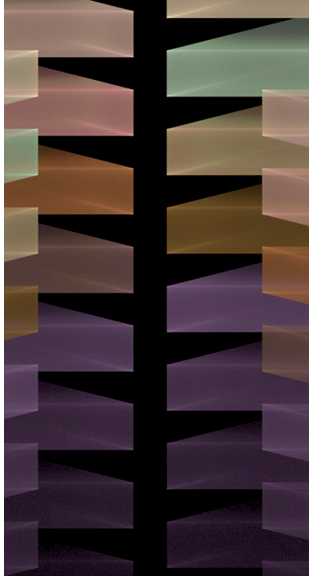

madrid institute
for advanced studies

annual
report
2012

institute
iMdea
software

annual report

foreword

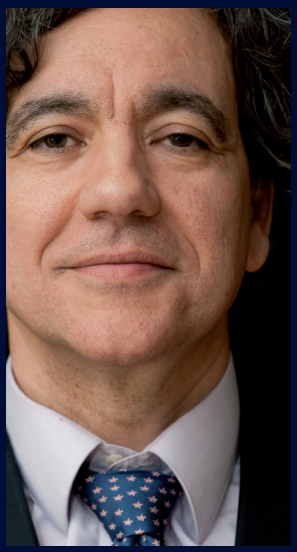

Manuel Hermenegildo
Director, IMDEA Software Institute

April 15, 2013

annual report
2012

The IMDEA Software Institute was created by the Regional Government of Madrid under the strong belief, more relevant currently than ever, that quality research in technology-related areas is the most successful and cost-effective way of generating knowledge, sustainable growth, and employment. Software-related technology indeed has an immense potential for raising industrial competitiveness, opening whole new business areas, creating high added-value jobs, and, ultimately, improving quality of life. Today, gathering the material for this 2012 annual report, the Institute manifests itself as a vibrant, exciting reality, making significant progress towards its goals of excellence in research and technology transfer.

Without any doubt, the main strength of the Institute is its people: its researchers and support staff. It has been very successful in attracting to Madrid top talent worldwide, including now 18 faculty (one half-time), 11 postdocs, 17 research assistants, and a number of interns, from 17 different nationalities. They joined after working at or obtaining their Ph.D. degrees from 32 different prestigious centers in 8 different countries, including Stanford U., Carnegie Mellon U., or Microsoft Research in the US, INRIA in France, U. of Cambridge in the UK, the Max Planck Software Institute in Germany, or ETH in Switzerland, to name just a few. In addition, 90 international researchers have visited and given talks at the Institute to date.

During 2012 Institute researchers have published 57 refereed articles (in some of the top venues in the field, such as POPL, ACM TOPLAS, CRYPTO, IEEE S&P, USENIX Security, CSF, JCS, FM, CAV, TPLP, ICSOC, ICFP, ICLP, ICALP, etc.), edited 2 proceedings, given 22 invited talks and 15 invited seminars and lectures, and participated in 56 program committees and 12 boards of journals and conferences, in addition to being conference and program chairs of 3 conferences. The Institute has received 6 best paper awards or mentions just in the last 2 years. Institute researchers received 2 of the 10 *Microsoft Software Engineering Innovation Foundation* awards given in 2012.

The Institute has also participated during 2012 in 19 funded research projects and contracts (9 by the EU and 1 by the US ONR and Stanford U.) and 16 fellowships, collaborating with a large number of companies including Atos, Siemens, Deimos, AbsInt, Microsoft, Fredhopper, Telefonica, Boeing, and Thales (and with many others in other recent projects, such as France Telecom, SAP, Trusted Logic, Airbus, Alcatel, Daimler, or EADS). The Institute is also working on the commercialization of the ActionGUI technology developed by its Modeling Lab in collaboration with ETH Zurich.

The Institute has also developed further its strategic partnership with Telefonica, Indra, Atos, or BBVA, including for example leading the presentation of a joint candidacy (also with UPM and BSC) for the extension of the European Institute of Technology (EIT) ICT Labs to Spain. The Institute is now an Associate Member of EIT ICT Labs and coordinates this node-building task from Madrid.

Finally, a major milestone in 2012 has been the completion of the construction of the Institute building. After obtaining all the certifications, the Institute moved in January 2013 to these permanent premises which offer an ideal environment for expanding the development of its multiple research and technology transfer objectives.

