

About Us

[Environment](#)
[Advisors](#)
[Doctoral students](#)
[Administrative office](#)

Program

[Structure](#)
[Preparatory phase](#)
[Dissertation phase](#)
[Admissions](#)
[Online Application](#)

Information

[Arrival](#)
[In Saarbrücken](#)
[On campus](#)
[Financial matters](#)

Events

[Social](#)
[Talks](#)
[Soft Skills](#)
[Guest program](#)



[Home](#) > [Doctoral students](#) > [PhD thesis defenses](#) > [PhD thesis defenses 2013](#)

2013

December

Sandro CASTRONOVO

The Pull Paradigm: Foundations of User-Centric Advanced Driver Assistance Systems Based on Bidirectional Car2X Communication

(Advisor: Prof. Wolfgang Wahlster)

Fri, 20 December 2013, 14:00h, building D3 2 (DFKI), Conference room Reuse -2,17

This thesis develops applications for vehicular ad-hoc networks that go far beyond the currently established areas of driving safety and traffic efficiency. The ad-hoc network is regarded as a dynamic information resource which is available to any vehicle at any time. In contrast to current state-of-the-art research, the proposed Pull Paradigm starts at the user's vehicle rather than at an information source somewhere in the network, e.g. a braking car. To access information from highly dynamic ad-hoc networks, bidirectional communication and information discovery and retrieval play a vital role. Therefore, in the course of the work, the applicability of the Pull Paradigm to established vehicular ad-hoc networks is thoroughly examined and missing aspects are identified. It turns out that a number of enhancements to almost all layers of the network stack are necessary in order to apply the Pull Paradigm using existing technology. The central elements here are two novel algorithms for managing information flow and dissemination in ad-hoc networks, which are at first formulated from the abstract perspective of graph theory. Using the knowledge gained leads to the development of PADE, a platform that supports development of vehicular ad-hoc network applications. The designed algorithms are then implemented as a routing scheme, integrated and evaluated in large, simulated city scenarios. Furthermore, PADE combines "real" and simulated communication technologies and abstracts from them, so that applications can be transferred from the lab into a test vehicle with minimal effort. In order to achieve this ambitious goal, PADE builds on a number of existing simulation and communication technologies. The practical applicability of the Pull approach is shown in two demonstrators that are integrated into a BMW 5 series test vehicle. The presentation module of the PADE platform was tested in the currently largest field operational test for vehicular ad-hoc communication. Over 400 drivers in 120 vehicles experienced the system on a daily basis.

Christian FEDERMANN

Hybrid Machine Translation using Binary Classification Models trained on Joint, Binarised Feature Vectors

(Advisor: Prof. Hans Uszkoreit)

Mon, 16 December 2013, 12:15h, building C7 4 (conference room)

We describe the design and implementation of a system combination method for machine translation output. It is based on sentence selection using binary classification models estimated on joint, binarised feature vectors. By contrast to existing system combination methods which work by dividing candidate translations into n-grams, i.e., sequences of n words or tokens, our framework performs sentence selection which does not alter the selected, best translation. First, we investigate the potential performance gain attainable by optimal sentence selection. To do so, we conduct the largest meta-study on data released by the yearly Workshop on Statistical Machine Translation (WMT). Second, we introduce so-called joint, binarised feature vectors which explicitly model feature value comparison for two systems A, B. We compare different settings for training binary classifiers using single, joint, as well as joint, binarised feature vectors. After having shown the potential of both selection and binarisation as methodological paradigms, we combine these two into a combination framework which applies pairwise comparison of all candidate systems to determine the best translation for each individual sentence. Our system is able to outperform other state-of-the-art system combination approaches; this is confirmed by our experiments. We conclude by summarising the main findings and contributions of our thesis and by giving an outlook to future research directions.

Thomas HELTEN

Processing and Tracking Human Motions Using Optical, Inertial, and Depth Sensors

(Advisors: Prof. Meinard Müller, Prof. Christian Theobalt)

Fri, 13 December 2013, 15:00h, building E1 4 (MPI-Inf), R 0.19

The processing of human motion data constitutes an important strand of research with many applications in computer animation, sport science and medicine. Currently, there exist various systems for recording human motion data that employ sensors of different modalities such as optical, inertial and depth sensors. Each of these sensor modalities have intrinsic advantages and disadvantages that make them suitable for capturing specific aspects of human motions as, for example, the overall course of a motion, the shape of the human body, or the kinematic properties of motions. In this thesis, we contribute with algorithms that exploit the respective strengths of these different modalities for comparing, classifying, and tracking



human motion in various scenarios. First, we show how our proposed techniques can be employed, e. g., for real-time motion reconstruction using efficient cross-modal retrieval techniques. Then, we discuss a practical application of inertial sensors-based features to the classification of trampoline motions. As a further contribution, we elaborate on estimating the human body shape from depth data with applications to personalized motion tracking. Finally, we introduce methods to stabilize a depth tracker in challenging situations such as in presence of occlusions. Here, we exploit the availability of complementary inertial-based sensor information.

Lizhen QU

Sentiment Analysis with Limited Training Data

(Advisor: Prof. Gerhard Weikum)

Wed, 4 December 2013, 10:30h, building E1 4 (MPI-Inf), R 0.24

Sentiments are positive and negative emotions, evaluations and stances. This dissertation focuses on learning based systems for automatic analysis of sentiments and comparisons in natural language text.

The proposed approach consists of three contributions:

1. Bag-of-opinions model: For predicting document-level polarity and intensity, we proposed the bag-of-opinions model by modeling each document as a bag of sentiments, which can explore the syntactic structures of sentiment-bearing phrases for improved rating prediction of online reviews.

2. Multi-experts model: Due to the sparsity of manually-labeled training data, we designed the multi-experts model for sentence-level analysis of sentiment polarity and intensity by fully exploiting any available sentiment indicators, such as phrase-level predictors and sentence similarity measures.

3. Senti-LSSVMrae model: To understand the sentiments regarding entities, we proposed Senti-LSSVMrae model for extracting sentiments and comparisons of entities at both sentence and subsentential level.

Different granularity of analysis leads to different model complexity, the finer the more complex. All proposed models aim to minimize the use of hand-labeled data by maximizing the use of the freely available resources. Our experimental results on real-world data showed that all models significantly outperform the state-of-the-art methods on the respective tasks.

