nodes. In practice, these payments require max flow computations on a social network graph, and they cannot be easily scaled to large network graphs. We designed Canal, a system that uses landmark routing-based techniques to effeciently approximate credit payments over large networks [203]. To demonstrate the effectiveness of our approach, we designed and evaluated, Genie, a Sybil-tolerant system that leverages social networks to thwart large-scale crawls of social networking sites [168].

Ongoing and Future Work: In ongoing work, we are planning to move from the problem of detecting Sybil identities to the problem of detecting *crowd computations* that have been tampered by Sybil identities. Many popular social and e-commerce sites like Twitter, YouTube, Amazon, eBay, and Yelp rely on the wisdom-of-crowds for rating and ranking information, users, products, and businesses. Today, these *crowd computations* are vulnerable to Sybil attacks. In ongoing work, we are investigating a different approach that can detect if a large-scale crowd computation has been tampered by Sybil participation, even when we cannot identify which of the participating identities are Sybils. Our key insight is that the distributions or *entropies* of many attributes in Sybil and non-Sybil user profiles tend to be very different in practice. Consequently, Sybil tampered and untampered computations tend to exhibit very different entropies in their participants' profiles. We leverage this insight to check if a given computation has been tampered by comparing the entropy of various attributes in the profiles participating in the given computation with that of profiles participating in known untampered computations.

9.2.3 Enabling users to better manage their data privacy

Past Work: Today users share their privacy-sensitive personal data (e.g. family videos) by uploading it to data centers owned and managed by social networking sites like Facebook. In the process, they lose control over their data to the site operators. In our previous work, we designed the Stratus system to enable users to regain control over their data sharing. In Stratus, users share their data directly from their home networks, which are under their control. We built inexpensive, low-power, always-on personal data servers at home using modern residential gateways and used them to share personal data.

Recent Work – during the review period: The growing popularity of content sharing over social websites requires end users to be content managers. Today, for every single piece of data shared on sites like Facebook – every wall post, photo, status update, friend request, and video – the uploader must decide which of his friends, group members, and other Facebook users should be able to access the data. In recent work [152], our overarching goal is to improve the defaults and provide better tools for managing privacy. To that end, we first attempted to quantify the magnitude of the problem of managing privacy by measuring the disparity between the desired and actual privacy settings of users in the Facebook social networking site. We deploued a survey, implemented as a Facebook application and recruited 200 Facebook users via Amazon Mechanical Turk. We found that, overall, privacy settings match users' expectations only 37% of content remains shared with the default settings.

Next we explored the potential to assist users in selecting appropriate privacy settings by examining the friend lists created by users to share their data with subsets of their friends [153]. We found that membership of friend lists is correlated with the social network. We implemented and deployed a Facebook application called FriendList Manager to show that we can automatically infer the membership of a user's friend lists by detecting groups within the 1-hop social network interconnecting the user's friends. Our application was used by more than 1000 Facebook users and data collected from the users suggests that FriendList Manager helps simplify the complex task of configuring privacy settings of users.

Ongoing and Future Work: Having tackled the challenge of configuring access control settings, in future work, we plan to examine the limits of access controls as a model for managing privacy in today's online world. We suspect that traditional access control model is 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 are currently exploring an alternate model for information privacy called *exposure*, which captures the set of people expected to learn an item of information eventually. We are investigating mechanisms to enable users to control the exposure of their personal data.

9.2.4 Finding relevant yet diverse information in the crowds

Many social networking sites have emerged as a popular platform for exchanging real-time information on the Web. For example, the Twitter microblogging site is used by hundreds of millions of users ranging from popular news organizations and celebrities to domain experts in fields like computer science and astrophysics and spammers. As a result, the quality of the *crowd-sourced* information exchanged on these sites is highly variable and finding the users that are authoritative sources of relevant and trust-worthy information on specific topics (i.e., topical experts) is a key challenge.

Recent Work – during the review period: In our recent work, we addressed this challenge by first leveraging the wisdom of crowds to identify expert or atuhoritative users on different topics and then using the experts to identify content important and relevant to a specific topic [189, 105, 107, 48]. Specifically, we tested these ideas in the context of the Twitter microblogging site. We showed that information gathered from the set of experts on a given topic is significantly more trustworthy and relevant to the topic than the information gathered from random crowds. One of the key findings of our study is that in microblogging sites like Twitter, wisdom of experts often trumps the wisdom of the crowds. This is in part because the wisdom of experts informs the wisdom of the crowds (i.e., information posted by experts is discussed widely by the crowds) and in part because experts interact with the crowds on these sites, extracting useful and important information (wisdom) from the crowds. Based on our findings, we designed and deployed search systems for finding topical experts and information on the Twitter site and our evaluation showed that our search systems outperform the official Twitter search system that is largely based on tapping the wisdom of the crowds.

Ongoing and Future Work: While we succeeded in designing systems to find relevant and trustworthy information from the crowds, a challenge for the future lies in retrieving *diverse* pieces of information or opinions on any given topic from the crowds. Retrieving diverse information is a much harder challenge because the goals of diversity are often at odds with the goal of retrieving relevant information on a topic. Finding a right trade-off between the two is the subject of our ongoing investigations.

9.2.5 Understanding information propagation over social networks

Past Work: In the past, we have conducted some of the earliest studies of information dissemination in online social networks by (i) collecting and analyzing large-scale traces of photo dissemination in the Flickr social network and URL dissemination in the Twitter social network, (ii) studying the role or influence of individual users in propagating the information in the social network and (iii) analyzing how users in social media sites like Twitter are exposed to news both directly from a large and diverse set of media sources and indirectly from the information forwarded by other users.

Recent Work – during the review period: In recent work, we focussed on understanding how social conventions emerge and propagate through social networks [133]. The way in which social conventions emerge in communities has been of interest to social scientists for decades. We studied the emergence of a particular social convention on Twitterthe way to indicate a tweet is being reposted and to attribute the content to its source. Initially, different variations were invented and spread through the Twitter network. We found that the inventors and early adopters were well-connected, active, core mem- bers of the Twitter community. The diffusion networks of these conventions were dense and highly clustered, so no single user was critical to the adoption of the conventions. Despite being invented at different times and having different adoption rates, only two variations came to be widely adopted. Our study describes this process in detail, highlighting insights and raising questions about how social conventions emerge. Our analysis yields useful insights on the extent to which one could predict the adoption of social conventions in the future [134].

Ongoing and Future Work: In ongoing and future work, we plan to investigate the role of recommendation systems and other site-operator controlled mechanisms (like featuring content on the front page of the site) in disseminating information. To date, most studies of information propagation on social networks, including our own, ignore such external factors. However, as social networking sites increasingly deploy such mechanisms, it is important to understand how they impact the viral propagation of information over social networks.

10 The Rigorous Software Engineering Group

10.1 Overview

The report covers the period from May 2011 – October 2013. 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 two graduate students (Zilong Wang and Johannes Kloos), and two postdoctoral researchers (Dmitry Chistikov and Rayna Dimitrova). With the departure of Ruzica Piskac, two students (Filip Niksic and Susanne van der Elsen) will join the group in October.

Two postdoctoral researchers (K.C. Shashidhar and Shahram Esmaeilsabzali) were associated with the group in the reporting period. Shashidhar accepted a software engineering position at Mathworks. Esmaeilsabzali accepted a postdoctoral researcher position at the University of Waterloo.

Majid Zamani and Indranil Saha graduated with Ph.D. degrees (from UC Los Angeles), in Summer 2012 and Summer 2013, respectively. Majid joined TU Delft as a postdoctoral researcher, and recently obtained an assistant professor position at TU Delft. Indranil joined UC Berkeley and University of Pennsylvania as a postdoctoral researcher in the ExCAPE project (with Sanjit Seshia and George Pappas).

Collaborations. The group has one joint publication each with the Synthesis, Analysis, and Automated Reasoning Group (led until recently by Ruzica Piskac) and the Dependable Systems Group (led by Rodrigo Rodrigues). There is an ongoing collaboration on infinite-state systems with Roland Meyer and his students at the University of Kaiserslautern (4 publications).

Externally, the group has ongoing collaborations with researchers at IMDEA Madrid (Pierre Ganty, Michael Emmi), Indian Institute of Science, Bangalore (Aditya Kanade), IST Austria (Krishnendu Chatterjee), TU Delft (Majid Zamani and Alessandro Abate), TU Munich (Javier Esparza), UC Los Angeles (Lei He, Todd Millstein, Jens Palsberg, and Paulo Tabuada), UC San Diego (Ranjit Jhala), and Toyota Motors (Jyo Deshmukh and Koichi Ueda).

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

In embedded systems and control, the group has published 5 papers in EMSOFT [164, 165, 185, 191, 82], 3 papers in CDC [209, 163, 208], 2 papers in HSCC [161, 184], 1 paper in CAV [167], 1 paper in DATE [131], 1 paper in NFM [162], 1 paper in Formal Methods in System Design [63], and 1 paper in Systems and Controls Letters [210].

In computer-aided verification, the group has published 4 papers in CAV [93, 132, 126, 127], 2 papers in CONCUR [166, 121], 2 papers in SAS [160, 147], 1 paper in LICS [92], 1 paper in FSTTCS [56], 1 paper in TACAS [108], 1 paper in FASE [154], 1 paper in CIAA [73], 1 paper in ATVA [62], 1 paper in Fundamenta Informaticae [72], 2 papers in Int. Journal on Foundations of Computer Science [65, 64], 1 paper in TOPLAS [98], 1 paper in Formal Methods in System Design [99], and 1 paper in Software Tools for Technology Transfer [113].

In programming languages and systems, the group has published one paper in OOPSLA [192], 1 paper in PLDI [95], and 1 paper in Eurosys [57].

We also published three invited papers [159, 158, 157].

Software, tools, data and technology transfer. Indranil Saha released a set of controller design and verification tools developed as part of his PhD: tools for equivalence verification between Simulink and C code, tools for type checking Simulink programs, and fixed point controller design for guaranteed stability and performance (COSTAN). Indranil's work on equivalence verification was used by Toyota Engineering North America's verification team.

Teaching. Rupak Majumdar taught the graduate course on Reactive Systems Verification at University of Kaiserslautern in Summer 2011. He is organizing a Verification Seminar Series with the University of Kaiserslautern to allow different researchers working in verification to come together at regular intervals.

External funding. The research of the group has been partially funded by industrial grants from Intel and Toyota.

Invited talks and awards. Rupak was an invited speaker at the following conferences and workshops: Ershov Memorial Conference 2011, Verified

Software: Theories, Tools, Experiments (VSTTE) 2012, Reachability Problems (RP) 2013. He lectured in the Marktoberdorf Summer School in 2011 and 2013.

He received the following best paper awards: ACM SIGPLAN PLDI "Test-of-time" award for his 2001 paper Automatic abstraction of C programs [38], EAPLS Best Paper award at the European Theory and Practice of Software (ETAPS) conference, 2012 (for his FASE 2012 paper [154]), ACM TODAES best paper award 2011 (for the best paper in all volumes of ACM Transactions on Design Automation of Embedded Systems in the year 2011, for the paper [76]).

Service. Internally to MPI-SWS, Rupak served as the Managing Director from July, 2012 (his term ends June, 2014). He also chairs the graduate program.

Rupak served on the program committees of ACM POPL 2012, CAV 2011-13, Fossacs 2012, FSTTCS 2012, ICALP 2013, LATA 2013, HSCC 2011-13, RTSS 2012 (Design and Verification track), and VMCAI 2012-13.

Rupak organized (along with Ranjit Jhala) an NSF Workshop on the Future of Formal Methods in 2012.

10.2 Research agenda

10.2.1 Current Directions

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

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 usual design flow for embedded control systems consists of two phases. In the first phase, the control engineer models the dynamics of the plant and, using tools from control theory, designs a real-valued control function that ensures the closed loop system (plant and controller) has certain desirable properties. In the second phase, the control function is implemented in software and integrated into the system along with several different implementation constraints. The first phase works in the realm of real analysis and ignores implementation issues such as fixed precision arithmetic or scheduling delays. The second phase implements the control function as a real-time software task, but does not usually consider the effects of implementation constraints on the behavior of the plant. Thus, it is not clear that the properties of the closed loop system proved at the level of control theory continue to hold for the implemented control system.

Our research in this area has focused on end-to-end arguments for control system design, where we add a perturbation term corresponding to implementation constraints while designing the controller, and use program analysis to bound the amount of perturbation. Our initial paper [11] showed how we can give a guarantee on the region of practical stability for a fixed-point implementation of a control law. In our analysis, we add a perturbation ϵ to the control input that models the possible deviation of the implementation from the "ideal" controller. Using tools from control theory, we bound the region of practical stability (the maximum deviation from the equilibrium point when the perturbed controller is used) as a function of the perturbation. Finally, using program analysis, we bound ϵ as the maximum error generated in the controller software due to fixed precision arithmetic. Together, this allows us to characterize the behavior of a closed loop system using the implemented controller when the perturbation is only due to fixed precision arithmetic.

Next, we study the effect of scheduling on controller design. Again, we take an end-to-end approach and jointly design a controller and a real-time scheduler. In designing the controller, we explicitly take into account that a fraction of control computations may not execute. This allows us to have a more relaxed scheduling problem (in which the scheduler can drop some tasks) while maintaining desirable properties of the closed loop system.

Finally, we considered the problem of *synthesizing* the best controller implementation. That is, out of many different controllers that stabilize the plant, we wish to find the one that guarantees the least implementation error for a given hardware budget. Our algorithm modifies the standard controller design methodology for LQR-LQG control and adds a secondary cost function to capture the effects of fixed precision arithmetic. We then find Pareto optimal controllers by exploring the space of controllers using particle swarm optimization. Unfortunately, the consideration of implementation errors makes the optimization problem non-convex, so standard LMI-based algorithms cannot be applied. Our experiments show that we can get over 4× reduction in implementation errors while remaining close to the LQR-LQG optimal controller.

The second direction of work has been on discrete abstractions of continuous dynamical systems. Here, we considered notions of ϵ-bisimulations for continuous dynamical systems. An ϵ-bisimulation of a continuous dynamical system is a discrete (finite-state) system such that every execution of the continuous system is matched by the discrete system to within a distance of ϵ. Discrete ϵ-bisimulations enable the use of discrete controller synthesis techniques on continuous systems. For continuous systems satisfying a certain well-formedness condition (called incremental input-to-state stability), finite ϵ-bisimulations exist for all $\epsilon > 0$. In [209, 210], we gave Lyapunov function characterizations of incremental stability. Additionally, we gave an algorithm to compute ϵ-bisimulations for digital control systems which takes variable time steps and can be more scalable than previous techniques [167].

Finally, we have worked on the foundations of mixed discrete and continuous systems (called hybrid systems). The main problem we considered is *robustness*. In continuous control theory, there is a well-understood notion of robustness, or input-to-state stability, that ensures that small perturbations in the inputs to a control system results in small perturbations in its outputs. Unfortunately, there is no good notion of robustness for hybrid systems. We made some progress towards with problem [161, 191], but the problem remains open.

Infinite-State Verification The second main focus of the group is in verification algorithms and tools for infinite-state systems. Here, we have worked on two main directions.

First, we have focused on the development of scalable tools for coverability analysis for Petri nets. Petri nets are a well-studied concurrency model. In previous work, we showed that the dataflow analysis problem for event-driven programs can be reduced to the analysis of Petri nets. Thus, tools for Petri net analysis can be used as a backend for analyzing event-driven programs.