Many thanks to all who have contributed to all these achievements, and very specially to the Madrid Regional Government for their continuing vision and support.

table of contents

table of contents

annual report

2012

General Presentation

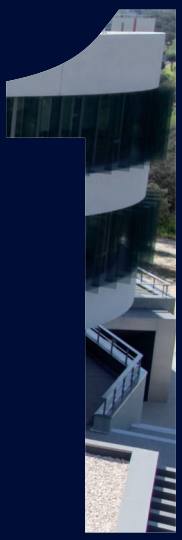

annual report

2012

1.1. Profile

The IMDEA Software Institute (Madrid Institute for Advanced Studies in Software Development Technologies) is a non-profit, independent research institute promoted by the Madrid Regional Government (CM) to perform research of excellence in the methods, languages, and tools that will allow the cost-effective development of software products with sophisticated functionality and high quality, i.e., safe, reliable, and efficient.

The IMDEA Software Institute is part of the Madrid Institute for Advanced Studies (IMDEA) network, an institutional framework created to foster social and economic growth in the region of Madrid by promoting research of excellence and technology transfer in a number of strategic areas (water, food, social sciences, energy, materials, nanoscience, networks, and software) with high potential impact.

1.2. Motivation and Goals

It is difficult to overstate the importance of software both for our everyday lives and for the industrial processes which, running behind the scenes, are necessary to sustain the modern world. Software is the enabling technology in many devices and services which are now an essential part of our world and on which we, to different degrees, depend: accounting, banking, cell phones, cars, flight control systems, behavior of the stock market, digital television, life support systems... not to mention tablets, computers, and the Internet itself. This pervasiveness explains the global figures around software and the IT services sector: according to the Global Industry Guide, the global software market had an estimated value of 293.000 M€ in 2011, which is estimated to grow to 396.000 M€ by 2016. It is one of the few sectors which, despite the economic turmoil, continues to grow in terms of turnover, profit, and jobs. According to European Commission data, in 2012 the ICT sector accounted for 6% of EU GDP and approximately eight million jobs.

Given the economic relevance of software and its pervasiveness, errors and failures in software can have high social and economic cost: the results of malfunctioning software can range from being annoying (e.g., having to reboot), through posing serious social and legal problems (e.g., privacy and security leaks), to having high economic cost (e.g., software failures in financial markets and software-related product recalls) or even being a threat to human lives (e.g., a malfunctioning airplane or medical device). Unfortunately, developing software of an appropriate level of reliability, security, and performance, at a reasonable cost is still a challenge today. Some degree of correctness can be achieved by careful human or machine-assisted inspection at very high monetary costs, but the risk of errors produced by human

mistakes will still be lurking in the dark. Reducing software errors in a cost-effective manner is therefore a task better left to automatic tools. These tools are, however, extremely hard to produce and pose scientific and technological challenges. At the same time, the ubiquity of software makes tacking these challenges a potentially highly profitable endeavor, since solutions to these challenges can have a significant and pervasive impact on productivity and on the general competitiveness of the economy.

Therefore, the main mission of the IMDEA Software Institute is to perform research of excellence in methods, languages, and tools that will allow the cost-effective development of software products with sophisticated functionality and high quality, i.e., secure, reliable, and efficient. This research focus includes all phases of the development cycle (analysis, design, implementation, validation, verification, maintenance); its distinguishing feature is the concentration on approaches that are rigorous and at the same time allow building practical tools.

In order to achieve its mission, the IMDEA Software Institute is gathering a critical mass of world-wide, top-class researchers, and is at the same time developing synergies between them and the already significant research base and industrial capabilities existing in the region. Indeed, most of the IT-related companies in Spain (and, specially, their research divisions) are located in the Madrid region, which facilitates collaboration and technology transfer. Thus, the IMDEA Software Institute brings about the opportunity of grouping a critical mass of researchers and industrial experts, which can allow for significant improvement in the impact of research.