November

Nuno SANTOS

Improving Trust in Cloud, Enterprise, and Mobile Computing Platforms

(Advisor: Dr. Rodrigo Rodrigues)

Wed, 27 November 2013, 18:00h, building E1 5 (MPI-SWS), R 0.29

Trust plays a fundamental role in the adoption of technology by society. Potential consumers tend to avoid a particular technology whenever they feel suspicious about its ability to cope with their security demands. Such a loss of trust could occur in important computing platforms, namely cloud, enterprise, and mobile platforms. In this thesis, we aim to improve trust in these platforms by (i) enhancing their security mechanisms, and (ii) giving their users guarantees that these mechanisms are in place.

To realize both these goals, we propose several novel systems. For cloud platforms, we present Excalibur, a system that enables building trusted cloud services. Such services give cloud customers the ability to process data privately in the cloud, and to attest that the respective data protection mechanisms are deployed. Attestation is made possible by the use of trusted computing hardware placed on the cloud nodes. For enterprise platforms, we propose an OS security model—the broker security model—aimed at providing information security against a negligent or malicious system administrator while letting him retain most of the flexibility to manage the OS. We demonstrate the effectiveness of this model by building BrokULOS, a proof-of-concept instantiation of this model for Linux. For mobile platforms, we present the Trusted Language Runtime (TLR), a software system for hosting mobile apps with stringent security needs (e.g., e-wallet). The TLR leverages ARM TrustZone technology to protect mobile apps from OS security breaches.

Tianxiang LU

Formal Verification of the Pastry Protocol

(Advisor: Prof. Christoph Weidenbach)

Wed, 27 November 2013, 10:30h, building E1 5 (MPI-SWS), R 0.02

Pastry is a structured P2P algorithm realizing a Distributed Hash Table (DHT) over an underlying virtual ring of nodes. Several implementations of Pastry are available and have been applied in practice, but no attempt has so far been made to formally describe the algorithm or to verify its properties. Since Pastry combines rather complex data structures, asynchronous communication, concurrency, resilience to churn, i.e. spontaneous join and departure of nodes, it makes an interesting target for verification.

This thesis focuses on Join protocol of Pastry and formally defines different statuses (from "dead" to "ready") of a node according to its stage during join. Only "ready" nodes are supposed to have consistent key mapping among each other and are allowed to deliver the answer message. The correctness property is identified by this thesis to be CorrectDelivery, stating that there is always at most one node that can deliver an answer to a lookup request for a key and this node is the numerically closest "ready" node to that key. This property is non-trivial to preserve in the presence of churn.

Through this thesis, unexpected violations of CorrectDelivery in previous versions of Pastry are discovered and analyzed using the TLA+ model checker TLC. Based on the analysis, Pastry is improved to a new design of the Pastry protocol IdealPastry, which is first verified using the interactive theorem prover TLAPS for TLA+.

IdealPastry is further improved to LuPastry, which is formally proved to be correct w.r.t. CorrectDelivery under the assumption that no nodes leave the network, which cannot be further relaxed due to possible network separation when particular nodes simultaneously leave the network.

Mohammed SHAHEEN

Cache based Optimization of Stencil Computations - An Algorithmic Approach

We are witnessing a fundamental paradigm shift in computer design. Memory has been and is becoming more hierarchical. Clock frequency is no longer crucial for performance. The on-chip core count is doubling rapidly. The quest for performance is growing. These facts have led to complex computer systems which bestow high demands on scientific computing problems to achieve high performance. Stencil computation is a frequent and important kernel that is affected by this complexity. Its importance stems from the wide variety of scientific and engineering applications that use it. The stencil kernel is a nearest-neighbor computation with low arithmetic intensity, thus it usually achieves only a tiny fraction of the peak performance when executed on modern computer systems. Fast on-chip memory modules were introduced as the hardware approach to alleviate the problem.

There are mainly three approaches to address the problem, cache aware, cache oblivious, and automatic loop transformation approaches. In this thesis, comprehensive cache aware and cache oblivious algorithms to optimize stencil computations on structured rectangular 2D and 3D grids are presented. Our algorithms observe the challenges for high performance in the previous approaches, devise solutions for them, and carefully balance the solution building blocks against each other. The many-core systems put the scalability of memory access at stake which has led to hierarchical main memory systems. This adds another locality challenge for performance.

We tailor our frameworks to meet the new performance challenge on these architectures. Experiments are performed to evaluate the performance of our frameworks on synthetic as well as real world problems.

October

Evgeny KRUGLOV **Superposition Modulo Theory**

(Advisor: Prof. Christoph Weidenbach)
Thu, 31 October 2013, 16:30h, building E1 5, R 0.02

This thesis is about the Hierarchic Superposition calculus $SUP(T)$ and its application to reasoning in hierarchic combinations $FOL(T)$ of the free first-order logic FOL with a background theory T . Particular hierarchic combinations covered in the thesis are the combinations of FOL and linear and non-linear arithmetic, LA and NLA resp. Recent progress in automated reasoning has greatly encouraged numerous applications in soft- and hardware verification and the analysis of complex systems. The applications typically require to determine the validity/unsatisfiability of quantified formulae over the combination of the free first-order logic with some background theories. The hierarchic superposition combines (i) reasoning in FOL equational clauses with universally quantified variables, which is based on the standard "flat" superposition calculus, and (ii) SMT-based reasoning techniques in such rich theories as, e.g., arithmetic, which are usually not (finitely) axiomatizable by FOL formulae. The thesis significantly extends previous results on $SUP(T)$, particularly: we introduce new substantially more effective sufficient completeness and hierarchic redundancy criteria turning $SUP(T)$ to a complete or a decision procedure for various $FOL(T)$ fragments; instantiate and refine $SUP(T)$ to effectively support particular combinations of FOL with the LA and NLA theories enabling a fully automatic mechanism of reasoning about systems formalized in $FOL(LA)$ or $FOL(NLA)$.

Tomasz JURKIEWICZ **Toward better Computation Models for Modern Machines**

(Advisor: Prof. Kurt Mehlhorn)
Wed, 30 October 2013, 14:00h, building E1 4, R 024

Modern computers are not random access machines (RAMs). They have a memory hierarchy, multiple cores, and a virtual memory. We address the computational cost of the address translation in the virtual memory and difficulties in design of parallel algorithms on modern many-core machines.