We have recently developed a new forward algorithm for Petri net coverability based on the incremental inductive reachability idea of Bradley. Our algorithm generalizes Bradley's procedure to downward-finite well-structured transition systems, a class that includes Petri nets and most other well-structured transition systems considered in practice. Experimental results show that our algorithm is competitive with the best implementations of Petri net coverability.

We are in the process of building a distributed model checker for coverability that can scale to large clusters.

Second, we have worked on algorithms for the analysis of infinite-state systems. While in general the analysis of infinite state systems is undecidable, one may still get decidability for particular models and algorithmic questions. Some of our main results in this area are as follows.

1. A complexity analysis for the expand-enlarge-and-check (EEC) algorithm for Petri nets. EEC is a generic procedure to check coverability. While it works well in practice for many examples, its complexity was unknown.

2. Results on bounded languages. A bounded regular language is of the form $w_1^* w_2^* \dots w_k^*$ for fixed strings $w_1, \dots, w_k$. We showed decidability and complexity results for checking emptiness of multi-headed automata relative to bounded languages. As an application, we show that verification for a large class of undecidable infinite-state models becomes decidable when restricted to bounded sets.

3. A general model for provenance analysis in asynchronous programs. We give a general model of provenance in message-passing programs. Intuitively, the provenance of a message is the history of principals who have forwarded this message. We show that static analysis of provenance is decidable for asynchronous programs when the provenance properties are expressed as regular sets.

10.2.2 Future Directions

We now outline some research directions over the next two years.

Verification of Asynchronous Programs One goal for the next two years is to build a suite of automatic verification tools for asynchronous programs. We are considering two classes of programs: low-level systems code, such as file system implementations, and Android applications.

We are also considering several different techniques. The first is to apply abstraction to the original programs and reduce the problem to Petri net coverability. Here, we can use our current work on scalable coverability checkers. However, we have to develop precise abstraction procedures. The second is to extend the idea of liquid types and come up with dependent type systems for asynchronous programs. We have an initial design for such a type system, and we are currently evaluating the type system on programs written in MirageOS, a library OS written in Ocaml. The third idea is to develop assume-guarantee reasoning principles for asynchronous programs and then infer the assumptions and guarantees using standard program analysis techniques.

We have already developed infrastructure to analyze Android applications. Initially, we are focusing on data race detection in the presence of multi-threading and event-based dispatch. But the same infrastructure should be useful in our further work.

Infinite-State Models The second main direction will be in the foundations of infinite-state verification. We shall continue our current work on parameterized verification and look at algorithmic and decidability questions for infinite-state systems. We are currently considering the following questions.

1. Complexity questions in language theory. The complexity of some basic questions in formal language theory remain open. For example, the complexity of equivalence of deterministic pushdown automata lie between P and primitive-recursive! We have recently obtained tight bounds on decision problems of *unary* deterministic pushdown automata, a very special case that was nevertheless open. We plan to investigate further complexity questions from language theory.

2. Parameterized verification for parameterized topologies. We are studying parameterized verification problems where the goal is to verify a property of a family of designs. We are motivated by verification problems in train control systems, where the system has to be verified against a parameterized family of track configurations. We have some preliminary results in the bounded case, but many verification questions remain open.

Marrying Markov Decision Processes with Probabilistic Databases In collaboration with Michael Backes, Peter Druschel, and Gerhard Weikum, we are planning a multi-year project to study the interactions between privacy, accountability, compliance, and trust for online interactions. Our proposal is currently under review with the European Research Council.

As part of the project, we plan to investigate extensions to probabilistic inference models studied in verification (such as Markov decision processes) with probabilistic databases. As a specific application, we plan to develop inference tools that evaluate a user's online history and produce advice about the effects of a user action on his or her privacy. There is a rich connection between deterministic models studied in verification and the theory of databases. A similar connection should be realizable in the probabilistic case as well.

11 The Software Analysis and Verification Group

11.1 Overview

The report covers the period from May 2011 to October 2013. The Software Analysis and Verification group works on specifying and verifying 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 is led by Viktor Vafeiadis, and consists of one Masters student (Mustafa Zengin), who joined the group after the departure of Umut Acar, and two doctoral students (Marko Doko and Soham Chakraborty), who joined the group in October 2013. Over the reporting period, the group had three interns (Marko Doko, Filip Sieczkowski, Chinmay Narayan) each staying for 2-3 months.

Collaborations. The group collaborates very closely with Derek Dreyer's group and to a lesser extent with Allen Clement's, Rodrigo Rodrigues's and Rupak Majumdar's groups. We also collaborate with researchers at Cambridge (Peter Sewell), Microsoft Research (Byron Cook, Matthew Parkinson), IMDEA (Alexey Gotsman, Aleksandar Nanevski), INRIA (Francesco Zappa Nardelli), IST Austria (Tom Henzinger, Ali Sezgin), Leuven (Bart Jacobs), TAU (Noam Rinetzky), and Purdue (Suresh Jagannathan).

Publications. The group has published regularly in the top conferences and journals of its field. During the reporting period, group members have co-authored one JACM [188] journal publication, two POPL [123, 125], one OOPSLA [199], one ICFP [213], one CONCUR [118], one ITP [198], one LICS [124], one MFPS [197], one SAS [200], and one TASE [211] conference publications.

Formal proof developments. All of the group's publications reported above except for [124] and [118] (i.e.,[188, 123, 125, 199, 213, 198, 197, 200, 211]) come with machine-checked proof developments in Coq or Isabelle/HOL, all of which are available online and some of which (paco, mtac, and adjustable references) are general-purpose libraries.

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).

Service. In the reporting period, Viktor has served on the program committees of POPL 2014, Coq 2013, ICFEM 2013, SEFM 2013, PCI 2013, TASE 2013, ICFEM 2012, SEFM 2012, CPP 2011, APLAS 2011, and ICFEM 2011, and on the external program committees of PLDI 2013 and POPL 2013. He chaired the MPI-SWS faculty recruiting committee in 2012 and served on the MPI-SWS student admissions committee in 2013.

11.2 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 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 focuses on reasoning about concurrent programs, especially ones running in a relaxed memory model, such as TSO or the C11 model. We have pursued two main goals in this direction. The first is to prove functional correctness properties, such as linearizability and refinement. As a means for achieving this goal, we develop suitable program logics for naturally expressing such reasoning. The second is to prove the correctness of program transformations, such as compiler optimizations, over concurrent programs running under relaxed memory models. As a step towards achieving this goal, we have also looked at compositional techniques for reasoning about program equivalence in the sequential setting.

Another important part of our work concerns the Coq interactive theorem prover and improving its applicability to reason about software. Specifically, we have proposed a lightweight approach for verifying stateful programs in Coq, whose use of mutable state is entirely local and hidden from the specification [198]. Next, we have developed a new typed tactic language, mtac [213], that is better suited for automating proofs than Coq's built-in untyped tactic language. Finally, we have also written a library for performing parametric coinduction (paco [125]) that outperforms the built-in support for coinduction in the kind of simulation proofs that are typical in software verification.

11.2.1 Reasoning about concurrency and compilers

Aspect-oriented linearizability proofs [118] Proving correctness of concurrent algorithms is typically done by monolithic simulation arguments that identify the so-called linearization points. These arguments are often quite complex and scale poorly to advanced non-blocking concurrency patterns, such as helping and optimistic updates.

Together with Tom Henzinger and Ali Sezgin, we proposed a more modular way of checking linearizability of concurrent queues by reducing the task of proving linearizability to establishing four basic properties, each of which can be proved independently by much simpler arguments.

In the future, we plan to extend this approach to handle other data structures such as sets and to develop a robust verifier for proving linearizability. Further ahead, we would like to extend this approach to the relaxed concurrency setting and check for the analogues of linearizability in that setting. Soham Chakraborty has started working on this problem.

Relaxed separation logic (RSL) [199] So far, concurrent program logics have been defined only for a sequentially consistent programming model. The adoption of relaxed consistency models by programming languages such as C/C++ and Java renders most of this work obsolete.

We developed a version of separation logic, which we call *relaxed separation logic* (RSL), that is sound under acquire-release concurrency, as provided by the recent C11 relaxed memory model. In particular, we show that the concept of resource ownership inherent to separation logic as well as a restricted form of ownership transfer are sound under C11.

We are currently working on extending the kind of reasoning provided by RSL in two ways: (1) to handle a larger fragment of the C11 model such as memory fences, and (2) to allow move advanced forms of reasoning, such rely/guarantee and ghost variables. Ultimately, we would also want to build automated verification tools for programs running under relaxed concurrency. Marko Doko is currently working on the first extension, whereas the second is in collaboration with Derek Dreyer and Aaron Turon.

Verified compilation for the TSO relaxed memory model [200, 188] In the context of the CompCertTSO project, which extended Leroy's CompCert compiler [149] to a concurrent setting, we implemented two new TSO-specific optimizations that remove redundant memory barrier instructions, and proved their soundness in Coq. While the optimizations we implemented are very cheap to perform by a standard thread-local control flow analysis, their correctness is much more subtle and required a different proof strategy from the optimizations we had encountered previously in the CompcertTSO project. To carry out the proof, we came up with a non-standard global simulation argument, which effectively incorporates a Boolean prophecy variable in an otherwise forward simulation.

Fault-tolerant parallelism [211] When running big parallel computations on thousands of processors, the probability that an individual processor will fail during the execution cannot be ignored. Computations should be replicated, or else failures should be detected at run-time and failed subcomputations reexecuted. We followed the latter approach and proposed a high-level operational semantics that detects computation failures, and allows failed computations to be restarted from the point of failure. We implemented this high-level semantics with a lower-level operational semantics that provides a more accurate account of processor failures, and proved in Coq the correspondence between these semantics.

Parametric bisimulations [123] In joint work with Chung-Kil Hur, Derek Dreyer, and Georg Neis, we introduced relational transition systems (a.k.a. parametric bisimulations, PBs), which are a powerful technique for reasoning about program equivalence in ML-like languages. In our ongoing work, we are exploring the generalization of PBs to the inter-language setting and their application to compositional compiler verification. For more details, please see Section 5.2.

11.2.2 Improving machine-assisted proving technology

Given the inherent undecidability of most interesting verification questions, and the desire to verify increasingly complicated software components, one has to resort to semi-manual proof. Such proofs are generally best performed interactively in proof assistants like Coq and Isabelle.

In our use of these proof assistants, however, we have found the current interactive theorem proving technology to be lacking in some respects, which we have attempted to rectify.

Paco: Parameterized coinduction [125] Coinduction is one of the most basic concepts in computer science underlying many verification problems, such as showing that two concurrent programs are equivalent. Nevertheless, its standard expositions are lacking in two key respects: they do not support compositional reasoning (i.e., breaking proofs into separate pieces that can be developed in isolation) nor incremental reasoning (i.e., developing proofs interactively by starting from the goal and generalizing the coinduction hypothesis repeatedly as necessary).

In joint work with Chung-Kil Hur, Georg Neis, and Derek Dreyer, we developed *parameterized coinduction*, which achieves both compositionality and incrementality, and have implemented it as Coq library called paco. Besides its compositionality benefits, using paco leads to faster and more robust proof developments than using Coq's builtin cofix tactic.

Mtac [213] Effective support for custom proof automation is essential for large-scale interactive proof development. However, existing languages for automation via tactics either (a) provide no way to specify the behavior of tactics statically within the logic of the theorem prover or (b) rely on advanced type-theoretic machinery that is not easily integrated into established theorem provers.

Developed jointly with Beta Ziliani, Derek Dreyer, Neel Krishnaswami, and Aleks Nanevski, Mtac is a lightweight but powerful extension to Coq for supporting dependently-typed tactic programming. Mtac tactics have access to all the features of ordinary Coq programming, as well as a new set of typed tactical primitives. We avoid the need to touch the trusted kernel typechecker of Coq by encapsulating uses of the new tactical primitives in a monad, and instrumenting Coq so that it executes monadic tactics during type inference.

Mtac has received significant attention already from Coq users and developers, in the hopes that it can eventually serve as a replacement for Coq's existing *ad hoc* tactic language, Ltac. The work on Mtac will serve as the foundation of Beta Ziliani's PhD thesis.

Adjustable references [198] Mutable state underlies important optimizations such as path compression in union-find algorithms and memoization, and is often crucial to achieving good performance. Nevertheless, verified programs rarely use mutable state because of its substantial verification cost: one must either commit to a deep embedding or follow a monadic style of programming. To avoid this cost, we extended Coq with a type of adjustable references. These are like ML references, except that the stored values are only partially observable and updatable only to values that are observationally indistinguishable from the old ones. As a result, verification can ignore the updates.

Part II
Adjunct Research Groups

12 The Information Security and Cryptography Group

12.1 Overview

The report covers the period from May 2011 – October 2013. The group's research interests are in theoretical foundations and applied aspects of information security and cryptography. Major research topics have been the design and verification of security protocols and implementations, privacy and anonymity, linking formal methods and cryptography, and using foundational approaches for mobile device security.

Personnel. The group is led by Max Planck Fellow Michael Backes, and currently has one postdoc (Dario Fiore). Aniket Kate was a postdoc in the group until mid-2012, when he become a junior research group leader in the cluster of excellence MMCI. Michael Backes additionally has the chair for information security and cryptography at Saarland University, and he is the director of the center for IT security, privacy, and accountability (CISPA). The works done in the MPI-SWS group are usually conducted jointly with the university group, which currently has seven graduate students (Fabian Bendun, Sebastian Gerling, Praveen Manoharan, Sebastian Meiser, Esfandiar Mohammadi, Raphael Reischuk, and Malte Skoruppa).

Collaborations. The group has a joint project with the distributed systems group (led by Peter Druschel) on achieving forgetfulness in the Internet. It has recently started collaborations with the foundations of computer security group (led by Deepak Garg) on the formalization of novel privacy regulations. Externally, the group has collaborated with researchers at IMDEA, Cornell, ETHZ, Waterloo, Stanford, CMU, Microsoft Research Cambridge, IBM Research Zurich, the Max Planck Institute for Informatics, Karlsruhe Institute of Technology, and Technische Universität Darmstadt.

Publications. The group has published regularly in the top conferences and journals of its field. During the reporting period (2011-2013), group members have co-authored three CSF [29, 12, 24], one S&P [28], two CCS [20, 35], three NDSS [18, 33, 32], one ESORICS [19], one WWW [177], one POST [13], one ASIACRYPT [31], one CT-RSA [17], two TOSCA [26, 34], two ASIACCS [25, 14], one TACAS [22], one ESSoS [204], one CPP [27], one ISVC [156], two DPM [36, 23], two WPES [21, 30], one NFM [15], and one IJIS [16] publications.

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 was AppGuard: within the last 12 months, more than 1,000,000 end users have downloaded AppGuard to protect themselves against malicious apps on their Android mobile devices.

Patents and technology transfer. After the tremendous response to our AppGuard technology (more than 1 Mio. downloads of the freely available, basic version), we decided to offer an extended version of AppGuard as a commercial product, which is sold by a spin-off company.

Press. Articles describing the group's work on AppGuard (protecting against malicious apps on Android) have appeared in numerous popular news media (Frankfurter Allgemeine Zeitung, Sueddeutsche Zeitung, C't, Heise, various TV news programs, etc.)