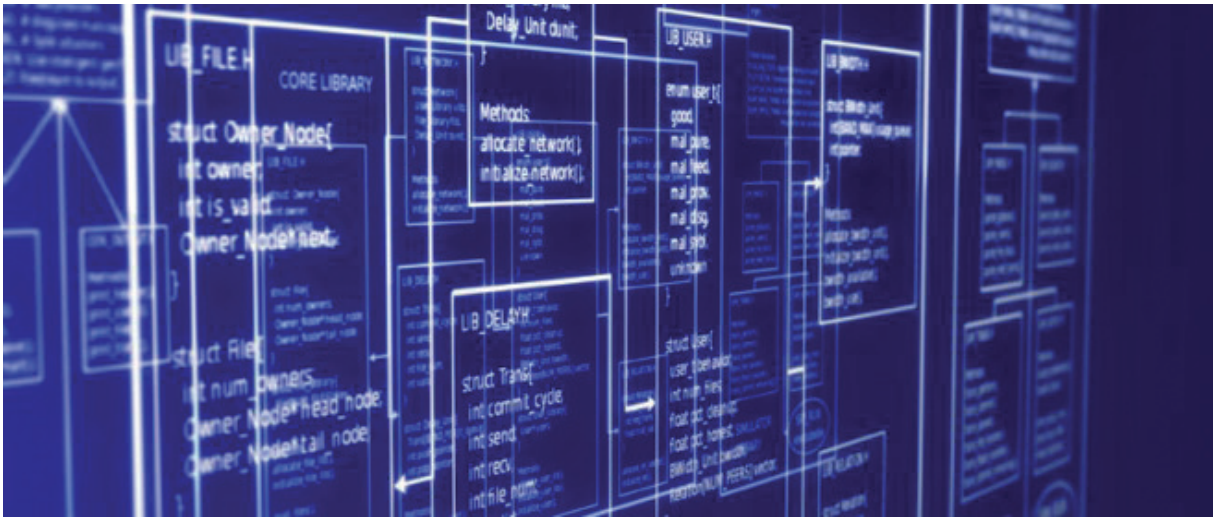

1.3. Legal Status, Governance, and Management

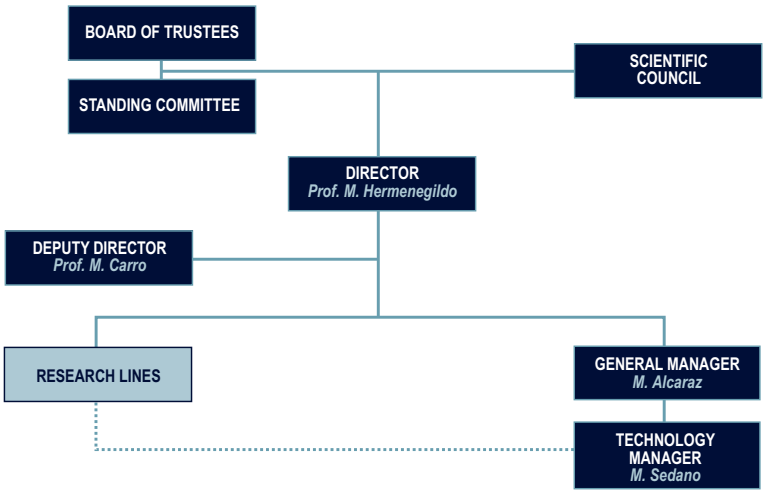

Figure 1.1: Governance and management structure of the IMDEA Software Institute.

The IMDEA Software Institute is a non-profit, independent organization, constituted as a Foundation. Its structure brings together the advantages and guarantees offered by a foundation with the flexible and dynamic management typical of a private body. This combination is deemed necessary to attain the goals of excellence in research, cooperation with industry, and attraction of talented researchers from all over the world, while managing a combination of public and private funds. The Institute was created officially on November 23, 2006 at the initiative of the Madrid Regional Government, following a design that was the result of a collaborative effort between industry and academia, and started its activities during 2007.

The main governing body of the Institute is the **Board of Trustees**. The Board includes representatives of the Madrid Regional Government, universities and research centers of Madrid, scientists with high international reputation in software development science and technology, and representatives of companies, together with independent experts.