Starting point for our work on virtual memory is the observation that the analysis of some simple algorithms (random scan of an array, binary search, heapsort) in either the RAM model or the EM model (external memory model) does not correctly predict growth rates of actual running times. We propose the VAT model (virtual address translation) to account for the cost of address translations and analyze the algorithms mentioned above and others in the model. The predictions agree with the measurements. We also analyze the VAT-cost of cache-oblivious algorithms. In the second part of the paper we present a case study of the design of an efficient 2D convex hull algorithm for GPUs. The algorithm is based on the ultimate planar convex hull algorithm of Kirkpatrick and Seidel, and it has been referred to as the first successful implementation of the QuickHull algorithm on the GPU by Gao et al. in their 2012 paper on the 3D convex hull. Our motivation for work on modern many-core machines is the general belief of the engineering community that the theory does not produce applicable results, and that the theoretical researchers are not aware of the difficulties that arise while adapting algorithms for practical use. We concentrate on showing how the high degree of parallelism available on GPUs can be applied to problems that do not readily decompose into many independent tasks.

Gernot GEBHARD **Static Timing Analysis Tool Validation in the Presence of Timing Anomalies**

(Advisor: Prof. Reinhard Wilhelm)
Tues, 22 October 2013, 14:00h, building E1 7, R 001

The validation of the timing behavior of a safety-critical embedded software system requires both safe and precise worst-case execution time bounds for the tasks of that system. Such bounds need to be safe to ensure that each component of the software system performs its job in time. Furthermore, the execution time bounds are required to be precise to ensure the (provable) schedulability of the software system. When trying to achieve both safe and precise bounds, timing anomalies are one of the greatest challenges to overcome. Almost every modern hardware architecture shows timing anomalies, which also greatly impacts the analyzability of such architectures with respect to timing. Intuitively spoken, a timing anomaly is a counterintuitive behavior of a hardware architecture, where a "good" event (leg a

cache hit) leads to an overall longer execution, whereas the corresponding "bad" event (in this case, a cache miss) leads to a globally shorter execution time. In the presence of such anomalies, the local worst-case is not always a safe assumption in static timing analysis. To compute safe timing guarantees, any (static) timing analysis has to consider all possible executions.

In this thesis we investigate the source of timing anomalies in modern architectures and study instances of timing anomalies found in rather simple hardware architectures. Furthermore we discuss the impact of timing anomalies on static timing analysis. Finally we provide means to validate the result of static timing analysis for such architectures through trace validation.

Aleksandar STUPAR

Soundtrack Recommendation for Images

(Advisor: Dr. Sebastian Michel)

Fri, 4 October 2013, 10:00h, building E1 4, R 0.24

The drastic increase in production of multimedia content has emphasized the research concerning its organization and retrieval. In this thesis, we address the problem of music retrieval when a set of images is given as input query, i.e., the problem of soundtrack recommendation for images.

To tackle this problem, we formulate a hypothesis that the knowledge appropriate for the task is contained in publicly available contemporary movies. Our approach, Picasso, employs similarity search techniques inside the image and music domains, harvesting movies to form a link between the domains. In addition to the proposed Picasso approach, we further investigate effectiveness and efficiency of the task at hand and present a novel benchmark collection to evaluate Picasso and related approaches.

Rüdiger EHLERS

Symmetric and Efficient Synthesis

(Advisor: Prof. Bernd Finkbeiner)

Wed, 2 October 2013, 15:00h, building E1 7, Room 001

Despite the many advantages of synthesis over the manual engineering of reactive systems, it is not yet a well-established part of today's system design flows. It is commonly agreed on that the main reasons for this discrepancy are the lack of scalability of current synthesis techniques and the insufficient quality of the implementations computed.

In this thesis, we tackle both of these problems for reactive synthesis. To improve the quality of synthesized systems, we analyse the problem of symmetric synthesis. In this alternative synthesis problem, the aim is to compute a solution that consists of multiple copies of the same process such that the overall system satisfies the specification. Such systems have no centralised control units, and are considered to be more robust and easier to maintain. We characterise undecidable and decidable cases of the problem, and provide a synthesis algorithm for rotation-symmetric architectures, which capture many practical cases.

To improve the scalability in synthesis, we start with a simple but scalable approach to reactive synthesis that has shown its principal applicability in the field, and extend its main idea both in terms of scope and usability. We enhance its expressivity in a way that allows the synthesis of robust systems, and remove its limitation to specifications of a very special form. Both improvements yield theoretical insight into the synthesis problem: we characterise which specification classes can be supported in synthesis approaches that use parity games with a fixed number of colours as the underlying computation model, and examine the properties of universal very-weak automata, on which we base a synthesis workflow that combines the ease of specification with a low complexity of the underlying game solving step. As a side-result, we also obtain the first procedure to translate a formula in linear-time temporal logic (LTL) to a computation tree logic (CTL) formula with only universal quantifiers, whenever possible.

September

Iftikhar AHMAD

Analysis of Algorithms for Online Uni-directional Conversion Problems

(Advisor: Prof. Günter Schmidt)

Mon, 30 September 2013, 15:00h, building B4 1, Fakultätssaal 0.17

In an online uni-directional conversion problem, an online player wants to convert an asset $\$D\$$ to a desired asset $\$Y\$$. The objective of the player is to obtain the maximum amount of the desired asset. Competitive analysis is used as a tool for the design, and analysis of online algorithms for conversion problems. Although widely used, competitive analysis has its own set of drawbacks when the applicability of online algorithms in real world is considered. In this work, we investigate online uni-directional conversion problems with the objective to suggest measures for improving the applicability of online conversion algorithms in real world. First, we study competitive ratio as a coherent measure of risk and conclude that as it satisfies all the desirable axioms of coherence, competitive ratio can be used in practice. Secondly, we evaluate a selected set of online algorithms on real world as well bootstrap data to highlight the gap between theoretically guaranteed and experimentally achieved competitive ratio. The third aspect of the study deals with generating synthetic data that truly represents all possible scenarios such as market crashes. We suggest the use of Extreme Value Theory (EVT) approach. Using EVT approach, we generate synthetic data and execute a selected set of non-preemptive uni-directional online algorithms on it. The fourth contribution of the thesis includes the design and analysis of risk-aware reservation price algorithms for conversion problems. The proposed algorithms are flexible to accommodate the risk level of the online player whereas guaranteeing a bounded worst case competitive ratio as well. We evaluate our algorithms using the competitive analysis approach as well as testing the algorithms on the real world data. The results will help to improve the applicability of online conversion algorithms in real world. We conclude the work by discussing a number of open questions that will provide new directions for future research.