Teaching. Michael Backes taught the undergraduate courses on Cryptography and Security in 2011 and 2012, respectively, and an advanced course on smartphone security in 2013. Additionally, he held two graduate student seminars in the last two years.

External funding. The research of the group has been partially funded by an ERC starting grant on end-to-end security, by the Excellence Cluster on Multimodal computing and interaction, as well as by the newly founded center for IT security, privacy, and accountability (CISPA).

Invited talks, awards, and honors Michael Backes was an invited speaker at ETAPS 2011. He was named the leading German scientist under the age of 40 by the Financial Times Germany, and one of the 100 most important IT people in Germany in 2011 by the Computerwoche Newspaper.

Service. Michael Backes is the program co-chair of the IEEE Symposium on Security & Privacy (S&P) in 2013 and 2014. He was the program co-chair of the IEEE Computer Security Foundations Symposium (CSF) in 2011. He currently serves on the steering committee of IEEE S&P, IEEE CSF and ESORICS. In the reporting period, Michael in particular served on the following program committees: ACM CCS 2011, IACR Crypto 2011, IEEE CSF 2011, ESORICS 2011, PETS 2011, NDSS 2011, ACM CCS 2012, ESORICS 2012, ESORICS 2013, IEEE S&P 2013, IEEE S&P 2014, IEEE CSF

2014, ESORICS 2014. He is furthermore an Associate Editor of Springer's International Journal on Information Security and Cryptography.

Dario Fiore served on the following program committees: PKC 2011, Pairing 2012, IWSEC 2012, Pairing 2013, IWSEC 2013, and WAHC13.

12.2 Research agenda

The group's research interests are in theoretical foundations and applied aspects of information security, privacy, and cryptography, such as the design and verification of security protocols and implementations, privacy and anonymity, linking formal methods and cryptography, and using foundational approaches for mobile device security. In the last two years, the group's research interest slightly shifted towards the analysis and development of practically deployable systems, in particular for mobile security.

The following sections describe four of our main research thrusts and highlight some contributions. We intentionally focus on that work that has been primarily conducted at MPI-SWS as part of the fellowship, in contrast to works that have primarily been conducted at Saarland University. A strict separation of these two does not exist though.

12.2.1 Design and analysis of security protocols and programs

Security proofs of cryptographic protocols and programs are known to be difficult, and work towards the automation of such proofs has started soon after the first protocols were developed.

Delegatable computation. We proposed a novel cryptographic technique for delegatable computation [20]. 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 propose a scheme that achieves these goals for computations of quadratic polynomials over multiple variables and allows for constant-time verification.

A framework for data-driven web-applications. We developed a novel method for enabling fast development and easy customization of interactive data-intensive web applications [177]. Our approach is based on a high-level hierarchical programming model that results in a very clean semantics of the application while at the same time creating well-defined interfaces for customization of application components. A prototypical implementation of a conference management system shows the efficacy of our approach.

Automated Synthesis of Secure Distributed Applications. We introduced a framework for the automated synthesis of security-sensitive distributed applications [33]. We provide the programmer with a high-level declarative language for specifying systems and the intended security properties (e.g., authorization and privacy), while abstracting away from any cryptographic details. A compiler takes as input a high-level specifications and produces the corresponding provably secure cryptographic implementations. We experimentally evaluate the feasiblity of our approach.

A type system for verifying of reference implementations. We introduced a novel type system for verifying the security of reference implementations of security protocols written in a core functional language [26]. The type system 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. This increased expressivity enables the analysis of important protocol classes that were previously out of scope, such as signatures of private data and encryptions of authenticated data,and in general applications based on zero-knowledge proofs. The type system comes with a mechanized proof of correctness and an efficient type-checker.

Miscellaneous. We developed a framework for deriving security protocols from a goal-driven language [34]. We developed a technique for automatically verifying typing constraints for a data processing language [27]. We proposed novel cryptographic protocols for verifiable secret sharing [31, 17]. We introduced a technique for analyzing quantitative information-flow w.r.t. non-uniform distributions [14]. We introduced a framework for analyzing of traffic side-channels of web applications [18]. We introduced a steganographic system based on diffusion-based image compression [156].

12.2.2 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 exclude that these abstraction miss attacks, so-called computational soundness results have been established. More recently, formal methods have been used to directly show the correctness of cryptographic proofs.

Formal analysis of the Merkle-Damgård construction. We presented the first machine-checkable proofs of collision-resistance and indifferentiability of Merkle-Damgaard construction for hash functions [12], which is a main building block of many hash functions, e.g., all SHA-3 finalists. Since practically used hash functions have known vulnerabilities but are core to many cryptographic constructions, there is an active search for a secure replacement. Our proofs are built and verified using an extension of the EasyCrypt framework, which relies on state-of-the-art verification tools such as automated theorem provers, SMT solvers, and interactive proof assistants.

Computational soundness without protocol restrictions. Computational soundness results prove that symbolic idealizations of cryptographic primitives characterize all attacks against the respective cryptographic realizations. Previous results came at the cost of imposing constraints on the set of permitted security protocols, e.g., dishonestly generated encryption keys must not be used. Identifying novel cryptographic definitions that the cryptographic realization needs to satisfy (which are satisfied by existing cryptographic), we established the first computational soundness result without protocol restrictions [35]. In particular, our result includes protocols that send, receive and use dishonestly generated encryption keys.

Improved computational soundness for ZK proofs. We improved on a previous computational soundness result of zero-knowledge (ZK) proofs, which had ad-hoc formalisms and required strong cryptographic realizations, leaving only highly inefficient realizations. We identified weaker cryptographic definitions that are sufficient for computational soundness of ZK proofs [13] and that are fulfilled by existing efficient ZK schemes.

Miscellaneous. We developed an extensible code generator for security protocols in Java [15].

12.2.3 Security of Mobile Devices

The pervasiveness of mobile devices makes mobile devices a prominent target for attackers. We highlight a recent work for securing mobile devices.

AppGuard – controlling third-party Android apps. We presented AppGuard, a powerful and flexible system for the enforcement of user-customizable security policies on untrusted Android applications [23, 22, 204]. The Android permission system turns out to be inadequate to protect the user against security and privacy threats. AppGuard offers complete mediation of security-relevant methods based on callee-site inline reference monitoring without requiring any changes to a smartphone's firmware or root access. We demonstrated the general applicability of AppGuard by several case studies. Our technique exhibits little overhead and its utility has been demonstrated by more than 1,000,000 downloads.

Miscellaneous. We proposed how to integrate mobile devices to strengthen the security of e-voting protocols [21].

12.2.4 Privacy and Anonymous Communication

In the presence of a strong economical interest in personal data the privacy of individuals is increasingly threatened.

ObliviAd: Privacy preserving online behavioral advertising. We proposed a privacy preserving architecture, called ObliviAd, for privacy preserving online behavioral advertising (OBA) [28]. OBA has become a rapidly increasing source of revenue for a number of web services, and it is typically conducted by third-party data analytics firms. This practice raises significant privacy concerns. Previous systems for privacy-preserving OBA either did not provide sufficient profile privacy and information accuracy, or required trusted third parties. Using a secure hardware token, ObliviAd meets these challenges with a privacy preserving advertisement distribution and a system for billing advertisers while preserving profile privacy. We presented the first formal security definitions for OBA systemand conducted a formal security analysis of ObliviAd using automated verification tools. Moreover, we experimentally evaluated the practicality of our approach.

Privacy-preserving social networks. We introduced a cryptographic framework to achieve access control, privacy of social relations, secrecy of resources, and anonymity of users in social networks [32]. We illustrated our technique on a core API for social networking, which includes methods for establishing social relations and for sharing resources. We analyzed the security of our protocols by formalizing the security properties and by verifying with an automated verification tool. We demonstrated the efficiency and scalability of our approach by building a prototypical implementation and conducting an experimental evaluation.

Privacy Preserving Accountable Computation. Accountability of distributed systems aims to irrefutably link malicious behavior to a malicious node and to allow honest nodes to disprove false accusations. Recent work shows how to achieve accountability, but previous approaches inevitably expose a node's private data. We argue that for randomized computations combining zero-knowledge proofs with previous approaches does not yield efficient solutions. We propose an alternative definition of accountable randomness and generalize it to a notion of privacy-preserving accountable randomized computation [19]. We present efficient instantiations for interesting classes of computations, in particular for digital signature schemes.

Analyzing and improving the onion routing network. The onion routing network Tor is undoubtedly the most widely employed technology for anonymous web access. We formalized the core onion routing (OR) protocol, based on the Tor-specification, and proved that this OR protocol is universally composable, a well-established cryptographic definition for composable security [24]. Our result greatly simplifies the process of analyzing OR anonymity metrics. We show that our results precisely characterizes the assumptions that a recent OR black-box model assumed in the analysis of the anonymity of OR. After analyzing Tor, we proposed an key-exchange protocol, called Ace, for Tor that we proved secure and that we showed to be more efficient that the current key-exchange protocol ntor [30]. We continued the analysis the anonymity guarantees of Tor by introducing novel quantitative anonymity notions for anonymous communication networks [29]. We showed that our anonymity notions implies previous anonymity notions, and we leveraged our previous analysis of OR to quantify the degree of anonymity that the OR protocol satisfies.

Miscellaneous. We developed a cryptographic construction to ensure query privacy of distributed hash tables, which have recently been considered for anonymous communication [25].

13 The Dependable Systems Group

13.1 Overview

In the beginning of 2012, the group leader joined the faculty of the Nova University of Lisbon, and remained affiliated with the institute as an adjunct faculty. Nonetheless, the research results presented in this report reflect the work that was conducted while the group leader was at the MPI-SWS. This report covers the period from May 2011 – October 2013.

During this period the group published in several top venues in the area of distributed computing and other areas of computer systems. The research conducted by the group focuses on several aspects of system dependability, from replication to testing and system security, with a strong focus in the area of cloud computing.

Personnel. The group is led by Rodrigo Rodrigues, and was comprised during the first several months of this period of one postdoc (Allen Clement) six doctoral students (Nuno Santos, Pedro Fonseca, Pramod Bhatotia, Cheng Li, and Daniel Porto) and one master's student (João Carreira). From these group members, Allen Clement joined as a faculty of MPI-SWS in mid-2012, Alexander Wieder joined the real-time systems group in early 2012, Daniel Porto left the MPI-SWS graduate program to join the Nova University of Lisbon graduate program in mid-2012, Nuno Santos has handed in his doctoral thesis and joined the faculty of the IST (the engineering school of the University of Lisbon) in mid-2013, and João Carreira concluded his master's in the end of 2011 and is now a research trainee with Edouard Bugnion at EPFL. The remaining students are still in the MPI-SWS graduate program.

Collaborations. The group has joint publications and ongoing submissions and collaborations with the following groups: Distributed Systems group (led by Peter Druschel), Programming Languages and Systems (led by Umut Acar), Networked Systems group (led by Krishna Gummadi), Rigorous Software Engineering group (led by Rupak Majumdar), Software Analysis and Verification (led by Viktor Vafeiadis), Real-Time Systems group (led by Björn Brandenburg), and the Robust Systems group (led by Allen Clement).

Externally, the group has collaborated with researchers at Cornell, Yale, CMU, EPFL, Nova University of Lisbon, Microsoft Research, IBM Research, and Federal University of Uberlândia (Brazil).

Publications. The group published in several top venues in the broad area of computer systems. In particular, group members have co-authored papers in the following top conferences: SOCC [47], FAST [45], NSDI [205], EuroSys [57], PODC [75], Usenix Security [187], OSDI [151], and Middleware [186]. In addition, the group has published a workshop paper at HotCloud [46], and two journal papers [180, 155].

Teaching. Rodrigo Rodrigues co-lectured the Operating Systems core course at Saarland University in the Summer semester of 2011. Cheng Li was a TA for that course.

External funding. The research of the group has been partially funded by a Yahoo! Faculty Research and Engagement Award and by an Amazon Web Services in Education Research Grant. Rodrigo Rodrigues was awarded an ERC Starting Grant in 2012.

Invited talks. Rodrigo Rodrigues gave a keynote talk at LADIS 2012: The 6th Workshop on Large Scale Distributed Systems and Middleware. He also gave invited seminar talks at IST Austria, IMDEA, TU Munich, and TU Wien.

Service. Internally to MPI-SWS, Rodrigo Rodrigues concluded his service as the graduate program chair. Nuno Santos was the student representative within several bodies of the MPI-SWS, and Pramod Bhatotia served as the student representative of MPI-SWS in the MPS.

Rodrigo Rodrigues served on the program committees of EuroSys 2012, DSN 2013, LADIS 2013, EuroSys 2013, and is currently serving on the PC for EuroSys 2014. He was the Shadow PC chair for EuroSys 2012. He was both workshop co-chair and publicity co-chair for SOSP 2011. He chaired the review committee of the 2013 EuroSys Roger Needham PhD award. In 2013, he also served on the doctoral thesis committee of Zarko Milosevic at EPFL.

Degrees. Nuno Santos handed in his doctoral thesis, and his final defense is scheduled for November 27, 2013.

João Carreira concluded his master's degree at IST (University of Lisbon), co-advised by Rodrigo Rodrigues.

13.2 Research agenda

The research agenda of the group is centered around three main vectors.

1. Visigoth fault tolerance. Existing fault models fail to capture the reality of data centers. On the one hand, the crash model is too optimistic since it assumes that all failures are silent, which does not capture arbitrary (commission) faults that happen at the dat a center scale. On the other hand, the Byzantine model captures arbitrary faults, but assumes a pessimistic scenario: a malicious adversary controlling system components at will. Similarly, the synchronous model optimistically assumes that all machines receive and process messages within certain time bounds; the asynchronous model pessimistically assumes that message transmission and processing may take arbitrarily long. The consequence of optimistic assumptions is putting system correctness at stake, while pessimistic assumptions increase replication costs and decrease performance. In this research we advocate a new approach called Visigoth fault tolerance (VFT) that more closely models the data center environment. Our new techniques represent a middle ground between crash and Byzantine fault tolerance, and also between synchrony and asynchrony: the number of replicas required is parameterized in a way that covers the spectrum between these traditional approaches.

2. Geo-replication. We are pursuing new ways to make geo-replicated systems fast, but without compromising the system correctness. To this end, we are researching automated ways to make geo-replicated systems fast if possible and consistent when necessary, by analyzing the code of operations in a geo-replicated service and assigning them appropriate consistency levels. In addition, we are researching new concurrency control protocols that achieve strong isolation levels within a reasonable latency envelope, namely a single cross data center round-trip.

3. Large scale parallel processing. We are researching new ways to improve large scale parallel processing, both in the context of distributed programming frameworks like MapReduce, Dryad, or Spark, and also in the context of multi-threaded computations running on multi-core machines. Currently, we are working on tools for performing pthread-based concurrent programs and sliding window computations based on the MapReduce paradigm incrementally. In the future, we intend to explore how these frameworks should be adjusted to better handle situations where data is continuously arriving at distant data centers and a global computations must be performed with requirements on latency, bandwidth, and CPU usage.

In addition, to these vectors, the group is wrapping up previous research projects on concurrency testing and cloud computing security.

Part III
Former Research Groups

14 The Programming Languages and Systems Group

14.1 Overview

This section covers the period from May 2011 – October 2012. The group's research interests are in parallel computation, incremental computation, and provenance.