The Board is in charge of guaranteeing the fulfillment of the foundational purpose and the correct administration of the funds and rights that make up the assets of the Institute, maintaining their appropriate returns and utility. The Board normally meets twice a year. In the interim Board-level decisions are delegated to the **Standing Committee** of the Board. The Board appoints the **Director**, who is the CEO of the Institute, among scientists with a well-established international reputation in software development science and technology. The Director fosters and supervises the activities of the Institute, and establishes the distribution and application of the available funds among the goals of the Institute,

within the patterns and limits decided by the Board of Trustees. The Director is assisted by the **Deputy Director** and the **General Manager**, who take care of the legal, administrative, and financial activities of the Institute. They are, in turn, helped by the **Technology Manager**, who is in charge of handling project preparation and development, industrial relations, and technology transfer. The current structure is depicted in Figure 1.1.

The Board of Trustees and the Director are assisted in their functions by the **Scientific Council** composed of high-level external scientists of international reputation with expertise in different areas of research covered by the Institute. The tasks of this scientific council include: to provide advice on and approve the selection of researchers; to provide advice and supervision in the preparation of yearly and longer-term strategic plans; to evaluate the performance with respect to those plans; and to give general advice on matters of relevance to the Institute.

1.4. Headquarters Building

A major milestone in 2012 has been the completion of the construction of the Institute's new building in the Montegancedo Science and Technology Park, along with the installation of all interior utilities and furnishings. After a period of final building checks, tests, and certifications, and obtaining the corresponding license, the Institute moved into this permanent location in January 2013. The new building offers an ideal environment for fulfilling its mission of research and technology transfer. It has allowed gathering all personnel, that was previously split in two locations, under one roof and with better and more efficient infrastructures, and allows tackling strategic initiatives. As an example of the latter, it is expected to host the headquarters of the Associate Partner Group in Spain of the European Institute of Technology ICT Labs, coordinated by the Institute, as well as joint research activities and units with industry and academia. Prior to the move, during 2012, the Institute continued to be temporarily located in a renovated floor of the nearby School of Computer Science of the Technical University of Madrid (UPM), also within the UPM's Montegancedo Science and Technology Park, and other temporary premises.

The location of the new IMDEA Software building provides excellent access to the UPM Computer Science Department as

well as to the other new research centers within the Montegancedo Park. These centers include the Madrid Center for Supercomputing and Visualization (CESVIMA), the Montegancedo Campus UPM company "incubator," the Institute for Home Automation (CEDINT), the Institute for Biotechnology and Plant Genomics (CBGP), the Center for Biomedical Technology (CTB), the Microgravity Institute, and the Spanish User Support and Operations Centre for ISS payloads (USOC). In particular, the CESVIMA houses the second largest supercomputer in Spain and one of the largest in Europe, as well as a state-of-the-art visualization cave, and is equipped for massive information storage and processing, high performance computing, and advanced interactive visualization. A number of additional research and technology transfer facilities are currently finishing construction in the campus.

The new site enjoys the benefits of all the convenient new infrastructures that have been completed recently around the campus, such as the recently opened "Montepríncipe" stop of the Madrid Underground. The campus recently obtained the prestigious "International Campus of Excellence" label, and is the only campus in Spain to receive a "Campus of Excellence in Research and Technology Transfer" award in the Information and Communications Technologies area from the Spanish government.

1.5. Appointments to the Board of Trustees

At the beginning of 2012, Jorge Sáinz moved to a position in the Ministry of Education of the Spanish Government and Juan Ángel Botas, professor at Rey Juan Carlos University, replaced him in his position as Deputy Director of Research at the Comunidad de Madrid and as member of the Board of Trustees and Chair of the Standing Committee.

Also during 2012, Prof. Jesús González Barahona was appointed to the Board as the new representative of *Universidad Rey Juan Carlos*, the seat held by David Ríos until the end of his term in late 2011.

1.6. Members of the Governing Bodies

Board of Trustees

Chairman of the Foundation

Prof. David S. Warren
State University of New York at Stony Brook, USA.

Vice-chairman of the Foundation

Excma. Sra. Dña. Alicia Delibes Liniers
Vice-counselor for Education, Madrid Regional Government, Spain.

Madrid Regional Government

Excma. Sra. Dña. Alicia Delibes Liniers
Vice-counselor for Education, Madrid Regional Government, Spain.

Ilmo. Sr. D. José María Rotellar García
Vice-counselor of the Treasury, Madrid Regional Government, Spain.