Esteban LÉON SOTO

Multi-agent Communication for the Realization of Business Processes

(Advisor: Prof. Jörg Siekmann)

As Internet and information technologies expand further into daily business activities, new solutions and techniques are required to cope with the growing complexity. One area that has gained attention is systems and organizations interoperability and Service Oriented Architectures (SOA).

Web Services have grown as a preferred technology in this area. Although these techniques have proved to solve problems of low level integration of heterogeneous systems, there has been little advance at higher levels of integration like how to rule complex conversations between participants that are autonomous and cannot depend on some ruling or orchestrating system. Multi-agent technology has studied techniques for content-rich communication, negotiation, autonomous problem solving and conversation protocols.

These techniques have solved some of the problems that emerge when integrating autonomous systems to perform complex business processes. The present research work intends to provide a solution for the realization of complex Business Process between heterogeneous autonomous participants using multi-agent technology. We developed an integration of Web Services and agent-based technologies along with a model for creating conversation protocols that respect the autonomy of participants. A modeling tool has been developed to create conversation protocols in a modular and reusable manner. BDI-Agents implementations that communicate over Web Services are automatically generated out of these models.

Cristián MADRIGAL-MORA

A model-driven approach for organizations in multi-agent systems

(Advisor: Prof. Jörg Siekmann)

Fri, 20 September 2013, 15:00h, building D3 2, Vis.-Center -1.63

This thesis introduces a new model-driven approach to agent-oriented software engineering in which agent organizations not only play a crucial role, but are also represented in every abstraction level. In our methodology, multiagent systems are modeled at a platform-independent level and transformed into a platform-specific level preserving the organizational structures. The approach has been refined through several years and has been used in two European Union projects.

Jochen MIROLL

Scalable and Rate Adaptive Wireless Multimedia Multicast

(Advisor: Prof. Thorsten Herfet)

Wed, 18 September 2013, 12:00h, building E1 7, Room 4.07

The methods that are described in this talk enable efficient multimedia Internet protocol streaming over wireless digital communication systems to an arbitrary number of receivers by multicast. A fundamental difference as compared to point-to-point connections between exactly two communicating stations is in conveying information about successful packet reception at the receiver side due to the multitude of receivers. This work considers time division multiple access systems, in which a single channel is used for data transmission and feedback. Therefore, the amount of time that should be spent for transmitting feedback is limited. Feedback about reception from the receiver(s) is necessary for efficient transmission. With respect to this, feedback for wireless multicast is evaluated and shown to be feasible. Aggregation of feedback in time is the mechanism proposed herein for physical layer bit rate adaptation. It is evaluated with respect to rate adaptation by example of orthogonal frequency division multiplex based IEEE 802.11 wireless networks. In proposed mechanisms, a constant amount of time is spent for multicast feedback, independent of the number of receivers (n). Therefore, also multimedia data throughput may remain independent of n . This may be taken for granted in case of statically configured, single purpose systems such as digital television. In the scope of this work are, however, multi-user and multi-purpose digital communication networks. Wireless LAN and LTE mobile networks are well-known examples. In suchlike systems, it is of great importance to remain independent of the number of receivers in order to maintain service quality. With a consumer hardware prototype for digital live-TV re-distribution in the local wireless network, the proposed mechanisms could be demonstrated.

Jens KERBER

Of Assembling Small Sculptures and Disassembling Large Geometry

(Advisor: Prof. Hans-Peter Seidel)

Tues, 17 September 2013, 14:00h, building E1 4 (MPII), Room 0.19

This thesis describes the research results and contributions that have been achieved during the author's doctoral work. It is divided into two independent parts, each of which is devoted to a particular research aspect. The first part covers the true-to-detail creation of digital pieces of art, so-called relief sculptures, from given 3D models. The main goal is to limit the depth of the contained objects with respect to a certain perspective without compromising the initial three-dimensional impression. Here, the preservation of significant features and especially their sharpness is crucial. Therefore, it is necessary to overemphasize fine surface details to ensure their perceptibility in the more complanate relief. Our developments are aimed at amending the flexibility and user-friendliness during the generation process. The main focus is on providing real-time solutions with intuitive usability that make it possible to create precise, lifelike and aesthetic results. These goals are reached by a GPU implementation, the use of efficient filtering techniques, and the replacement of user defined parameters by adaptive values. Our methods are capable of processing dynamic scenes and allow the generation of seamless artistic reliefs which can be composed of multiple elements. The second part addresses the analysis of repetitive structures, so-called symmetries, within very large data sets. The automatic recognition of components and their patterns is a complex correspondence problem which has numerous applications ranging from information visualization over compression to automatic scene understanding. Recent algorithms reach their limits with a growing amount of data, since their runtimes rise quadratically. Our aim is to make even massive data sets manageable. Therefore, it is necessary to abstract features and to develop a suitable, low-dimensional descriptor which ensures an efficient, robust, and purposive search. A simple inspection of the proximity within the descriptor space helps to significantly reduce the number of necessary pairwise comparisons. Our method scales quasi-linearly and allows a rapid analysis of data sets which could not be handled by prior

approaches because of their size.

Martin SUNKEL

Statistical Part-based Models for Object Detection in Large 3D Scans

(Advisor: Dr. Michael Wand)

Tues, 17 September 2013, 12:00h, building E1 4 (MPII), Room 0.19

3D scanning technology has matured to a point where very large scale acquisition of high resolution geometry has become feasible. However, having large quantities of 3D data poses new technical challenges. Many applications of practical use require an understanding of semantics of the acquired geometry. Consequently scene understanding plays a key role for many applications.

This thesis is concerned with two core topics: 3D object detection and semantic alignment. We address the problem of efficiently detecting large quantities of objects in 3D scans according to object categories learned from sparse user annotation. Objects are modeled by a collection of smaller sub-parts and a graph structure representing part dependencies. The thesis introduces two novel approaches: A part-based chain structured Markov model and a general part-based full correlation model. Both models come with efficient detection schemes which allow for interactive run-times.