Personnel. The leader of the group, Umut Acar, moved on to a faculty position at Carnegie Mellon University. Postdocs Arthur Charguéraud and Mike Rainey both accepted researcher positions at INRIA. Postdoc Joshua Dunfield continued on in Derek Dreyer's group and subsequently in Rupak Majumdar's group. PhD student Matthew Hammer completed his degree and moved on to a postdoc position at the University of Maryland. Another member of the group, Roly Perera, who was a visiting researcher for several years, moved on to a postdoc position first in Deepak Garg's group and subsequently at the University of Edinburgh. The remaining member of the group, Yan Chen, is working on completing his PhD thesis by the end of 2014. Yan Chen has already produced good work and should be able to complete on time. Pramod Bhatotia and Ezgi Cicek, who were jointly advised by Rodrigo Rodrigues and Deepak Garg respectively, are continuing their studies at MPI-SWS.

Collaborations. During the review period, collaborators included Deepak Garg, Rodrigo Rodrigues, and Alexander Wieder (MPI-SWS), Amal Ahmed (Northeastern), Guy Blelloch (CMU), James Cheney (Edinburgh), Matthew Fluet (Rochester), Ruy Ley-Wild (IMDEA Software Institute), Paul Levy (Birmingham), and John Reppy (University of Chicago).

Publications. During the reporting period, group members have co-authored one PLDI [69], three ICFP [60, 70, 173], two OOPSLA [116, 3], one ESOP [150], one POST [2], two PPoPP (the Symposium on Principles and Practice of Parallel Programming [5, 42]), four journal [61, 71, 190, 43] and two workshop [4, 6] publications.

Software. Hammer and Acar released CEAL, a C-based language for self-adjusting computation, which is available open-source from the MPI-SWS PLV (programming languages/verification) group web page (`http://plv.mpi-sws.org`).

External funding. Acar was awarded an ERC Starting Grant in 2012. Acar was also awarded a GIF grant, which was turned over to another German scientist.

Ezgi Cicek won an Anita Borg Scholarship from Google in 2012.

Invited talks. Acar was an invited speaker at the Workshop on the Theory and Practice of Provenance (2011) and the Dagstuhl Seminar on Principles of Provenance (2012).

Service. Acar served as program chair of the Workshop on the Theory and Practice of Provenance (2012), as general chair (and program committee member) of the Workshop on Declarative Aspects and Applications of Multicore Programming (2012), and as a member of the External Review Committee of POPL 2012.

14.2 Research agenda

The group's research investigates a broad range of issues in the design, implementation, and applications of programming languages and systems. The group's work often involves developing and applying techniques from the study of algorithms, as well as the theoretical study of programming languages.

Incremental computation. Input to programs often changes slowly or incrementally over time. In some applications, incremental changes to input result in only small changes in output, making it possible to respond to such changes asymptotically more efficiently than by re-running the whole computation. Traditionally, realizing such asymptotic efficiency improvements requires designing problem-specific algorithms known as dynamic or incremental algorithms, which are often significantly more complicated than conventional algorithms.

The group's work on *implicit type-directed* incremental computation [70, 69] made significant progress towards genuinely automatic incrementalization: given a few type annotations describing what can change over time, we can compile a conventional Standard ML program to an incremental program that is often asymptotically more efficient, leading to orders of magnitude speedups in practice.

To further improve the time and space performance of the implicit type-directed approach, Chen and Acar are developing techniques for parallel computation and granularity control, with the goal of extending the implicit type-directed approach to extremely large dynamic datasets.

Parallel computation. A classic problem in parallel computing is determining whether to execute a task in parallel or sequentially. If small tasks are executed in parallel, the overheads due to task creation can be overwhelming. If large tasks are executed sequentially, some processors will be underutilized. Charguéraud, Rainey, and Acar developed an approach to this *granularity problem* [3, 4] that combines programmer-provided *complexity annotations* with run-time profiling. By combining static information (the asymptotic information from the annotations) with dynamic information (the hardware-specific constant factors derived from profiling), the approach can accurately estimate the time needed to execute parallel tasks. An empirical evaluation showed that the approach can reduce the run-time overheads due to task creation to 3–13% of the sequential time and can obtain scalable speedups when running on multiple processors.

Provenance. Provenance—meta-information about the origin, history, or derivation of an object—is a central challenge in establishing trust and providing security in computer systems. Provenance is needed to provide satisfactory accountability, reproducibility, and trust for scientific and other high-value data. The group developed a core calculus for provenance in a higher-order functional language [2], going beyond special-purpose languages such as workflows and database queries, and using a notion of *trace* to explore aspects of provenance such as the disclosure and obfuscation problems.

15 The Synthesis, Analysis and Automated Reasoning Group

15.1 Overview

The Synthesis, Analysis and Automated Reasoning Group started in January 2012. The group existed until August 2013, when the group leader Ruzica Piskac left the institute to join Yale University as an assistant professor. The group is broadly interested in programming languages, software verification, automated reasoning, and code synthesis. A common thread in the group's research is improving software reliability and trustworthiness using formal techniques. During the reporting period, the group's activities have focused on decision procedures for synthesis and verification. More generally, the group investigated a translation from separation logic to a first-order logic fragment, the type inhabitation problem which was then applied to synthesis of code snippets, the reachability problem for the certain classes of well-structured transition systems, and simplification and redundancy techniques in first-order theorem provers.

Personnel. The group was led by Ruzica Piskac and had two graduate students (Filip Niksic and Susanne van den Elsen). Upon Ruzica's leaving, Filip and Susanne joined the Rigorous Software Engineering group led by Rupak Majumdar.

Collaborations. The group has joint publications with the Rigorous Software Engineering group led by Rupak Majumdar and with the Automation of Logic group at the Max Planck Institute for Computer Science (led by Christoph Weidenbach). Externally, the group has collaborated with researchers at the Lab for Automated Reasoning and Analysis, EPFL, Switzerland (led by Viktor Kuncak), IST Austria, and the Analysis of Computer Systems group at NYU.

Publications. The group has published in the top conferences and journals of its field. In the time frame 2012–2013 group members have coauthored two CAV [132, 174], one PLDI [114], and two journal [144, 143] publications. Additionally, the group has published a paper in the memorial volume dedicated to Harald Ganzinger [120].

Software, tools, and data. Ruzica is involved in the development of the InSynth tool. The tool is publicly available for download.

Teaching. Ruzica Piskac taught the seminar on Decision Procedures at Saarland University in the fall semester in 2012.

Invited talks and honors. Ruzica was an invited speaker at SYNT 2013, a workshop co-located with CAV 2013. Ruzica also lectured at the Winter School MOVEP 2012 as well as the Frühjahrsakademie Papenburg 2012, organized by the Studienstiftung des deutschen Volkes. The paper on complete functional synthesis [143] was invited and published as a "Research Highlight" in the *Communications of the ACM*. In 2012 Ruzica participated in the invitational workshop "Rising Stars in EECS: An Academic Career Workshop for Women" organized by MIT. For her doctoral thesis Ruzica was awarded the Patrick Denantes Prize in 2012 at EPFL.

Service. Internally to MPI-SWS, Ruzica served as a member of the Faculty Recruiting Committee for the 2013 hiring season, as well as a member of the Student Recruiting Committee in 2012 and 2013. Externally, Ruzica served on the program committees of VMCAI 2014, FMCAD 2013, the ESEC/FSE 2013 Tool Demonstrations Track, CSL'13, INFINITY 2012, PAAR-2012. Ruzica was also an external reviewer for the Handbook of Model Checking.

15.2 Research agenda

Although research into software correctness has a long history, software reliability is still difficult to obtain. The complexity of the problem is reflected in the fact that most tools developed for software verification focus on only one specific aspect of software reliability: some tools focus on proving termination, some tools focus on proving that the program corresponds to its specification, and some tools focus on detecting possible executions that will lead to an error state. One commonality among these tools is that they translate programs into mathematical formalisms.

In the last decade we have witnessed amazing progress in automated reasoning, i.e. the problem of automatically checking satisfiability of logical formulas. We refer to algorithms solving this problem as *decision procedures*. New insights into the problem and improved decision procedures have now culminated in solvers that scale up to industrial benchmarks. The goal of the Synthesis, Analysis and Automated Reasoning group was to investigate new decision procedures that can be applied in software verification. Additionally, our research was focused on new application areas for decision procedure, such as software synthesis or the coverability problem, as well as on new synergies of different formalisms.

Software Synthesis [144, 143, 114] Software synthesis is a rigorous, yet elegant approach to increase software reliability. The essence of software synthesis is that the programmer only states *what* should be done, and not *how* it should be done. Instead of writing code directly, the programmer provides a specification and the synthesis tool automatically generates code that satisfies this specification. Consequently, the generated code is correct by construction. Although first attempts at automated software synthesis began in the sixties, they were unsuccessful, primarily due to the lack of algorithms and systems that could handle such complex tasks. With solvers becoming more powerful, code synthesis has again shifted into the focus of current programming language research. Recently a number of synthesis tools have emerged that deploy solvers to generate code for non-trivial specifications. However, these tools still offer a limited degree of automation. They use solvers as a black-box for the extensive search of the program space. Since the search space is very large, the user still needs to provide additional hints such as the basic control structure of the desired code.

The goal of our research was to overcome these limitations by devising new synthesis algorithms that integrate automated reasoning on a foundational level. That is, instead of using solvers as a black box, we developed new decision procedures that are dedicated to solving synthesis problems. Instead of searching the program space, these procedures directly construct code from the specification.

We developed two synthesis tools based on such procedures: *Comfusy* [144, 143] and *InSynth* [114]. These tools have a high degree of automation: in particular users do not need to provide any hints other than the specification. Comfusy takes as input a specification that can contain program variables, whose exact value is not known at compile time, and returns a code fragment that satisfies the specification. Furthermore, it also returns a program assertion expressing preconditions that those input variables have to satisfy so that there is a solution program. InSynth can be thought of as an improved version of the autocompletion feature that is implemented in many code editors. By invoking InSynth, the user asks our tool to suggest a list of suitable code fragments for the given program point. InSynth displays a ranked list of suggested code snippets for that program point and the user can choose the best solution. Unlike previous tools, we use type constraints to derive complex code snippets, i.e. the specification is implicitly derived from the program context. Finding a code snippet of the right type is closely related to the type inhabitation problem, and our algorithm can be seen as a new constructive algorithm to address this important problem in type theory. In a usability study we showed that the desired snippet was among top 5 suggested solutions in 94% of the test cases. Both Comfusy and InSynth are integrated in the programming language Scala. In fact, Typesafe, the company behind Scala, has expressed interest in making InSynth a part of the standard Scala distribution.

An IC3-style Algorithm for the Coverabilty Problem [132] There are various well known formalisms of concurrent and distributed systems, such as Petri nets, lossy channel systems, dynamic process networks, etc. All those formalism belong to the class of so called *well-structured transition systems (WSTS)*.

An important question in concurrent systems is whether two processes can access the critical section at the same time. This problem belongs to the class called safety properties and in WSTS the verification of a large class of safety properties can be reduced to the coverability problem. The coverability problem is decidable for most WSTS but of very high complexity (e.g., the coverabilty problem for Petri nets is EXPSPACE-complete). As the verification of concurrent programs becomes increasingly important, we are also witnessing an increased interest in tackling the coverability problem for WSTS. In the past few years, many new results were published at top conferences and journals, leading to new tools for solving the coverabilty problem.

Around the same time, Aaron Bradley developed a new algorithm, called IC3, for checking safety properties of hardware designs. This algorithm provided new insights on the problem of hardware verification and significantly improved the techniques and tools in this problem domain.

Based on the IC3 algorithm we developed a new algorithm and a tool for checking coverability problems for WSTS. The original IC3 algorithm deals with finite state circuits, while WSTS are used for modeling infinite-state systems. Therefore, the main theoretical challenge for this new approach was an extension of the IC3 algorithm to infinite state systems. while still guaranteeing that the algorithm terminates. We also implemented the algorithm, with the specific focus on checking coverability of Petri nets. Running the experiments on the standard Petri net benchmarks, we demonstrated that our tool is competitive with state-of-the-art implementations for coverability checking, both in time and space usage.

Separation Logic Reasoning Using SMT Solvers [174] Separation logic is well known formalism for verification of heap-manipulating programs. Separation logic (SL) has gained widespread popularity because of its ability to succinctly express complex invariants of a programs heap configurations. Several specialized provers have been developed for decidable SL fragments. However, these provers cannot be easily extended or combined with solvers for other theories that are important in program verification, e.g., linear arithmetic.

To make SL amenable for combinations with other theories, we developed a reduction of decidable SL fragments to a decidable first-order theory that fits well into the satisfiability modulo theories (SMT) framework. Using this reduction we are able to automatically reason using SMT solvers about most of relevant problems in separation logic such satisfiability, entailment, frame inference, and abduction problems.

Our algorithm integrates separation logic into existing verification tools that provide SMT backends. We implemented this approach in a verification tool and applied it to heap-manipulating programs whose verification involves reasoning in theory combinations. What makes our tool unique is its decidable specification language, which supports mixing of assertions expressed in separation logic and first-order logic. The user of the tool can thus take advantage of the succinctness of separation logic specifications and the discipline of local reasoning. Yet, at the same time, the user can revert to classical logic in the cases where decidable separation logic fragments are less suited, such as reasoning about complex constraints on data and heap structures with complex sharing.

Redundancy and Simplification for FOL Theorem Provers [120] Harald Ganzinger's contributions to the research on first-order theorem proving are significant: he formulated and proved the correctness of the theory which is used as the basis of modern resolution-based first-order logic(FOL) theorem provers. FOL theorem provers try to construct a proof of unsatisfiabilty by deriving fresh clauses. As the number of newly derived clauses can be very large, one of the key contributions of Harald's work were his simplification and redundancy techniques for cutting down the search space. In particular, transitive predicates have to be carefully handled, since they can easily lead to a large number of freshly derived clauses, most of which are irrelevant for the proof search. To solve this problem, Harald introduced the chaining calculus for efficient reasoning about the transitive predicates. He implemented this calculus in his experimental theorem prover called Saturate.

We used the Saturate system for the formal verification of a checker for priority queues. We structured the correctness proof into roughly a hundred separate lemmas. Each of these lemmas we proved using Saturate and the state-of-the-art prover SPASS. Saturate was an experimental tool, and it was outperformed on most of the tasks by SPASS. However, there were lemmas where transitive relations played an important role. One of the main operations in a priority queue is accessing the minimal element. The ordering on the elements is a transitive relation. We experimentally showed that the simplification and redundancy techniques implemented in Saturate are important for the performance of a theorem prover: in those lemmas where the chaining calculus could be applied, Saturate significantly outperformed SPASS by several orders of magnitude, in the most extreme example Saturate proved one of the lemmas in four seconds, while SPASS needed one hour.

Part IV
Details

16 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.

16.1 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.

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

Leadership 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 Rupak Majumdar), assisted by the head of the administrative department, Volker 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 Geiss. The department provides personnel, finance, and purchasing services. Volker Geiss also handles much of the public relations, relations with local governments, and relations with other research institutions in Kaiserslautern and Saarbrücken. As of 2012, the core IT support group we share with the MPI for Informatics, and the institute's own user-facing IT support team report to Geiss as well.