Prof. Jon Juaristi
Deputy Director for Research, Madrid Regional Government, Spain.

Prof. Juan Ángel Botas
Assistant Director of Research, Department of Education, Madrid Regional Government, Spain. Chairman of the Standing Committee.

Universities and Public Research Bodies

Prof. Narciso Martí Oliet
Professor, Universidad Complutense de Madrid, Spain.

Prof. Javier Segovia Pérez
Dean of the School of Computer Science, Universidad Politécnica de Madrid, Spain.

Prof. Carmen Peláez Martínez
Consejo Superior de Investigaciones Científicas, Spain.

Prof. Jesús M. González Barahona
Subdirector de Infraestructura Tecnológica, Universidad Rey Juan Carlos, Madrid, Spain.

Scientific Trustees

Prof. David S. Warren
State University of New York at Stony Brook, USA. Chairman of the Board of Trustees.

Prof. Patrick Cousot
École Normale Supérieure de Paris (ENS), France and Courant Institute, New York University, USA.

Prof. Luis Moniz Pereira
Universidade Nova de Lisboa, Portugal.

Prof. José Meseguer
University of Illinois at Urbana Champaign, USA.

Prof. Roberto Di Cosmo
Université Paris 7, France.

Expert Trustees

Mr. José de la Sota Rius
Managing Director, Fundación para el Conocimiento (madri+d), Madrid, Spain.

Mr. Eduardo Sicilia Cavanillas
Escuela de Organización Industrial, Madrid, Spain.

Industrial Trustees

Board meetings have been attended, as invitees, by representatives of the following companies:

BBVA
Ms. María Carmen López Herranz. Innovation - Global Observatory & Portfolio Director at BBVA.

Telefónica I+D
Mr. Francisco Jariego, Director for Technology Strategy at Telefónica R&D.

Deimos Space
Mr. Miguel Belló Mora, General Director and Mr. Carlos Fernández de la Peña.

Atos
Mr. José María Cavanillas, Director Research & Innovation, and Ms. Clara Pezuela.

Secretary

Mr. Alejandro Blázquez Lidoy

Scientific Council

Prof. David S. Warren
State University of New York at Stony Brook, USA. Chairman of the Board.

Prof. María Alpuente
Universidad Politécnica de Valencia, Spain.

Prof. Roberto Di Cosmo
Universitè Paris 7, France.

Prof. Patrick Cousot
École Normale Supérieure de Paris (ENS), France and Courant Institute, New York University, USA.

Prof. Veronica Dahl
University Simon Fraser, Vancouver, Canada.

Prof. Herbert Kuchen
Universität Münster, Germany.

Prof. José Meseguer
University of Illinois at Urbana Champaign, USA.

Prof. Luis Moniz Pereira
Universidade Nova de Lisboa, Portugal.

Prof. Martin Wirsing
Ludwig-Maximilians-Universität, München, Germany.

Industrial and Institutional Partnerships

annual report

2012

2.1. Industrial Partnerships

As mentioned before, one of the most successful and cost-effective ways of increasing the competitiveness of industry, and thus contributing to sustainable growth and employment, is by incorporating into processes and products new scientific results and technology. As a generator of such knowledge and technology, in an area with high potential economic impact, IMDEA Software collaborates with industry in a variety of ways in order to foster technology transfer.

A very important way in which this is carried out is through focused collaborations with companies in the framework of *competitive funded collaborative projects* and *direct contracts*. These instruments represent an excellent vehicle for performing joint research and pushing it down the path to products. Figure 2.1 lists some of the companies with which the IMDEA Software Institute has collaborated to date in such projects and contracts (the currently active ones are described further in Chapter 5).