Kristina SCHERBAUM

Data Driven Analysis of Faces from Images

(Advisor: Prof. Hans-Peter Seidel)

Tues, 17 September 2013, 10:00h, building E1 4 (MPII), Room 0.19

This talk proposes three new data-driven approaches to detect, analyze, or modify faces in images. All presented approaches are inspired by the use of prior knowledge and they derive information about facial appearances from pre-collected databases of images or 3D face models. First, we show an approach that extends a widely-used monocular face detector by an additional classifier that evaluates disparity maps of a passive stereo camera. The algorithm runs in real-time and significantly reduces the number of false positives compared to the monocular approach. Next, with a many-core implementation of the detector, we train view-dependent face detectors based on tailored views which guarantee that the statistical variability is fully covered. These detectors are superior to the state of the art on a challenging dataset and can be trained in an automated procedure. Finally, we present a model describing the relation of facial appearance and makeup. The approach extracts makeup from before/after images of faces and allows to modify faces in images. Applications such as machine-suggested makeup can improve perceived attractiveness as shown in a perceptual study. In summary, the presented methods help improve the outcome of face detection algorithms, ease and automate their training procedures and the modification of faces in images. Moreover, their data-driven nature enables new and powerful applications arising from the use of prior knowledge and statistical analyses.

Miguel GRANADOS

Advanced Editing Methods for Image and Video Sequences

(Advisor: Prof. Christian Theobalt)

Tues, 10 September 2013, 9:00h, building E1 4 (MPII), Room 0.19

In the context of image and video editing, this thesis proposes methods for modifying the semantic content of a recorded scene. Two different editing problems are approached: First, the removal of ghosting artifacts from high dynamic range (HDR) images recovered from exposure sequences, and second, the removal of objects from video sequences recorded with and without camera motion. These editings need to be performed in a way that the result looks plausible to humans, but without having to recover detailed models about the content of the scene, e.g. its geometry, reflectance, or illumination.

The proposed editing methods add new key ingredients, such as camera noise models and global optimization frameworks, that help achieving results that surpass the capabilities of state-of-the-art methods. Using these ingredients, each proposed method defines local visual properties that approximate well the specific editing requirements of each task. These properties are then encoded into a energy function that, when globally minimized, produces the required editing results. The optimization of such energy functions corresponds to Bayesian inference problems that are solved efficiently using graph cuts.

The proposed methods are demonstrated to outperform other state-of-the-art methods. Furthermore, they are demonstrated to work well on complex real-world scenarios that have not been previously addressed in the literature, i.e., highly cluttered scenes for HDR deghosting, and highly dynamic scenes and unconstrained camera motion for object removal from videos.

Matthias BÖHMER

Understanding and Supporting Mobile Application Usage

(Advisor: Prof. Antonio Krüger)

Fri, 6 September 2013, 14:00h, building D3 2, Room -2.17 (Reuse)

In recent years mobile phones have evolved significantly. While the very first cellular phones only provided functionality for conducting phone calls, smartphones nowadays provide a rich variety of functionalities. Additional hardware capabilities like new sensors (e.g. for location) and touch screens as new input devices gave rise to new use cases for mobile phones, such as navigation support, taking pictures or making payments. Mobile phones not only evolved with regard to technology, they also became ubiquitous and pervasive in people's daily lives by becoming capable of supporting them in various tasks. Eventually, the advent of mobile application stores for the distribution of mobile software enabled the end-users themselves to functionally customize their mobile phones for their personal purposes and needs.

So far, little is known about how people make use of the large variety of applications that are available. Thus, little support exists for end-users to make effective and efficient use of their smartphones given the huge numbers of applications that are available. This dissertation is motivated by the evolution of mobile phones from mere communication devices to multi-functional tool sets, and the challenges that have arisen as a result. The goal of this thesis is to contribute systems that support the use of mobile applications and to ground these systems' designs in an understanding of user behavior gained through empirical observations.

The contribution of this dissertation is twofold: First, this work aims to understand how people make use of, organize, discover and multitask between the various functionalities that are available for their smartphones. Findings are based on observations of user behavior by conducting studies in the wild. Second, this work aims to assist people in leveraging their smartphones and the functionality that is available in a more effective and efficient way. This results in tools and improved user interfaces for end-users. Given that the number of available applications for smartphones is rapidly increasing, it is crucial to understand how people make use of such applications to support smartphone use in everyday life with better designs for smartphone user interfaces.

Avishek ANAND

Indexing Methods for Web Archives

(Advisor: Dr. Klaus Berberich)

Fri, 6 September 2013, 9:00h, building E1 4 (MPII), Room 0.24

There have been numerous efforts recently to digitize previously published content and preserving born-digital content leading to the widespread growth of large text repositories. Web archives are such continuously growing text collections which contain versions of documents spanning over long time periods. Web archives present many opportunities for historical, cultural and political analyses. Consequently there is a growing need for tools which can efficiently access and search them.

In this work, we are interested in indexing methods for supporting text-search workloads over web archives like time-travel queries and phrase queries. To this end we make the following contributions:

- Time-travel queries are keyword queries with a temporal predicate, e.g., “mpii saarland” @ [06/2009], which return versions of documents in the past. We introduce a novel index organization strategy, called index sharding, for efficiently supporting time-travel queries without incurring additional index-size blowup. We also propose index-maintenance approaches which scale to such continuously growing collections.
- We develop query-optimization techniques for time-travel queries called partition selection which maximizes recall at any given query-execution stage.
- We propose indexing methods to support phrase queries, e.g., “to be or not to be that is the question”. We index multi-word sequences and devise novel query-optimization methods over the indexed sequences to efficiently answer phrase queries.

We demonstrate the superior performance of our approaches over existing methods by extensive experimentation on real-world web archives.

August

Matthias LANG

Foundations of Realistic Rendering - A Mathematical Approach

(Advisor: Prof. Philipp Slusallek)

Wed, 28 August 2013, 15:00h, building D3 2, DFKI, Visualisierungszentrum (-1.63, NB)

New, more efficient and more elegant algorithms for computing at least approximate solutions to the light transport equation and its different variants can be developed more easily in the context of a deeper understanding of the light transport equation. Since the problems of realistic renderings are deeply rooted in various mathematical disciplines, the complete understanding of the global illumination problem requires knowledge of several areas of mathematics.