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 Laubscher and Susanne Girard 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

Group Name	Group Leader	Start	Status	Retirement
Real-Time Systems	Brandenburg	2011	tenure-track	—
Robust Systems	Clement	2012	tenure-track	—
Social Information Systems	Danescu-Niculescu-Mizil	2013	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	—
Programming Languages and Systems	Acar	2010	left in 2012	—
Synthesis, Analysis and Automated Reasoning	Piskac	2012	left in 2013	—

Figure 1: MPI-SWS research groups

provided by a team headed by Jörg Herrmann. This team (currently 7 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 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 hired four such developers, Michael Ohlmann (private advertising project), Sebastian Probst-Eide and Matthias Kretschmer (private analytics), and William Caldwell (network anonymity). We have also used a fifth developer, Jeff Hoye, on a consulting basis for both research projects and institute administrative tools such as our admissions system.

16.2 Research program and groups

This information is provided in Sections 2–8.

16.3 Personnel structure

Currently, the institute has 81 members (excluding interns and visitors). Among these, there are 58 researchers versus 23 non-research staff. During the reporting period, 83 members joined MPI-SWS, and 51 departed. The percentage of women among the MPI-SWS members is 21.0% (17 female, 65 male). Among researchers, the percentage is only 13.8% (8 female, 50 male).

Eight researchers are covered by external funding.

The ratio of institute members with permanent versus temporary contracts is as follows (excluding interns and visitors):

Permanent faculty	5
Tenure-track faculty	5
Permanent staff	11
Temporary contracts (incl. 1 Fellow and 4 undergrad. assistants)	14
Postdocs	7
PhD students	39

16.4 Structure of the budget

The institute's total yearly budget is EUR 10.51M per year. Of that, the institute's yearly expenditure (as of December 2012) amounted to EUR 7.14M per year, including EUR 2.21M for material expenses, EUR 1.86M special financing for basic scientific equipment, EUR 139K for investment in major equipment, EUR 2.21M for personnel expenses (excluding stipends) and EUR 834K for graduate and postdoctoral stipends. (Personnel funds can be used to fund additional stipends but not vice versa.)

The institute has 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 a fraction of its full budget, since only 10 faculty positions have been filled (including 3 director positions).

16.5 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 40TB 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 two clusters for research. One cluster has sixty single- and dual-core opteron systems that are connected to the institute's intranet and have direct access to the storage services. A smaller experimental cluster of ten isolated opteron nodes is outside the security perimeter and has an open connection to the Internet connection for networking research. 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 During the review period, construction of the institute buildings in both Saarbrücken and Kaiserslautern was completed, and we have now comfortably moved into these new buildings. This is a marked improvement over our previous situation, especially in Saarbrücken, where we were previously located off-campus (in the Wartburg building). We are very happy to return to campus, where we can engage fully and interactively with our colleagues in the Saarbrücken Graduate School.

Overall, the transition to our new buildings has been viewed very positively by the members of the institute. The new buildings are light and spacious, encouraging transparent and open communication between researchers within each location and between the two locations. A number of lecture rooms and meeting rooms in both locations are equipped with state-of-the-art videoconferencing technology, which we use regularly to enhance inter-site communication and collaboration, as well as to teach courses that can be cross-listed between the two sites.

16.6 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.

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 seven postdocs from six countries. A list of our current and past 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 have had good success in attracting a talented and diverse set of students, using very proactive recruiting. (We currently have 39 doctoral students from sixteen 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 has hosted 14 undergraduate interns and 36 graduate interns from 13 countries. A full list of current and past interns can be found online at `http://www.mpi-sws.org/index.php?n=people&s=function&c=interns`.

Faculty members from other institutions frequently come for research visits. Recent short-term visitors include Flavio Junqueira (Yahoo!), Yuri Gurevich (Microsoft Research Redmond), Uday Khedkar (IIT Bombay), Animesh Kumar (IIT Bombay), Amey Karkare (IIT Kanpur), Aditya Kanade (IISc, Bangalore), Thomas Wies (NYU), Pierre Ganty (IMDEA Software Institute), Michael Emmi (IMDEA Software Institute), Ahmed Bouajjani (LIAFA), Javier Esparza (TU Munich), Jyotirmoy Deshmukh (Toyota), Jean-Francois Raskin (Free Univ. Brussels), Daniel Kroening (Oxford), Ram Sewak Sharma (Unique Identification Authority of India), Claude Castelluccia (INRIA), Petr Kuznetsov (Telecom ParisTech), Y. Charlie Hu (Purdue), Winter Mason (Yahoo!), Stefan Saroiu (Microsoft Research Redmond), Manuel Gomez Rodrigues (MPI for Intelligent Systems), Balachander Krishnamurthy (AT&T), Ratul Mahajan (Microsoft Research Redmond), Alice Oh (KAIST), Alan Mislove (Northeastern), Kevin Elphinstone (Univ. New South Wales), Doug Terry (Microsoft Research Silicon Valley), Daniel Peek (Facebook), Alexey Gotsman (IMDEA Software Institute), Uday Reddy (Birmingham), Neal Glew (Intel Labs), Martin Hofmann (Ludwig-Maximilians-Universität München), Andrew Pitts (Cambridge), Andrej Bauer (University of Ljubljana), Amal Ahmed (Northeastern), Brigitte Pientka (McGill), and Peter Müller (ETH Zürich).

We have also had several longer-term visitors, including Lorenzo Alvisi (UT Austin), Meeyoung Cha (KAIST), Bryan Ford (Yale), Niloy Ganguly (IIT Kharagpur), Johannes Gehrke (Cornell), and Robert Harper (CMU).

16.7 Equal opportunity

Ensuring gender diversity is a well-known perennial problem in computer science departments worldwide. By the end of our past review period, MPI-SWS was no exception: only 2 out of 48 researchers at the institute—a mere 4%—were women.

In the past review period, however, we have turned the situation around considerably. Of the researchers who joined the institute since May 2011 (i.e. during the present review period), 6 out of 26 doctoral students, 1 out of 9 postdocs, and 1 out of 5 faculty were women. All in all, 20% of the researchers who joined the institute during this review period are women. The proportion of new female doctoral students (23%) is particularly encouraging, given that the average proportion of female doctoral students in CS at German universities is around 12.9%.

Unfortunately, Ruzica Piskac, our new female faculty member who joined the institute in January 2012, ended up leaving earlier this year for an assistant professor position at Yale. With Piskac's departure, the institute currently has one female postdoc, six female doctoral students, one female Master's student, and one female research support staff member. We will continue our efforts to increase the representation of women among its research staff and student population, and in particular to actively recruit excellent female faculty.

16.8 Relations with domestic and foreign research institutions

Domestic The institute seeks to conclude cooperation agreements with the TU Kaiserslautern and Saarland University. The goal is to provide a framework for a close cooperation among the institute and the CS departments at both universities. The centerpiece of such a cooperation would be a joint graduate program. We would like to have all MPI-SWS faculty members appointed as adjunct professors. Currently, we are gaining experience with a joint graduate program. We hope that we can agree on a cooperation agreement within the next two years.

At this point, we are working towards integration with the universities in several directions. Professors from the two departments are present in our faculty recruitment committees and our graduate student admission committees. MPI-SWS faculty have taught courses in their area of expertise. There are some joint research projects (e.g., Garg and Hammer (UdS), Majumdar and Meyer (TUKL), Francis and Nebel (TUKL)). Druschel is a PI in CISPA and the Excellence Cluster at UdS. We have also engaged with TU KL in their faculty recruitment efforts, providing additional resources in certain cases to help recruit excellent faculty.

During this reporting period, faculty at MPI-SWS have taught the following courses:

- Reactive Systems Verification, TU Kaiserslautern, 2011
- Selected Topics in Information Security, Saarland University, 2011
- Concurrent Program Logics, Saarland University, 2011
- Cryptography and Security, Saarland University, 2011 and 2012
- Operating Systems, Saarland University, 2011 and 2013
- Distributed Systems, Saarland University and TU Kaiserslautern, 2012
- Hot Topics in Security and Privacy, Saarland University, 2012
- Introductory Proof Theory, Saarland University, 2012
- Decision Procedures, Saarland University, 2012
- Real-Time Scheduling and Synchronization, Saarland University, 2012
- Parametricity and Modular Reasoning, Saarland University, 2012
- Social Computing Systems, Saarland University, 2013
- Fault-Tolerant Distributed Real-Time Systems, Saarland University, 2013
- Operating System Design and Implementation, Saarland University, 2013

The institute is part of the “Excellence Cluster on Multi-Model Computing and Interaction” and the “Saarbrücken Graduate School of Computer Science” in Saarbrücken, which were awarded by the German federal government as part of its Initiative for Excellence in 2008 and renewed in 2012, and which provide funds of about EUR 85M over ten years. (Virtually all of these funds go to Saarland University.) In addition, the institute participates in the Intel Visual Computing Institute (Intel VCI) in Saarbrücken, which is a collaborative effort between Intel, Saarland University, DFKI, MPI-SWS, and the Center for IT Security, Privacy and Accountability (CISPA). The institute is also part of the Science Alliance Kaiserslautern.

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

Universities Aarhus University, Birmingham, Carnegie Mellon University, Cornell University, Duke University, EPFL, ETHZ, Federal University of Uberlândia (Brazil), Grenoble Institute of Technology, IIT Bangalore, IIT Kharagpur, ITU-Copenhagen, Indiana, KAIST, Karlsruhe Institute of Technology, La Sapienza (University of Rome), Leuven, NYU, Northeastern University, NOVA University of Lisbon, Purdue, Rochester, Saarland University, Stanford University, Stevens Institute of Technology, TU Delft, TU Kaiserslautern, TU Munich, Technische Universität Darmstadt, UC Los Angeles, UC San Diego, UFMG, UNC Chapel Hill, Tel Aviv University, University of Bucharest, University of Cambridge, University of Chicago, University of Edinburgh, University of Helsinki, University of Illinois at Urbana-Champaign, University of Luxembourg, University of Maryland, University of Pennsylvania, University of Texas at Austin, University of Torino, University of Washington, Waterloo, and Yale University.

Research Institutes Microsoft Research, IBM Research, MPI-INF, AT&T, IMDEA, INRIA, and IST Austria.

Industry Akamai, Facebook, Google, SYSGO AG, and Toyota.

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

16.9 Activities regarding the transfer of knowledge/relations with industry

Francis' group has created a spinoff company based on its privacy research over the last four years. The initial service of the spinoff is private analytics. The spinoff is funded by the German government as part of a program called EXIST. The level of funding supports four people for 18 months. The spinoff is also supported by MPI-SWS in the form of office space, IT support, equipment, and some salary support early on. Currently the spinoff is still administratively part of the MPI-SWS, and is not a distinct legal entity. The Max Planck Society retains some equity and royalty rights, and a third of this is given to the MPI-SWS. The spinoff has been active since September 2012, and the EXIST funding started in October 2013.

The goal of the spinoff, however, goes well beyond tech transfer. Privacy is not a purely technical problem. It is influenced by economics, personal perceptions, national laws and reguations, and social norms. The CS research community does not have a good track record of developing privacy technologies that are useful in practice. We believe that this is in large part because the CS research community addresses privacy as a narrow technical problem.

The spinoff is being treated as a kind of research instrument by Francis' group. There is virtually no proprietary technology within the spinoff. This is because, as a privacy company that needs to build and maintain trust, the spinoff technology is completely open (to the point of remote attestation of the spinoff service through TPMs). To succeed, the spinoff must successfully navigate technical, economic, legal, and public perception problems. Our research thesis is that this is possible through 1) a combination of strong privacy mechanisms that never-the-less allow for good analytics utility, 2) completely transparent operation, and 3) enforceable oversight by privacy organizations. The spinoff itself is a long-term research experiment where the results and insights will be openly published.

Druschel's group has had a project involving Akamai Technologies, Inc., involving data and research problems from Akamai's content distribution network. The results had a direct impact on Akamai's service.

A number of groups maintain and distribute data sets, tools and other software artifacts for use by researchers and practitioners (and in some cases, the public) and encourage technology transfer to industry.

The institute is a founding member of the VICCI cloud computing research testbed. This testbed comprises seven compute clusters, four of them in US universities (Princeton, Stanford, Georgia Tech and the University of Washington), one in Japan (University of Tokyo) and an additional cluster in Europe (ETH Zurich).

16.10 Symposia, conferences, etc.

The institute organized the third institute retreat in June 2012 at the European Academy of Otzenhausen. 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, 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 26 distinguished lecturers: Dawn Song (UC Berkeley), Johannes Gehrke (Cornell University), Donald Kossmann (ETH Zürich), Ion Stoica (University of California, Berkeley), Jon Crowcroft (University of Cambridge, UK), Thomas Anderson (University of Washington), Geoffrey Voelker (University of California, San Diego), Adrian Perrig (Carnegie Mellon University), Alexander Smola (Yahoo! Research & UC Berkeley & ANU), Monica Lam (Stanford University), Val Tannen (University of Pennsylvania and EPFL), Gernot Heiser (NICTA, Kensington, Australia), Jan Vitek (Purdue University), Christopher Kruegel (University of California, Santa Barbara), Lorenzo Alvisi (University of Texas at Austin), Jeffrey Mogul (HP Labs, Palo Alto), Michael Hicks (University of Maryland), Timothy Roscoe (ETH Zurich), Nikolaj Bjorner (Microsoft Research Redmond), Vitaly Shmatikov (University of Texas, Austin), Anastasia Ailamaki (EPFL), David Walker (Princeton University), Divesh Srivastava (AT & T Labs), Peter Sewell (University of Cambridge), Frank McSherry (Microsoft Research Silicon Valley), and John Wilkes (Google).

Previous distinguished lecturers include Dan Suciu (UW), Robert Kraut (CMU), M. Angela Sasse (UCL), Yuri Gurevich (MSR), and Nick Benton (MSR), Lydia Kavraki (Rice), Bob Harper (CMU), Mike Ernst (MIT/UW), Gerard Berry (Esterel Technologies), Patrick Cousot (ENS), Rajeev Alur (UPenn), Maurice Herlihy (Brown), Byron Cook (MSRC), Tom Reps (Wisconsin), Rupak Majumdar (UCLA, now MPI-SWS), Tom Ball (MSR), Patrice Godefroid (MSR), Greg Morrisett (Harvard), Johannes Gehrke (Cornell), Paul Francis (Cornell, now MPI-SWS), Moshe Vardi (Rice), Andy Gordon (MSRC), Ed Lee (Berkeley), Jennifer Rexford (Princeton), Eric Brewer (Berkeley), Henning Schulzrinne (Columbia), Erik Sandewall (Linkoeping), Jean-Raymond Abrial (ETHZ), Matthias Felleisen (Northeastern), Anja Feldmann (TU Berlin/Deutsche Telekom Labs), Luca Cardelli (MSRC), and Rustan Leino (MSR). (A list of lecture abstracts and titles is available online at `http://www.mpi-sws.org/index_flash.php?n=lectures/dlseries/program`.) We feel that 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.

16.11 Committee work of the faculty

MPI-SWS researchers have served on the program committees of over 100 conferences and workshops, and have chaired or co-chaired the PCs of four conferences and five workshops. This information is provided in detail in the individual research group sections.

In this section, we document the steering committee and advisory/editorial board work that institute members have done.

Francis serves as the MPI-SWS representative to Science Alliance, an organization consisting of representatives of local universities and research laboratories.