The Institute has also established long-term, *strategic partnerships* with the main stakeholders in the sector which facilitate longer-term collaboration across projects. In particular, the Institute has established close ties with Telefonica, Indra, Atos, and BBVA, which have led to a number of strategic initiatives. For example, the Institute has led the process of presentation of a joint candidacy (with UPM, BSC, Atos, Telefónica, INDRA and BBVA) for the extension of the European Institute of Technology (EIT) ICT Labs to Spain. In 2012, the Institute completed the formal application process for becoming an Associate Member of EIT ICT Labs in order to start in 2013 the coordination of the tasks towards building the corresponding Associate Partner Group in Madrid. The Institute is also working on the establishment of *Joint Research Units* with its strategic partners (for example, such a unit is in development with Telefonica). Also along the strategy line, the Institute participates jointly with industry in Spanish and EU *Technology Platforms*, such as the Technology Clusters of the Madrid Region, the INES Spanish Platform for Software and Services, the Internet of the Future Es.Internet Spanish platform, and the European Technology Platform for Software and Services. All these activities contribute towards aligning research agendas and facilitate joint participation in projects.

Another important form of technology transfer is the *commercialization of technology* developed at the Institute. Given the controversy around software patents (and the impossibility of filing software patents in Europe) the Institute is combining the protection of its intellectual property (for example, a *software registration* has been made for the ActionGUI technology developed by the Institute's Modeling Lab) with other innovative business models, such as those based on open-source or free software licenses, together with the licensing of such technology and the *creation of technology-based companies*. In this line the Institute is actively working on the commercialization of the ActionGUI technology, in collaboration with ETH Zurich.

Project/Contract	Funding Entity	Industrial Partners
MOBIUS	FP6: IP	France Telecom, SAP AG, Trusted Labs
HATS	PF7: IP	Fredhopper
NESSoS	PF7: NoE	Siemens, ATOS
ES_PASS (Through associated group at UPM)	ITEA2, MITyC	Airbus France, Thales Avionics, CS Systèmes d'Information, Daimler AG, PSA Peugeot Citroen, Continental, Siemens VDO Automotive, EADS Astrium, GTD, Thales Transportation, and IFB Berlin
EzWeb	MITyC	Telefónica I+D, Alimerka, Integrasys, Codesyntax, TreeLogic, Intercom Factory, GESIMDE, Yaco, SoftTelecom
DESAFIOS-10	MICINN	BBVA-GlobalNet, LambdaStream, Deimos Space
PROMETIDOS	Madrid Regional Government	Deimos Space, Atos Origin, BBVA-GlobalNet, Thales, Telefónica I+D
MTECTEST	Madrid Regional Government	Deimos Space
SEIF award (Gotsman)	Microsoft SEIF	Microsoft Research
SEIF award (Marron)	Microsoft SEIF	Microsoft Research
PhD Funds	Microsoft	Microsoft Research
ENTRA	FP7: STREP	XMOS
VARIES	FP7: ARTEMIS	Barco NV, HI iberia, IntegraSys, Tecnalia, Sirris, Spicer, Fraunhofer Gesellschaft, Pure-Systems Gmbh, STiftelsen Sintef, Autronica, Franders Mechatronics Technology Centre, Atego Systems, Teknologia Tutkimuskeskus VTT, Mobisoft, HIQ Finland, Softkinetic Sensors, Macq, Metso Automation.
4CaaST	FP7: IP	Telefónica I+D, SAP, France Telecom, Telecom Italia, Ericsson, Nokia-Siemens, Bull SAS, 2nd Quadrant, Flexiant
Contract	AbsInt	AbsInt GmbH
Contract	Boeing	Boeing Research & Technology Europe
Contract	Telefonica	Telefonica Digital

Figure 2.1: Some companies with which the IMDEA Software Institute has collaborated through projects and contracts to date.

Other forms of industry collaboration include the *funding by industry of research assistantships* at the Institute (doctorate or masters work) on industry-relevant topics (for example, Microsoft funding for systems software verification), *transfer of research personnel trained at the Institute to companies* (IMDEA Software-trained personnel has already been transferred to companies such as Atos, Microsoft, Google, or Logicblox), funding by industry of *research stays of Institute researchers at company premises* (for example, Institute researchers have made industrially-funded extended stays at Deimos Space, Microsoft Redmond in the US, or Microsoft Cambridge in the UK, and a framework agreement has been signed with Microsoft for this purpose), *access to the Institute's technology and scientific results* (for example, researchers of the Institute have met with personnel from BBVA, Telefónica I+D, Ericsson, GMV, INDRA, IBM, Canal de Isabel II, Interligare, or Lingway, among many others, to present their main research results), access to the Institute's researchers as consultants, participation of company staff in Institute activities, etc.