Our objective will be to strictly mathematically formulate the global illumination problem based on principles of functional analysis, the theory of integral equations as well as of measure, integration, and probability theory. Additionally, we will try to show the interaction of all of these mathematical disciplines interwoven into the realistic rendering and to represent it in a comprehensible manner, especially for students and other interested people.

Bilyana TANEVA

Automatic Population of Knowledge Bases with Multimodal Data about Named Entities

(Advisor: Prof. Gerhard Weikum)

Mon, 12 August 2013, 14:00h, building E1 4 (MPI-Inf), Room 024

Knowledge bases are of great importance for Web search, recommendations, and many Information Retrieval tasks. However, maintaining them for not so popular entities is often a bottleneck. Typically, such entities have limited textual coverage and only a few ontological facts. Moreover, these entities are not well populated with multimodal data, such as images, videos, or audio recordings.

The goals in this thesis are (1) to populate a given knowledge base with multimodal data about entities, such as images or audio recordings, and (2) to ease the task of maintaining and expanding the textual knowledge about a given entity, by recommending valuable text excerpts to the contributors of knowledge bases.

The thesis makes three main contributions. The first two contributions concentrate on finding images of named entities with high precision, high recall, and high visual diversity. Our main focus are less popular entities, for which the image search engines fail to retrieve good results. Our methods utilize background knowledge about the entity, such as ontological facts or a short description, and a visual-based image similarity to rank and diversify a set of candidate images.

Our third contribution is an approach for extracting text contents related to a given entity. It leverages a language-model-based similarity between a short description of the entity and the text sources, and solves a budget-constraint optimization program without any assumptions on the text structure. Moreover, our approach is also able to reliably extract entity related audio excerpts from news podcasts. We derive the time boundaries from the usually very noisy audio transcriptions.

Matthias BERG

Formal Verification of Cryptographic Security Proofs

(Advisor: Prof. Michael Backes)

Fri, 9 August 2013, 15:00h, building E1 1, Room 407

Verifying cryptographic security proofs manually is inherently tedious and error-

prone. The game-playing technique for cryptographic proofs advocates a modular proof design where cryptographic programs called games are transformed stepwise such that each step can be analyzed individually. This code-based approach has rendered the formal verification of such proofs using mechanized tools feasible. In the first part of this dissertation we present Verypto: a framework to formally verify game-based cryptographic security proofs in a machine-assisted manner. Verypto has been implemented in the Isabelle proof assistant and provides a formal language to specify the constructs occurring in typical cryptographic games, including probabilistic behavior, the usage of oracles, and polynomial-time programs. We have verified the correctness of several game transformations and demonstrate their applicability by verifying that the composition of 1-1 one-way functions is one-way and by verifying the IND-CPA security of the ElGamal encryption scheme. In a related project Barthe et al. developed the EasyCrypt toolset, which employs techniques from automated program verification to validate game transformations. In the second part of this dissertation we use EasyCrypt to verify the security of the Merkle-Damgaard construction - a general design principle underlying many hash functions. In particular we verify its collision resistance and prove that it is indistinguishable from a random oracle.

Ralf OSBILD

General Analysis Tool Box for Controlled Perturbation Algorithms and Complexity and Computation of Θ -Guarded Regions

(Advisor: Prof. Kurt Mehlhorn)

Fri, 2 August 2013, 14:00h, building E1 4 (MPI-Inf), Room 024

Diese Dissertation auf dem Gebiet der Algorithmischen Geometrie beschäftigt sich mit den folgenden zwei Problemen.

1. Die Implementierung von verlässlichen und effizienten geometrischen Algorithmen ist eine herausfordernde Aufgabe. Controlled Perturbation verknüpft die Geschwindigkeit von Fließ-komma-Arithmetik mit einem Mechanismus, der die Verlässlichkeit garantiert. Wir präsentieren einen allgemeinen „Werkzeugkasten“ zum Analysieren von Controlled Perturbation Algorithmen. Dieser Werkzeugkasten ist in unabhängige Komponenten aufgeteilt. Wir präsentieren drei alternative Methoden für die Herleitung der wichtigsten Schranken. Des Weiteren haben wir alle Prädikate, die auf Polynomen und rationalen Funktionen beruhen, sowie Objekt-erhaltende Perturbationen in die Theorie miteinbezogen. Darüber hinaus wurde der Werkzeugkasten so entworfen, dass er das tatsächliche Verhalten des untersuchten Algorithmus ohne vereinfachende Annahmen widerspiegelt.
2. Illumination und Guarding Probleme stellen ein breites Gebiet der Algorithmischen und Kombinatorischen Geometrie dar. Hierzu tragen wir die Komplexität und Berechnung von Θ -bewachten Regionen bei. Sie stellen eine Verallgemeinerung der konvexen Hülle dar und sind mit α -hulls und Θ -maxima verwandt. Die Schwierigkeit beim Studium der Θ -bewachten Regionen ist die Abhängigkeit ihrer Form und Komplexität von Θ . Für alle Winkel Θ beweisen wir grundlegende Eigenschaften der Region, leiten untere und obere Schranken ihrer worst-case Komplexität her und präsentieren einen Algorithmus, um die Region zu berechnen.

July

Jeremias RÖBLER

From Software Failure to Explanation

(Advisor: Prof. Andreas Zeller)

Fri, 12 July 2013, 10:00h, building E1 1, Room 407

“Why does my program crash?”—This ever recurring question drives the developer both when trying to reconstruct a failure that happened in the field and during the analysis and debugging of the test case that captures the failure.

This is the question this thesis attempts to answer. For that I will present two approaches which, when combined, start off with only a dump of the memory at the moment of the crash (a core dump) and eventually give a full explanation of the failure in terms of the important runtime features of the program such as critical branches, state predicates or any other execution aspect that is deemed helpful for understanding the underlying problem.

The first approach (called RECORE) takes a core dump of a crash and by means of search-based test case generation comes up with a small, self-contained and easy to understand unit test that is similar to the test as it is attached to a bug report and reproduces the failure. This test case can serve as a starting point for analysis and manual debugging. Our evaluation shows that in five out of seven real cases, the resulting test captures the essence of the failure. But this failing test case can also serve as the starting point for the second approach (called BUGEX). BUGEX is a universal debugging framework that applies the scientific method and can be implemented for arbitrary runtime features (called facts). First it observes those facts during the execution of the failing test case. Using state-of-the-art statistical debugging, these facts are then correlated to the failure, forming a hypothesis. Then it performs experiments: it generates additional executions to challenge these facts and from these additional observations refines the hypothesis. The result is a correlation of critical execution aspects to the failure with unprecedented accuracy and instantaneously point the developer to the problem. This general debugging framework can be implemented for any runtime aspects; for evaluation purposes I implemented it for branches and state predicates. The evaluation shows that in six out of seven real cases, the resulting facts pinpoint the failure.