Dreyer serves as the elected staff representative for MPI-SWS in the Chemistry, Physics & Technology (CPT) Section of the Max Planck Society, which entails participation in CPT Section meetings three times a year. He was also elected to the position of member-at-large on the ACM SIGPLAN Executive Committee (since July 2012), where he serves as "awards chair."

Backes currently serves on the steering committee of IEEE S&P, IEEE CSF and ESORICS, and he is an Associate Editor of Springer's International Journal on Information Security and Cryptography.

Gummadi is a founding member and steering committee member of M-Lab, a measurement lab founded by the New America Foundation's Open Technology Institute, the PlanetLab Consortium, Google Inc. and academic researchers. He also serves on the technical advisory board of the newly formed ACM conference on online social networks (COSN).

Druschel has served on the Technical Advisory Boards of Microsoft Research, Cambridge since 2011 and Microsoft Research, Bangalore, since 2013. He has served on the scientific committee of the Laboratory on Information, Networking and Communication Sciences (LINCS), Paris, and on the steering committee of the ACM SIGOPS Asia-Pacific Workshop on Systems (APSys) since 2009, and on the the steering committee of the EuroSys/INRIA Winter School on Hot Topics in Distributed Systems (HTDC) since 2008.

He served on the SIGOPS EuroSys steering committee through 2012. In addition, Peter serves on the editorial board of the ACM Communications of the ACM (CACM).

16.12 Publications

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

During the reporting period, the institute produced 177 peer-reviewed conference, workshop, and journal publications, of which 20 are collaborative across research groups: [69, 60, 70, 173, 116, 3, 150, 2, 5, 42, 61, 71, 190, 43, 4, 6, 29, 12, 24, 28, 20, 35, 18, 33, 32, 19, 177, 13, 31, 17, 26, 34, 25, 14, 22, 204, 27, 156, 36, 23, 21, 30, 15, 16, 129, 52, 51, 112, 53, 206, 55, 59, 54, 151, 128, 75, 10, 168, 81, 80, 176, 37, 123, 138, 125, 195, 109, 135, 140, 89, 213, 193, 136, 90, 124, 116, 137, 74, 139, 170, 86, 110, 85, 94, 182, 41, 130, 145, 183, 212, 146, 168, 8, 202, 115, 58, 88, 66, 146, 9, 179, 67, 196, 68, 101, 83, 100, 103, 140, 169, 102, 104, 175, 142, 97, 152, 181, 203, 141, 133, 48, 105, 134, 189, 164, 165, 185, 191, 82, 209, 163, 208, 161, 184, 167, 131, 162, 63, 210, 93, 132, 126, 127, 166, 121, 160, 147, 92, 56, 108, 154, 73, 62, 72, 65, 64, 98, 99, 113, 192, 132, 174, 114, 144, 143, 120, 47, 45, 205, 57, 75, 187, 151, 186, 46, 180, 155, 188, 123, 125, 199, 213, 118, 198, 124, 197, 200, 211].

16.13 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.

16.14 Appointments, scientific awards and memberships

Ruzica Piskac's paper on complete functional synthesis [143] was invited and published as a "Research Highlight" in the *Communications of the ACM*. She was also awarded the Patrick Denantes Prize in 2012 for her doctoral dissertation at EPFL.

Cristian Danescu-Niculescu-Mizil's work won the best paper award at WWW 2013 [81] and was nominated for the best paper award at ACL 2013 [176].

Björn Brandenburg's ECRTS'13 paper [112] was awarded an "outstanding paper award" and his SIES'13 paper [206] and EMSOFT'11 papers both won the "best paper" award at their respective conferences. In addition, his dissertation was recognized with three dissertation awards: UNC Chapel Hill's 2012 *Dean Linda Dykstra Distinguished Dissertation Award* in the category "Mathematics, Physical Sciences and Engineering," the 2012 *Council of Graduate Schools / ProQuest Distinguished Dissertation Award* in the category "Mathematics, Physical Sciences and Engineering," and the 2012 *Distinguished Dissertation Award* of the European Design and Automation Association in the area "New directions in embedded systems design and embedded software."

Peter Druschel won the 2011 ACM/IFIP/Usenix Middleware 10-year Best Paper Award, together with Ant Rowstron (Microsoft Research, Cambridge).

Paul Francis was awarded the 2011 SIGCOMM Test of Time award, together with his co-authors Mark Handley, Richard Karp, Sylvia Ratnasamy, and Scott Shenker.

Krishna Gummadi received the Best Paper Award at the AAAI's ICWSM 2012 conference [133].

Rupak Majumdar was awarded the 2011 SIGPLAN PLDI Test of Time award, together with co-authors Tom Ball, Todd Millstein, and Sriram Rajamani. He won the EAPLS Best Paper award at the 2012 ETAPS conference (for his FASE 2012 paper [154],) and the 2011 ACM TODAES best paper award (for the best paper in all volumes of ACM Transactions on Design Automation of Embedded Systems in the year 2011) [76].

Michael Backes was named the leading German scientist under the age of 40 by the Financial Times Germany, and one of the 100 most important IT people in Germany in 2011 by the Computerwoche Newspaper.

16.15 External funding

During the present review period, two faculty—Umut Acar and Rodrigo Rodrigues—successfully obtained ERC Starting Grants. Acar also obtained a GIF grant and Rodrigues was awarded an Amazon Web Services in Education Research Grant. Unfortunately, both of them left the institute (for positions at CMU and NOVA University of Lisbon, respectively) soon after obtaining their grants, but we are nonetheless very proud of their accomplishments.

Derek Dreyer secured two industry fellowships: a 2012 Google European Doctoral Fellowship for his student Georg Neis, and a 2013 Microsoft Research PhD Scholarship for his student David Swasey. Each provides funding for a doctoral student for 3 years.

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

(MPI for Informatics) are co-PIs on an ERC Synergy Grant proposal, which has made it to the final round of competition, with a reverse site visit taking place in Brussels on Nov 5, 2013.

Viktor Vafeiadis was awarded a 2013 young explorer ADVENT grant from the European Commission's FP7 FET.

Deepak Garg (along with Christian Hammer from the University of Saarland) secured funding from the German Research Foundation, Deutsche Forschungsgemeinschaft (DFG).

Majumdar's group has been awarded industrial grants from Intel and from Toyota.

Gummadi's group won a Google European Doctoral Fellowship as well as an Anita Borg Scholarship. Ezgi Cicek, currently a student in Deepak Garg's group, and formerly in Umut Acar's group, won an Anita Borg Scholarship from Google in 2012.

Gummadi's and Majumdar's groups also receive funding under an IMPECS grant (Indo-Max Planck Center for Computer Science) for research collaborations with IITs in India.

The research of Michael Backes's group has been partially funded by an ERC starting grant on end-to-end security, by the Excellence Cluster on Multimodal Computing and Interaction, as well as by the newly founded center for IT security, privacy, and accountability (CISPA).

Druschel is a co-PI in the successful renewal of both Saarland University's MMCI Cluster of Excellence and the Saarbrúcken 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).

Paul Francis's group was awarded an EXIST startup grant from the German Government (Bundesministerium für Wirtschaft und Technologie) for roughly EUR500K over 18 months.

16.16 Public relations work

Articles describing Michael Backes's group's work on AppGuard (protecting against malicious apps on Android) have appeared in numerous popular news media (Frankfurter Allgemeine Zeitung, Sueddeutsche Zeitung, C't, Heise, various TV news programs, etc.)

Cristian Danescu-Niculescu-Mizil's research on language and social computing has been featured in popular-media outlets such as the New Scientist, NBC's The Today Show, NPR and the New York Times. Cristian's opinion on external social computing research was solicited and quoted by ABC News.

Krishna Gummadi's research has been covered in numerous popular news media and technology blogs worldwide including the New York Times, Harvard Business Review, MIT Technology Review, New Scientist, Wired magazine, Slashdot, Businessweek, MaxPlanck Research (Germany), Sueddeutsche Zeitung (Germany), Science TV (Korea), and MTV (Brazil).

Stevens Le Blond's work on security flaws in Skype and other peer-to-peer applications has also received global media attention, with articles appearing in the Wall Street Journal, Le Monde (France), die Zeit (Germany), Daily Mail, Slashdot, Wired, and New Scientist.

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

- An open house held in cooperation with the Max Planck Institute for Informatics and the Department of Computer Science at Saarland University.
- Together with the TU Kaiserslautern, the Institute participated in the "Long Night of the Sciences" with a scientific presentation in layman's language. As part of this event, members of the Institute and 600 visitors from the Palatinate region also took part in an activity that demonstrated the close connection between science and sports.
- Participation in a national computer science competition with an award presentation ceremony in Kaiserslautern.
- Participation in the media project for Rhineland-Palatinate as a business location.
- Publication of articles and book chapters, with the goal of making the Institute well known to the broad public:
  - "Innovative Saarland"
  - "*Research* Magazine"
  - "Business Location Rhineland-Palatinate"
  - Monograph "Kaiserslautern – Region with a Promising Future"
  - Monograph "Saarland"
  - TU Kaiserslautern research brochures "Research and International Partnership"

- Co-organizer of the coordination of the Journalist Prize in Computer Science, a prize for the best scientific articles by journalists.
- Presentation of Science prizes, together with the Minister of Science.
- Co-organizer and co-coordinator for the kick-off event for the EU framework program Horizon 2020 by the Saarland state government at the institute building in Saarbrücken.
- Presentation for the official visit of the Consul General of India at the institute.

References

[1] Aircloak Analytics: Anonymized User Data without Data Loss. `http://aircloak.com/resources/white-paper-1.1.pdf`.

[2] U. A. Acar, A. Ahmed, J. Cheney, and R. Perera. A core calculus for provenance. In *Proceedings of the Conference on Principles of Security and Trust*, pages 410–429. Springer, 2012.

[3] U. A. Acar, A. Charguéraud, and M. Rainey. Oracle scheduling: Controlling granularity in implicitly parallel languages. In *Proceedings of the 26th ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, pages 499–518. ACM, Oct. 2011.

[4] U. A. Acar, A. Charguéraud, and M. Rainey. Efficient primitives for creating and scheduling parallel computations. In *Proceedings of the Workshop on Declarative Aspects and Applications of Multicore Programming (DAMP)*, Jan. 2012.

[5] U. A. Acar, A. Charguéraud, and M. Rainey. Scheduling parallel programs by work stealing with private deques. In *Proceedings of the 18th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 219–228. ACM, Feb. 2013.

[6] U. A. Acar and Y. Chen. Streaming big data with self-adjusting computation. In *Proceedings of the Workshop on Data-Driven Functional Programming*, pages 15–18. ACM, 2013.

[7] P. Aditya, V. Erdelyi, M. Lentz, , E. Shi, B. Bhattacharjee, and P. Druschel. EnCore: Private, Context-based Communication for Mobile Social Apps. Submitted for publication.

[8] P. Aditya, M. Zhao, Y. Lin, A. Haeberlen, P. Druschel, B. Maggs, and B. Wishon. Reliable Client Accounting for Hybrid Content-Distribution Networks. In *Proceedings of the 9th USENIX Symposium on Networked Systems Design and Implementation (NSDI'12)*, Apr. 2012.

[9] I. E. Akkus, R. Chen, M. Hardt, P. Francis, and J. Gehrke. Non-tracking web analytics. In *ACM Conference on Computer and Communications Security*, pages 687–698, 2012.

[10] L. Alvisi, A. Clement, A. Epasto, S. Lattanzi, and A. Panconesi. SoK: The Evolution of Sybil Defense via Social Networks. *2012 IEEE Symposium on Security and Privacy*, 0:382–396, 2013.

[11] A. Anta, R. Majumdar, I. Saha, and P. Tabuada. Automatic verification of control system implementations. In *EMSOFT 10: Embedded Software*. ACM, 2010.

[12] M. Backes, G. Barthe, M. Berg, B. Grégoire, C. Kunz, M. Skoruppa, and S. Z. Béguelin. Verified Security of Merkle-Damgård. In *Proc. 25th IEEE Computer Security Foundations Symposium (CSF)*, pages 354–368, 2012.

[13] M. Backes, F. Bendun, and D. Unruh. Computational Soundness of Symbolic Zero-Knowledge Proofs: Weaker Assumptions and Mechanized Verification. In *Proc. 2nd International Conference on Principles of Security and Trust (POST)*, pages 206–225, 2013.

[14] M. Backes, M. Berg, and B. Köpf. Non-uniform distributions in quantitative information-flow. In *Proc. 6th ACM Symposium on Information, Computer and Communications (ASIACCS)*, pages 367–375, 2011.

[15] M. Backes, A. Busenius, and C. Hritcu. On the development and formalization of an extensible code generator for real life security protocols. In *Proc. 4th International NASA Symposium on Formal Methods*, pages 371–387, 2012.

[16] M. Backes, I. Cervesato, A. D. Jaggard, A. Scedrov, and J.-K. Tsay. Cryptographically sound security proofs for basic and public-key Kerberos. *International Journal of Information Security*, 10(2):107–134, 2011.

[17] M. Backes, A. Datta, and A. Kate. Asynchronous Computational VSS with Reduced Communication Complexity. In *Proc. RSA Conference – The Cryptographers' Track (CT-RSA)*, pages 259 – 276, 2013.

[18] M. Backes, G. Doychev, and B. Köpf. Preventing Side-Channel Leaks in Web Traffic: A Formal Approach. In *Proc. 20th Annual Network and Distributed System Security Symposium (NDSS)*, 2013.

[19] M. Backes, D. Fiore, and E. Mohammadi. Privacy-Preserving Accountable Computation. In *Proc. 18th European Symposium on Research in Computer Security (ESORICS)*, pages 38–56, 2013.

[20] M. Backes, D. Fiore, and R. M. Reischuk. Verifiable Delegation of Computation on Outsourced Data. In *Proc. 20th ACM Conference on Computer and Communications Security (CCS)*, 2013.

[21] M. Backes, M. Gagné, and M. Skoruppa. Using Mobile Device Communication to Strengthen E-voting Protocols. In *Proc. 12th ACM Workshop on Privacy in the Electronic Society (WPES)*, 2013.

[22] M. Backes, S. Gerling, C. Hammer, M. Maffei, and P. von Styp-Rekowsky. AppGuard – Enforcing User Requirements on Android Apps. In *Proc. 19th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pages 543–548, 2013.

[23] M. Backes, S. Gerling, C. Hammer, M. Maffei, and P. von Styp-Rekowsky. AppGuard – Fine-grained Policy Enforcement for Untrusted Android Applications. In *Proc. 8th International Workshop on Data Privacy Management (DPM)*, 2013.

[24] M. Backes, I. Goldberg, A. Kate, and E. Mohammadi. Provably Secure and Practical Onion Routing. In *Proc. 25th IEEE Computer Security Foundations Symposium (CSF)*, pages 369–385, 2012.

[25] M. Backes, I. Goldberg, A. Kate, and T. Toft. Adding query privacy to robust DHTs. In *Proc. 7th ACM Symposium on Information, Computer and Communications (ASIACCS)*, pages 30–31, 2012.

[26] M. Backes, C. Hritcu, and M. Maffei. Union and Intersection Types for Secure Protocol Implementations. In *Joint Workshop on Theory of Security and Applications (TOSCA)*, pages 1–28, 2011.

[27] M. Backes, C. Hritcu, and T. Tarrach. Automatically Verifying Typing Constraints for a Data Processing Language. In *Proc. 1st International Conference on Certified Programs and Proofs (CPP)*, pages 296–313, 2011.

[28] M. Backes, A. Kate, M. Maffei, and K. Pecina. ObliviAd: Provably Secure and Practical Online Behavioral Advertising. In *Proc. 33rd IEEE Symposium on Security and Privacy (S& P)*, pages 257–271, 2012.

[29] M. Backes, A. Kate, P. Manoharan, S. Meiser, and E. Mohammadi. AnoA: A Framework for Analyzing Anonymous Communication Pro-

tocols. In *Proc. IEEE 26th Computer Security Foundations Symposium (CSF)*, pages 163 – 178, 2013.

[30] M. Backes, A. Kate, and E. Mohammadi. Ace: an efficient key-exchange protocol for onion routing. In *Proc. 11th Annual ACM Workshop on Privacy in the Electronic Society (WPES)*, pages 55–64, 2012.

[31] M. Backes, A. Kate, and A. Patra. Computational Verifiable Secret Sharing Revisited. In *Proc. 17th International Conference on Theory and Application of Cryptography and Information Security (ASIACRYPT)*, pages 590–609, 2011.

[32] M. Backes, M. Maffei, and K. Pecina. A Security API for Distributed Social Networks. In *Proc. 18th Network and Distributed System Security Symposium (NDSS)*, 2011.

[33] M. Backes, M. Maffei, and K. Pecina. Automated Synthesis of Secure Distributed Applications. In *Proc. 19th Annual Network and Distributed System Security Symposium (NDSS)*, 2012.

[34] M. Backes, M. Maffei, K. Pecina, and R. M. Reischuk. G2C: Cryptographic Protocols from Goal-Driven Specifications. In *Joint Workshop on Theory of Security and Applications (TOSCA)*, pages 57–77, 2011.

[35] M. Backes, A. Malik, and D. Unruh. Computational soundness without protocol restrictions. In *Proc. 19th ACM Conference on Computer and Communications Security (CCS)*, pages 699–711, 2012.

[36] M. Backes and S. Meiser. Differentially Private Smart Metering with Battery Recharging. In *Proc. 8th International Workshop on Data Privacy Management (DPM)*, 2013.

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

[38] T. Ball, R. Majumdar, T. Millstein, and S. K. Rajamani. Automatic predicate abstraction of C programs. In *PLDI 01: Programming Languages Design and Implementation*, pages 203–213. ACM, 2001.

[39] S. Baruah and B. Brandenburg. Multiprocessor feasibility analysis of recurrent task systems with specified processor affinities. In *Pro-*

ceedings of the 34th IEEE Real-Time Systems Symposium, 2013 (to appear).

[40] N. Benton, L. Cardelli, and C. Fournet. Modern concurrency abstractions for C$^\sharp$. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 26(5):769–804, 2004.

[41] N. Benton, C.-K. Hur, A. Kennedy, and C. McBride. Strongly typed term representations in Coq. *Journal of Automated Reasoning*, 49(2):141–159, Aug. 2012.

[42] L. Bergstrom, M. Fluet, M. Rainey, J. Reppy, S. Rosen, and A. Shaw. Data-only flattening for nested data parallelism. In *Proceedings of the 18th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 81–92. ACM, 2013.

[43] L. Bergstrom, M. Fluet, M. Rainey, J. Reppy, and A. Shaw. Lazy tree splitting. *Journal of Functional Programming*, 22:382–438, Sept. 2012.

[44] K. Bhargavan, C. Fournet, and A. D. Gordon. Modular verification of security protocol code by typing. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, 2010.

[45] P. Bhatotia, R. Rodrigues, and A. Verma. Shredder: GPU-accelerated incremental storage and computation. In *Proceedings of the 10th USENIX conference on File and Storage Technologies*, FAST'12, Berkeley, CA, USA, 2012. USENIX Association.

[46] P. Bhatotia, A. Wieder, I. E. Akkus, R. Rodrigues, and U. A. Acar. Large-scale Incremental Data Processing with Change Propagation. In *Proceedings of the 3rd USENIX conference on Hot topics in cloud computing (HotCloud'11)*, Portland, OR, June 2011.

[47] P. Bhatotia, A. Wieder, R. Rodrigues, U. A. Acar, and R. Pasquin. Incoop: MapReduce for incremental computations. In *Proceedings of the 2nd ACM Symposium on Cloud Computing*, SOCC '11, pages 7:1–7:14, New York, NY, USA, 2011. ACM.

[48] 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.

[49] B. Blanchet. An efficient cryptographic protocol verifier based on Prolog rules. In *IEEE Computer Security Foundations Workshop (CSFW)*, 2001.

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

[51] 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.

[52] B. Brandenburg. Improved analysis and evaluation of real-time semaphore protocols for P-FP scheduling. In *Proceedings of the 19th IEEE Real-Time and Embedded Technology and Applications Symposium*, pages 141–152, 2013.

[53] B. Brandenburg and J. Anderson. Real-time resource-sharing under clustered scheduling: mutex, reader-writer, and k-exclusion locks. In *Proceedings of the 9th ACM International Conference on Embedded Software*, 2011.

[54] B. Brandenburg and J. Anderson. The OMLP family of optimal multiprocessor real-time locking protocols. *Design Automation for Embedded Systems*, online first, DOI 10.1007/s10617-012-9090-1, July 2012.

[55] B. Brandenburg and A. Bastoni. The case for migratory priority inheritance in Linux: bounded priority inversions on multiprocessors. In *Proceedings of the 14th Real-Time Linux Workshop*, 2012.

[56] G. Calin, E. Derevenetc, R. Majumdar, and R. Meyer. A theory of partitioned global address spaces. In *FSTTCS 2013*. LIPICS, 2013.

[57] J. a. Carreira, R. Rodrigues, G. Candea, and R. Majumdar. Scalable testing of file system checkers. In *Proceedings of the 7th ACM european conference on Computer Systems*, EuroSys '12, pages 239–252, New York, NY, USA, 2012. ACM.

[58] C. Castelluccia, P. Druschel, S. F. Hübner, S. Gorniak, D. Ikonomou, A. Pasic, B. Preneel, H. Tschofening, and R. Tirtea. Privacy, Accountability and Trust—Challenges and Opportunities. In *European Network and Information Security Agency (ENISA) report*, 2011. Invited publication.

[59] F. Cerqueira and B. Brandenburg. A comparison of scheduling latency in Linux, PREEMPT-RT, and LITMUS$^{\mathrm{RT}}$. In *Proceedings of the 9th Annual Workshop on Operating Systems Platforms for Embedded Real-Time applications*, pages 19–29, 2013.

[60] A. Charguéraud. Characteristic formulae for the verification of imperative programs. In *Proceedings of the 16th ACM SIGPLAN International Conference on Functional programming (ICFP)*, pages 418–430. ACM, 2011.

[61] A. Charguéraud. The locally nameless representation. *Journal of Automated Reasoning*, pages 1–46, May 2011.

[62] K. Chatterjee, M. Chmelik, and R. Majumdar. Equivalence of games with probabilistic uncertainty and partial-observation games. In *ATVA*, volume 7561 of *Lecture Notes in Computer Science*, pages 385–399. Springer, 2012.

[63] K. Chatterjee, L. de Alfaro, M. Faella, R. Majumdar, and V. Raman. Code aware resource management. *Formal Methods in System Design*, 42(2):146–174, 2013.

[64] K. Chatterjee, L. de Alfaro, and R. Majumdar. The complexity of coverage. *Int. J. Found. Comput. Sci.*, 24(2):165–186, 2013.

[65] K. Chatterjee and R. Majumdar. Discounting and averaging in games across time scales. *Int. J. Found. Comput. Sci.*, 23(3):609–625, 2012.

[66] R. Chen, I. E. Akkus, and P. Francis. SplitX: high-performance private analytics. In *SIGCOMM*, pages 315–326, 2013.

[67] R. Chen, A. Reznichenko, P. Francis, and J. Gehrke. Towards statistical queries over distributed private user data. In *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*, NSDI'12, pages 13–13, Berkeley, CA, USA, 2012. USENIX Association.

[68] R. Chen, A. Shaikh, J. Wang, and P. Francis. Address-based Route Reflection. In *Proc. ACM CoNEXT*, December 2011.

[69] Y. Chen, J. Dunfield, and U. A. Acar. Type-based automatic incrementalization. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation*, June 2012.

[70] Y. Chen, J. Dunfield, M. A. Hammer, and U. A. Acar. Implicit self-adjusting computation for purely functional programs. In *Proceedings of the 16th ACM SIGPLAN International Conference on Functional programming (ICFP)*, pages 129–141, Sept. 2011.

[71] J. Cheney, A. Ahmed, and U. A. Acar. Provenance as dependency analysis. *Mathematical Structures in Computer Science*, 21(6):1301–1337, Dec. 2011.

[72] D. Chistikov, V. Fedorova, and A. Voronenko. Certificates of non-membership for classes of read-once functions. *Fundamenta Informaticae*, 2013.

[73] D. V. Chistikov and R. Majumdar. A uniformization theorem for nested word to word transductions. In *CIAA*, volume 7982 of *Lecture Notes in Computer Science*, pages 97–108. Springer, 2013.

[74] G. Claret, L. del Carmen Gonzáles Huesca, Y. Régis-Gianas, and B. Ziliani. Lightweight proof by reflection using a posteriori simulation of effectful computation. In *Conference on Interactive Theorem Proving (ITP)*, 2013.

[75] A. Clement, F. Junqueira, A. Kate, and R. Rodrigues. On the (limited) power of non-equivocation. In *Proceedings of the 2012 ACM symposium on Principles of distributed computing*, PODC '12, pages 301–308, New York, NY, USA, 2012. ACM.

[76] J. Cong, B. Liu, R. Majumdar, and Z. Zhang. Behavior-level observability analysis for operation gating in low-power behavioral synthesis. *ACM Trans. Design Autom. Electr. Syst.*, 16(1):4, 2010.

[77] 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.

[78] 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.

[79] 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.

[80] 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.

[81] 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.

[82] E. Darulova, V. Kuncak, R. Majumdar, and I. Saha. Synthesis of fixed-point programs. In *EMSOFT 2013*. ACM, 2013.

[83] A. Datta, J. Blocki, N. Christin, H. DeYoung, D. Garg, L. Jia, D. Kaynar, and A. Sinha. Understanding and protecting privacy: Formal semantics and principled audit mechanisms. In *International Conference on Information Systems Security (ICISS)*, 2011.

[84] A. Datta, A. Derek, J. C. Mitchell, and A. Roy. Protocol composition logic (PCL). In *Electronic Notes in Theoretical Computer Science (Gordon D. Plotkin Festschrift)*. 2007.

[85] D. Dreyer, A. Ahmed, and L. Birkedal. Logical step-indexed logical relations. *Logical Methods in Computer Science*, 7(2:16):1–37, June 2011. Special issue devoted to selected papers from LICS 2009.

[86] 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.

[87] D. Dreyer, G. Neis, A. Rossberg, and L. Birkedal. A relational modal logic for higher-order stateful ADTs. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, 2010.

[88] P. Druschel, M. Backes, and R. Tirtea. The right to be forgotten—between expectations and practice. In *European Network and Information Security Agency (ENISA) report*, 2012. Invited publication.

[89] J. Dunfield. Elaborating intersection and union types. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*, 2012.

[90] J. Dunfield and N. R. Krishnaswami. Complete and easy bidirectional typechecking for higher-rank polymorphism. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*, 2013.

[91] 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.

[92] J. Esparza, P. Ganty, and R. Majumdar. A perfect model for bounded verification. In *LICS*, pages 285–294. IEEE, 2012.

[93] J. Esparza, P. Ganty, and R. Majumdar. Parameterized verification of asynchronous shared-memory systems. In *CAV*, volume 8044 of *Lecture Notes in Computer Science*, pages 124–140, 2013.

[94] M. Fiore and C.-K. Hur. On the mathematical synthesis of equational logic. *Logical Methods in Computer Science*, 7(3:12):1–24, Sept. 2011.

[95] J. Fischer, R. Majumdar, and S. Esmaeilsabzali. Engage: a deployment management system. In *PLDI*, pages 263–274. ACM, 2012.

[96] C. Fournet and G. Gonthier. The reflexive chemical abstract machine and the join calculus. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, 1996.

[97] E. Fragkaki, L. Bauer, L. Jia, and D. Swasey. Modeling and enhancing Android's permission system. In *European Symposium on Research in Computer Security (ESORICS)*, 2012.

[98] P. Ganty and R. Majumdar. Algorithmic verification of asynchronous programs. *ACM Trans. Program. Lang. Syst.*, 34(1):6, 2012.

[99] P. Ganty, R. Majumdar, and B. Monmege. Bounded underapproximations. *Formal Methods in System Design*, 40(2):206–231, 2012.

[100] D. Garg, V. Genovese, and S. Negri. Countermodels from sequent calculi in multi-modal logics. In *ACM/IEEE Symposium on Logic in Computer Science (LICS)*, 2012.

[101] D. Garg, L. Jia, and A. Datta. Policy auditing over incomplete logs: Theory, implementation and applications. In *ACM Conference on Computer and Communications Security (CCS)*, 2011.

[102] D. Garg and F. Pfenning. Stateful authorization logic — Proof theory and a case study. *Journal of Computer Security*, 20(4), 2012.

[103] V. Genovese, D. Garg, and D. Rispoli. Labeled sequent calculi for access control logics: Countermodels, saturation and abduction. In *IEEE Symposium on Computer Security Foundations (CSF)*, 2012.

[104] V. Genovese, D. Garg, and D. Rispoli. Labeled goal-directed search in access control logic. In *International Workshop on Security and Trust Management (STM)*, 2012.

[105] 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.

[106] 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.

[107] 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.

[108] M. Gligoric and R. Majumdar. Model checking database applications. In *TACAS*, volume 7795 of *Lecture Notes in Computer Science*, pages 549–564. Springer, 2013.

[109] G. Gonthier, B. Ziliani, A. Nanevski, and D. Dreyer. How to make ad hoc polymorphism less ad hoc. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*, 2011.

[110] G. Gonthier, B. Ziliani, A. Nanevski, and D. Dreyer. How to make ad hoc proof automation less ad hoc. *Journal of Functional Programming*, 23(4):357–401, July 2013. Special issue devoted to selected papers from ICFP 2011.

[111] S. Guha, B. Cheng, and P. Francis. Privad: Practical Privacy in Online Advertising. In *NSDI*, 2011.

[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. Gupta, R. Majumdar, and A. Rybalchenko. From tests to proofs. *STTT*, 15(4):291–303, 2013.

[114] T. Gvero, V. Kuncak, I. Kuraj, and R. Piskac. Complete completion using types and weights. In *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '13, Seattle, WA, USA*, pages 27–38, 2013.

[115] A. Haeberlen, P. Fonseca, R. Rodrigues, and P. Druschel. Fighting Cybercrime with Packet Attestation. Technical Report MPI-SWS-2011-002, Max Planck Institute for Software Systems technical report, July 2011.

[116] M. A. Hammer, G. Neis, Y. Chen, and U. A. Acar. Self-adjusting stack machines. In *ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages and Applications (OOPSLA)*, 2011.

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

[118] T. A. Henzinger, A. Sezgin, and V. Vafeiadis. Aspect-oriented linearizability proofs. In P. R. D'Argenio and H. C. Melgratti, editors, *CONCUR*, volume 8052 of *Lecture Notes in Computer Science*, pages 242–256. Springer, 2013.