Both approaches are independent from one another and each automates a tedious and error prone task. When being combined, they automate a large part of the debugging process, where the remaining manual task—fixing the defect—can never be fully automated.

May

Kaustubh PATIL

Genome signature based sequence comparison for taxonomic assignment and tree inference

(Advisor: Prof. Alice McHardy)

Wed, 29 May 2013, 15:00h, building E1 5 (MPI-SWS), HS 002

In this work we consider the use of the genome signature for two important

bioinformatics problems; the taxonomic assignment of metagenome sequences and tree inference from whole genomes. We look at those problems from a sequence comparison point of view and propose machine learning based methods as solutions. For the first problem, we propose a novel method based on structural support vector machines that can directly predict paths in a tree implied by evolutionary relationships between taxa. For the second problem we propose a distance metric learning method. Based on the assumption that for different groups of prokaryotes different oligonucleotide weights can be more informative, our method learns group-specific distance metrics. In the outlook, we expect that for the addressed problems the work of this thesis will complement and in some cases even outperform alignment-based sequence comparison at a considerably reduced computational cost, allowing it to keep up with advancements in sequencing technologies.

Sabine SCHMALTZ

Towards the Pervasive Formal Verification of Multi-Core Operating Systems and Hypervisors Implemented in C

(Advisor: Prof. Wolfgang Paul)

Fri, 10 May 2013, 14:00h, building E1 7, Room 001

The Verisoft XT project had the goal of verifying correctness of the Microsoft Hyper-V hypervisor and achieved great code verification results using the concurrent C verification tool VCC developed by our project partners during the project. A sound mathematical theory to support code verification was not established.

To remedy this shortcoming, we sketch a model stack for a simplified multi-core architecture based on a simplified MIPS model for system programmers and illustrate on a high level of abstraction how to obtain a simulation between neighboring models. We survey the current state of theory development and outline missing links and parts.

As part of the dissertation, a hardware model for our architecture is formalized at a detailed level of abstraction of the model stack. In addition, we provide operational semantics for a quite simple intermediate language for C as well as an extension of this semantics with specification (ghost) state and code which can serve as a basis for arguing the soundness of VCC.

Kim HERZIG

TLB Virtualization in the Context of Hypervisor Verification

(Advisor: Prof. Andreas Zeller)

Mon, 6 May 2013, 16:00h, building E1 1, Room 407

Developers change source code to add new functionality, fix bugs, or refactor their code. Many of these changes have immediate impact on quality or stability. However, some impact of changes may become evident only in the long term. This thesis makes use of change genealogy dependency graphs modeling dependencies between code changes capturing how earlier changes enable and cause later ones. Using change genealogies, it is possible to:

(a) Apply formal methods like model checking on version archives to reveal temporal process patterns. Such patterns encode key features of the software process and can be validated automatically:

In an evaluation of four open source histories, our prototype would recommend pending activities with a precision of 60-72%.

(b) Classify the purpose of code changes. Analyzing the change dependencies on change genealogies shows that change genealogy network metrics can be used to automatically separate bug fixing from feature implementing code changes.

(c) Build competitive defect prediction models. Defect prediction models based on change genealogy network metrics show competitive prediction accuracy when compared to state-of-the-art defect prediction models.

As many other approaches mining version archives, change genealogies and their applications rely on two basic assumptions: code changes are considered atomic and bug reports are considered to refer to corrective maintenance tasks. In a manual examination of more than 7,000 issue reports and code changes from bug databases and version control systems of open-source projects, we found 34% of all issue reports to be mis-classified and that up to 15% of all applied issue fixes consist of multiple combined code changes serving multiple developer maintenance tasks. This introduces bias in bug prediction models confusing bugs and features. To partially solve these issues and to measure its impact we present an approach to untangle such combined changes with a mean success rate of 58-90% after the fact.

March

Mikhail KOVALEV

Mining and Untangling Change Genealogies

(Advisor: Prof. Wolfgang Paul)

Wed, 27 Mar 2013, 15:00h, building E1 7, Room 0.01

In this thesis we address the challenges of hypervisor verification for multicore processors. As a first contribution we unite different pieces of hypervisor verification theory into a single theory comprising the stack of highly nontrivial computational models used. We consider multicore hypervisors for x86-64 architecture written in C. To make code verification in a C verifier possible, we define a reduced hardware model and show that under certain safety conditions it simulates the full model. We introduce an extension of the C semantics, which takes into consideration possible MMU and guest interaction with the memory of a program. We argue that the extended C semantics simulates the hardware machine, which executes compiled hypervisor code, given that the compiler is correct.

The second contribution of the thesis is the formal verification of a software TLB and memory virtualization approach, called SPT algorithm. Efficient TLB virtualization is one of the trickiest parts of building correct hypervisors. An SPT algorithm maintains dedicated sets of "shadow" page tables, ensuring memory separation and correct TLB abstraction for every guest. We use our extended C semantics to specify correctness criteria for TLB virtualization and to verify a simple SPT algorithm written in C. The code of the algorithm is formally verified in Microsoft's VCC automatic verifier, which is ideally suited for proofs performed on top of our semantic stack.

Oana CIOBOTARU

Rational Cryptography: Novel Constructions, Automated Verification and Unified Definitions

(Advisor: Prof. Michael Backes)

Tues, 26 Mar 2013, 14:00h, building E1 7, Room 0.01

Rational cryptography has recently emerged as a very promising field of research by combining notions and techniques from cryptography and game theory, because it offers an alternative to the rather inflexible traditional cryptographic model. In contrast to the classical view of cryptography where protocol participants are considered either honest or arbitrarily malicious, rational cryptography models participants as rational players that try to maximize their benefit and thus deviate from the protocol only if they gain an advantage by doing so.