[119] M. P. Herlihy and J. M. Wing. Linearizability: a correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 12(3):463–492, 1990.

[120] T. Hillenbrand, R. Piskac, U. Waldmann, and C. Weidenbach. From search to computation: Redundancy criteria and simplification at work. In *Programming Logics - Essays in Memory of Harald Ganzinger*, pages 169–193, 2013.

[121] R. Hüchting, R. Majumdar, and R. Meyer. A theory of name boundedness. In *CONCUR*, volume 8052 of *Lecture Notes in Computer Science*, pages 182–196. Springer, 2013.

[122] 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.

[123] 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.

[124] C.-K. Hur, D. Dreyer, and V. Vafeiadis. Separation logic in the presence of garbage collection. In *IEEE Symposium on Logic in Computer Science (LICS)*, 2011.

[125] C.-K. Hur, G. Neis, D. Dreyer, and V. Vafeiadis. The power of parameterization in coinductive proof. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, 2013.

[126] R. Jhala, R. Majumdar, and A. Rybalchenko. Hmc: Verifying functional programs using abstract interpreters. In *CAV*, volume 6806 of *Lecture Notes in Computer Science*, pages 470–485. Springer, 2011.

[127] M. Jose and R. Majumdar. Bug-assist: Assisting fault localization in ansi-c programs. In *CAV*, volume 6806 of *Lecture Notes in Computer Science*, pages 504–509. Springer, 2011.

[128] 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.

[129] C. Kenna, J. Herman, B. Brandenburg, A. Mills, and J. Anderson. Soft real-time on multiprocessors: Are analysis-based schedulers really worth it? In *Proceedings of the 32nd IEEE Real-Time Systems Symposium*, pages 93–103, 2011.

[130] 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.

[131] J. Kloos and R. Majumdar. Supervisor synthesis for controller upgrades. In *DATE*, pages 1105–1110. EDA Consortium San Jose, CA, USA / ACM DL, 2013.

[132] J. Kloos, R. Majumdar, F. Niksic, and R. Piskac. Incremental, inductive coverability. In *CAV*, volume 8044 of *Lecture Notes in Computer Science*, pages 158–173. Springer, 2013.

[133] 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.

[134] 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.

[135] N. Krishnaswami and N. Benton. A semantic model for graphical user interfaces. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*, 2011.

[136] N. R. Krishnaswami. Higher-order functional reactive programming without spacetime leaks. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*, 2013.

[137] N. R. Krishnaswami and N. Benton. Adding equations to System F types. In *European Symposium on Programming (ESOP)*, 2012.

[138] N. R. Krishnaswami, N. Benton, and J. Hoffmann. Higher-order functional reactive programming in bounded space. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, 2012.

[139] N. R. Krishnaswami and D. Dreyer. Internalizing relational parametricity in the Extensional Calculus of Constructions. In *EACSL Annual Conference on Computer Science Logic (CSL)*, 2013.

[140] N. R. Krishnaswami, A. Turon, D. Dreyer, and D. Garg. Superficially substructural types. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*, 2012.

[141] 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.

[142] A. Kumar, V. Rajani, and D. Janakiram. Psi-CAL: Foundations of a programming language for services computing. In *IEEE International Conference on Services Computing (SCC)*, 2013.

[143] V. Kuncak, M. Mayer, R. Piskac, and P. Suter. Software synthesis procedures. *Commun. ACM*, 55(2):103–111, 2012.

[144] V. Kuncak, M. Mayer, R. Piskac, and P. Suter. Functional synthesis for linear arithmetic and sets. *STTT*, 15(5-6):455–474, 2013.

[145] 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.

[146] S. Le Blond, D. Choffnes, W. Zhou, P. Druschel, H. Ballani, and P. Francis. Towards efficient traffic-analysis resistant anonymity networks. In *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*, SIGCOMM '13, 2013.

[147] J. K. Lee, J. Palsberg, R. Majumdar, and H. Hong. Efficient may happen in parallel analysis for async-finish parallelism. In *SAS*, volume 7460 of *Lecture Notes in Computer Science*, pages 5–23. Springer, 2012.

[148] M. Lentz, V. Erdelyi, P. Aditya, E. Shi, P. Druschel, and B. Bhattacharjee. SDDR: Light-Weight Cryptographic Discovery for Mobile Encounters. Submitted for publication.

[149] X. Leroy. Formal verification of a realistic compiler. *Communications of the ACM*, 52(7):107–115, 2009.

[150] R. Ley-Wild, U. A. Acar, and G. E. Blelloch. Non-monotonic self-adjusting computation. In *21st European Symposium on Programming*, pages 476–496. Springer, 2012.

[151] C. Li, D. Porto, A. Clement, J. Gehrke, N. Preguiça, and R. Rodrigues. Making geo-replicated systems fast as possible, consistent when necessary. In *Proceedings of the 10th USENIX conference on Operating Systems Design and Implementation*, OSDI'12, pages 265–278, Berkeley, CA, USA, 2012. USENIX Association.

[152] 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.

[153] 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.

[154] Z. Long, G. Calin, R. Majumdar, and R. Meyer. Language-theoretic abstraction refinement. In *FASE*, volume 7212 of *Lecture Notes in Computer Science*, pages 362–376. Springer, 2012.

[155] R. Lopes Pereira, T. Vazão, and R. Rodrigues. Adaptive Search Radius - Using hop count to reduce P2P traffic. *Computer Networks*, 56(2):642–660, Feb. 2012.

[156] M. Mainberger, C. Schmaltz, M. Berg, J. Weickert, and M. Backes. Diffusion-Based Image Compression in Steganography. In *Proc. 8th International Symposium on Advances in Visual Computing (ISVC)*, pages 219–228, 2012.

[157] R. Majumdar. End-to-end guarantees in embedded control systems - (abstract). In E. M. Clarke, I. Virbitskaite, and A. Voronkov, editors, *Ershov Memorial Conference*, volume 7162 of *Lecture Notes in Computer Science*, page 40. Springer, 2011.

[158] R. Majumdar. The marriage of exploration and deduction. In *VSTTE*, volume 7152 of *Lecture Notes in Computer Science*, page 162. Springer, 2012.

[159] R. Majumdar, R. Meyer, and Z. Wang. Provenance verification. In *RP*, volume 8169 of *Lecture Notes in Computer Science*, pages 21–22. Springer, 2013.

[160] R. Majumdar, R. Meyer, and Z. Wang. Static provenance verification for message passing programs. In *SAS*, volume 7935 of *Lecture Notes in Computer Science*, pages 366–387. Springer, 2013.

[161] R. Majumdar, E. Render, and P. Tabuada. Robust discrete synthesis against unspecified disturbances. In *HSCC*, pages 211–220. ACM, 2011.

[162] R. Majumdar, I. Saha, K. C. Shashidhar, and Z. Wang. Clse: Closed-loop symbolic execution. In *NASA Formal Methods*, volume 7226 of *Lecture Notes in Computer Science*, pages 356–370. Springer, 2012.

[163] R. Majumdar, I. Saha, K. Ueda, and H. Yazarel. Compositional equivalence checking for models and code of control systems. In *CDC 13*, 2013.

[164] R. Majumdar, I. Saha, and M. Zamani. Performance-aware scheduler synthesis for control systems. In *EMSOFT*, pages 299–308. ACM, 2011.

[165] R. Majumdar, I. Saha, and M. Zamani. Synthesis of minimal-error control software. In *EMSOFT*, pages 123–132. ACM, 2012.

[166] R. Majumdar and Z. Wang. Expand, enlarge, and check for branching vector addition systems. In *CONCUR*, volume 8052 of *Lecture Notes in Computer Science*, pages 152–166. Springer, 2013.

[167] R. Majumdar and M. Zamani. Approximately bisimilar symbolic models for digital control systems. In *CAV*, volume 7358 of *Lecture Notes in Computer Science*, pages 362–377. Springer, 2012.

[168] 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.

[169] A. Nanevski, A. Banerjee, and D. Garg. Dependent type theory for verification of information flow and access control policies. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 35(2):6:1–6:41, July 2013.

[170] G. Neis, D. Dreyer, and A. Rossberg. Non-parametric parametricity. *Journal of Functional Programming*, 21(4&5):497–562, Sept. 2011. Special issue devoted to selected papers from ICFP 2009.

[171] K. G. Niederhoffer and J. W. Pennebaker. Linguistic style matching in social interaction. *Journal of Language and Social Psychology*, 21(4):337–360, 2002.

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

[173] R. Perera, U. A. Acar, J. Cheney, and P. B. Levy. Functional programs that explain their work. In *Proceedings of the 17th ACM SIGPLAN International Conference on Functional Programming*, pages 365–376. ACM, 2012.

[174] R. Piskac, T. Wies, and D. Zufferey. Automating separation logic using smt. In *Computer Aided Verification - 25th International Conference, CAV 2013, Saint Petersburg, Russia. Proceedings*, pages 773–789, 2013.

[175] V. Rajani, H. Mehta, S. J. Balaji, and D. Janakiram. KAAS: Kernel as a service. In *IEEE World Congress on Services (SERVICES)*, 2012.

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

[177] R. M. Reischuk, M. Backes, and J. Gehrke. SAFE extensibility of data-driven web applications. In *Proc. 21st World Wide Web Conference (WWW)*, pages 799–808, 2012.

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

[179] A. Reznichenko, S. Guha, and P. Francis. Auctions in do-not-track compliant internet advertising. In *ACM Conference on Computer and Communications Security*, pages 667–676, 2011.

[180] R. Rodrigues, B. Liskov, K. Chen, M. Liskov, and D. Schultz. Automatic Reconfiguration for Large-Scale Reliable Storage Systems. *IEEE Trans. Dependable Secur. Comput.*, 9(2):145–158, Mar. 2012.

[181] 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.

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

[183] A. Rossberg, C. Russo, and D. Dreyer. F-ing modules. Extended version of TLDI 2010 paper. Conditionally accepted to the Journal of Functional Programming (JFP), 2013.

[184] P. Roy, P. Tabuada, and R. Majumdar. Pessoa 2.0: a controller synthesis tool for cyber-physical systems. In *HSCC*, pages 315–316. ACM, 2011.

[185] I. Saha and R. Majumdar. Trigger memoization in self-triggered control. In *EMSOFT*, pages 103–112. ACM, 2012.

[186] N. Santos, R. Rodrigues, and B. Ford. Enhancing the OS against security threats in system administration. In *Proceedings of the 13th International Middleware Conference*, Middleware '12, pages 415–435, New York, NY, USA, 2012. Springer-Verlag New York, Inc.

[187] N. Santos, R. Rodrigues, K. P. Gummadi, and S. Saroiu. Policy-sealed data: a new abstraction for building trusted cloud services. In *Proceedings of the 21st USENIX conference on Security symposium*, Security'12, Berkeley, CA, USA, 2012. USENIX Association.

[188] J. Sevcík, V. Vafeiadis, F. Z. Nardelli, S. Jagannathan, and P. Sewell. Compcerttso: A verified compiler for relaxed-memory concurrency. *J. ACM*, 60(3):22, 2013.

[189] N. Sharma, S. Ghosh, F. Benevenuto, N. Ganguly, and K. P. Gummadi. Inferring Who-is-Who in the Twitter Social Network. In *Proceedings of the 4th ACM SIGCOMM Workshop On Social Networks (WOSN'12)*, 2012.

[190] O. Sümer, U. A. Acar, A. T. Ihler, and R. R. Mettu. Adaptive exact inference in graphical models. *J. Machine Learning Research*, 12:3147–3186, Nov. 2011.

[191] P. Tabuada, A. Balkan, S. Y. Caliskan, Y. Shoukry, and R. Majumdar. Input-output robustness for discrete systems. In *EMSOFT*, pages 217–226. ACM, 2012.

[192] S. Tetali, M. Lesani, R. Majumdar, and T. Millstein. Mrcrypt: Static analysis for secure cloud computations. In *OOPSLA*. ACM, 2013.

[193] 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.

[194] A. Turon and C. Russo. Scalable join patterns. In *ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages and Applications (OOPSLA)*, 2011.

[195] 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.

[196] Z. A. Uzmi, M. Nebel, A. Tariq, S. Jawad, R. Chen, A. Shaikh, J. Wang, and P. Francis. SMALTA: Practical and Near-Optimal FIB Aggregation. In *Proc. ACM CoNEXT*, December 2011.

[197] V. Vafeiadis. Concurrent separation logic and operational semantics. *Electr. Notes Theor. Comput. Sci.*, 276:335–351, 2011.

[198] V. Vafeiadis. Adjustable references. In S. Blazy, C. Paulin-Mohring, and D. Pichardie, editors, *ITP*, volume 7998 of *Lecture Notes in Computer Science*, pages 328–337. Springer, 2013.

[199] V. Vafeiadis and C. Narayan. Relaxed separation logic. In *OOPSLA*. ACM, 2013.

[200] V. Vafeiadis and F. Z. Nardelli. Verifying fence elimination optimisations. In E. Yahav, editor, *SAS*, volume 6887 of *Lecture Notes in Computer Science*, pages 146–162. Springer, 2011.

[201] A. Vahldiek, E. Elnikety, A. Mehta, D. Garg, P. Druschel, J. Gehrke, R. Rodrigues, and A. Post. Guardat: A foundation for policy-protected persistent data, July 2013. Submitted for publication.

[202] 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.

[203] 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.

[204] P. von Styp-Rekowsky, S. Gerling, M. Backes, and C. Hammer. Callee-Site Rewriting of Sealed System Libraries. In *Proc. 5th International Symposium on Engineering Secure Software and Systems (ESSoS)*, pages 33–41, 2013.

[205] A. Wieder, P. Bhatotia, A. Post, and R. Rodrigues. Orchestrating the deployment of computations in the cloud with Conductor. In *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*, NSDI'12, Berkeley, CA, USA, 2012. USENIX Association.

[206] A. Wieder and B. Brandenburg. Efficient partitioning of sporadic real-time tasks with shared resources and spin locks. In *Proceedings of the 8th IEEE International Symposium on Industrial Embedded Systems*, pages 49–58, 2013.

[207] 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*, 2013 (to appear).

[208] M. Zamani, P. Esfahani, R. Majumdar, A. Abate, and J. Lygeros. Bisimilar finite abstractions of stochastic control systems. In *CDC 13*, 2013.

[209] M. Zamani and R. Majumdar. A Lyapunov approach in incremental stability. In *CDC-ECE*, pages 302–307. IEEE, 2011.

[210] M. Zamani, N. van de Wouw, and R. Majumdar. Backstepping controller synthesis and characterizations of incremental stability. *Systems & Control Letters*, 62(10):949–962, 2013.

[211] M. Zengin and V. Vafeiadis. A programming language approach to fault tolerance for fork-join parallelism. In *TASE*. IEE, 2013.

[212] M. Zhao, P. Aditya, A. Chen, Y. Lin, A. Haeberlen, P. Druschel, B. Maggs, B. Wishon, and M. Ponec. Peer-Assisted Content Distribution in Akamai NetSession. In *Proceedings of the 13th ACM SIGCOMM Conference on Internet Measurement (IMC'13)*, 2013.

[213] B. Ziliani, D. Dreyer, N. R. Krishnaswami, A. Nanevski, and V. Vafeiadis. Mtac: a monad for typed tactic programming in Coq. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*, 2013.