The main research goals for rational cryptography are the design of more efficient protocols when players adhere to a rational model, the design and implementation of automated proofs for rational security notions and the study of the intrinsic connections between game theoretic and cryptographic notions. In this thesis, we address all these issues. First we present the mathematical model and the design for a new rational file sharing protocol which we call RatFish. Next, we develop a general method for automated verification for rational cryptographic protocols and we show how to apply our technique in order to automatically derive the rational security property for RatFish. Finally, we study the intrinsic connections between game theory and cryptography by defining a new game theoretic notion, which we call game universal implementation, and by showing its equivalence with the notion of weak stand-alone security.

Stefan WARWAS

A Model-driven Framework for Engineering Multiagent Systems

(Advisor: Prof. Philipp Slusallek)

Mon, 11 Mar 2013, 14:00h, building D3 4 (access via D3 2), VIS-Room -1.63

Since the invention of computer systems, the level of abstraction of software languages has been steadily increased from op-codes to object-oriented languages. Agent technology promises to embody high-level concepts that go beyond those of object-oriented approaches. This dissertation presents the Bochica framework for Agent-Oriented Software Engineering (AOSE). The framework's task in the software development process is (i) to capture the design decisions for a system under consideration on a platform-independent level of abstraction and (ii) to project this design to a target platform. Bochica goes beyond the state-of-the-art in AOSE as it combines the benefits of a platform-independent approach with the possibility to address concepts of custom application domains and execution environments. Several extension interfaces are specified to enable the customization of the underlying modeling language to the engineer's needs. Bochica is accompanied by an iterative adaptation process to gradually incorporate extensions. Conceptual mappings for projecting Bochica models to executable code are specified. In order to enable Bochica for modeling agents that inhabit semantically-enhanced virtual worlds, an according extension model is proposed. Finally, a model-driven reverse engineering approach for lifting the underlying design of already implemented Multiagent System (MAS) to the platform-independent layer is introduced. The framework has been successfully evaluated for designing intelligent agents that operate a virtual production line as well as for extracting the underlying design of an already implemented MAS. The evaluation results show that the Bochica approach to AOSE contributes to overcome the gap between design and code.

February

Ya-Fang WANG

Methods and Tools for Temporal Knowledge Harvesting

(Advisor: Prof. Gerhard Weikum)

Mon, 25 Feb 2013, 15:00h, building E1 4 (MPI-Inf), Rm 0.24

To extend the traditional knowledge base with temporal dimension, this thesis offers methods and tools for harvesting temporal facts from both semi-structured and textual sources. Our contributions are briefly summarized as follows.

1. Timely YAGO: A temporal knowledge base called Timely YAGO (T-YAGO) which extends YAGO with temporal attributes is built. We define a simple RDF-style data model to support temporal knowledge.
2. PRAVDA: To be able to harvest as many temporal facts from free-text as possible, we develop a system PRAVDA. It utilizes a graph-based semi-supervised learning algorithm to extract fact observations, which are further cleaned up by an Integer Linear Program based constraint solver. We also attempt to harvest spatio-temporal facts to track a person's trajectory.
3. PRAVDA-live: A user-centric interactive knowledge harvesting system, called PRAVDA-live, is developed for extracting facts from natural language free-text. It is built on the framework of PRAVDA. It supports fact extraction of user-defined relations from ad-hoc selected text documents and ready-to-use RDF exports.
4. T-URDF: We present a simple and efficient representation model for time-dependent uncertainty in combination with first-order inference rules and recursive queries over RDF-like knowledge bases. We adopt the common possible-worlds semantics known from probabilistic databases and extend it towards histogram-like confidence distributions that capture the validity of facts across time. All of these components are fully implemented systems, which together form an integrative architecture. PRAVDA and PRAVDA-live aim at gathering new facts (particularly temporal facts), and then T-URDF reconciles them.

Finally these facts are stored in a (temporal) knowledge base, called T-YAGO. A SPARQL-like time-aware querying language, together with a visualization tool, are designed for T-YAGO. Temporal knowledge can also be applied for document summarization.

Christoph CULLMANN

Cache Persistence Analysis for Embedded Real-Time Systems

(Advisor: Prof. Reinhard Wilhelm)

Thurs, 14 Feb 2013, 15:00h, building E1 7, Rm 001

To compute a worst-case execution time (WCET) estimate for a program running on

a safety-critical hard real-time system, the effects of the architecture of the underlying hardware have to be modeled. The classical cache analysis distinguishes three categories for memory references to cached memory: always-hit, always-miss and not-classified. The cache persistence analysis tries to classify memory references as persistent thereby improving the classical cache analysis by limiting the number of misses for not-classified memory references.

We present several new abstract interpretation based cache persistence analyses.

Two are based on the concept of conflict counting, one on the may cache analysis, and one combines both concepts. All analyses also fix a correctness issue of the original cache persistence analysis by Ferdinand and Wilhelm.

For non-fully-timing-compositional architectures using the persistence information is not straightforward. A novel path analysis enables the use of persistence information also for state-of-the-art architectures that exhibit timing anomalies / domino effects.

The new analyses are practically evaluated within the industrially used WCET analyzer aiT on a series of standard benchmark programs and a series of real avionic examples.

January

Jens HAUPERT

DOMeMan: Repräsentation, Verwaltung und Nutzung von digitalen Objektgedächtnissen

(Advisor: Prof. Wolfgang Wahlster)

Wed, 16 Jan 2013, 16:15h, building D3 2 (DFKI), Rm -2.17 (Reuse)

This thesis addresses the research question, how an infrastructure for digital object memories has to be designed. Primary goal of this thesis is to identify and develop components and processes of an architecture concept particularly suited to represent, manage, and use digital object memories. In order to leverage acceptance and deployment of this novel technology, the envisioned infrastructure has to include tools for integration of new systems, and for migration with existing systems.

Special requirements to object memories result from the heterogeneity of data in so-called open-loop scenarios. On the one hand, it has to be flexible enough to handle different data types. On the other hand, a simple and structured data access is required. Depending on the application scenario, the latter one needs to be complemented with concepts for a rights- and role-based access and version control.

First, this thesis provides a data model for structuring object memory data by means of meta data, which enables cross-domain data exchange in heterogeneous scenarios. Then, a software architecture concept will be introduced, which provides means of storing and accessing memory data, and integrates an activity component into object memories. Finally, the concept is completed by a toolset that enables developers to migrate existing datasets, to create new structures based on semantic information, and to support user interaction with this data by means of application-specific visualizations.