2.2. Cooperation with Research Institutions

The Institute also has strong collaborations with universities and other research centers, in the Madrid region and beyond. Again, an important way in which such cooperation happens is through focused collaborations in the framework of *competitive funded collaborative projects*. However, the Institute has also established *longer-term, strategic partnerships through agreements* with a number of institutions in order to reach objectives that go beyond individual projects. At present the Institute has already signed agreements with the following universities and research centers:

- Universidad Politécnica de Madrid (from November 2007).
- Universidad Complutense de Madrid (from November 2007).
- Universidad Rey Juan Carlos (from January 2008).
- Roskilde University, Denmark (from June 2008).
- Consejo Superior de Investigaciones Científicas (from November 2008).
- Swiss Federal Institute of Technology (ETH) Zurich (from November 2012).
- Microsoft Research (from December 2012).

These agreements establish a framework for the development of collaborations and include the joint participation in and development of graduate programs; the joint use of resources, equipment, and infrastructure; exchange of staff; joint participation in research projects; the association of researchers and research groups with the Institute; or the joint commercialization of technology. In particular, research assistants at the IMDEA Software Institute can follow graduate studies at any of the cooperating Institutions, while funded by the IMDEA Software Institute and Institute faculty collaborate in those graduate programs.

To illustrate the scope and importance for the Institute of these agreements, we offer here some highlights. The agreement with the Universidad Politécnica de Madrid includes provisions for the location of the Institute building in its Montegancedo Science and Technology Park as well as a joint graduate program, instrumented currently as a separate track on *Software Development through Rigorous Methods* in an existing Masters / PhD program at UPM ("MUSS / DSS"). Under the agreement with the Consejo Superior de Investigaciones Científicas, two of its researchers —Cesar Sánchez and Pedro López— are also part of the research staff of the Institute. Under the agreement with Roskilde University, one of its full professors —John Gallagher— is also part-time senior researcher at the Institute. As mentioned before, the agreement with ETH Zurich includes the joint development and commercialization of the ActionGUI technology, from the Institute's Modeling Lab. The Institute also manages for the Regional Government REDIMadrid, the Madrid academic network which connects all the public universities and other research institutions in the Madrid region to each other and to the national and international backbones, and has signed agreements in this context

with all such institutions. Finally, the Institute of course also collaborates with the other Institutes in the IMDEA network. As an example, the Institute has secured and coordinates the AMAROUT-I and AMAROUT-II COFUND programs of Marie Curie fellowships, which fund personnel in all the IMDEA Institutes, and provides other services to the IMDEA network.

The Institute also has a strong presence in national and international bodies. It is a member of *Informatics Europe*, the organization of all deans, chairs, and directors of the leading Departments and Institutes of Computer Science in Europe, similar to the CRA in the US. In addition, the Institute is a member of ERCIM, the *European Research Consortium for Informatics and Mathematics* through SpaRCIM, the Spanish representative in ERCIM, where Manuel Hermenegildo, IMDEA Software Institute Director, is also the President of the Executive Board.

Research

As briefly explained in Section 1.2, the cost-effective development of complex, safe, reliable, and efficient software is not a simple task, and it cannot be solved by simple "magic bullets" or more enlightened management. The problem affects all stages in the development lifecycle (analysis, design, implementation, verification, maintenance). The IMDEA Software Institute performs research on these aspects along a number of dimensions which include *Methodologies* (the development and industrial adoption of mathematically rigorous methodologies can improve the software process further), *Languages* (the basis for expressing software functionality, behavior, and properties), *Verification and Validation* (semantically well-founded, tool-supported methods to validate code or designs with respect to specifications), and *Adequacy/Optimization* (the optimal use of resources to achieve a desired goal). To this end:

- The research vision materializes in a number of *High-level Research Lines*. The current main lines are depicted as rows in Figure 3.1.
- The vision includes also a number of focusing *Areas of Application*: areas of engineering where the Institute aims and expects to make an impact and which have been identified as priorities in collaboration with industry. The main current areas of application are depicted as columns in Figure 3.1.

These areas of application and high-level research lines are explained further in the rest of the chapter. Finally, two fundamental, cross-cutting issues pervade the vision:

- *Tools*: well-founded and cost-effective (prototypes of) tools are fundamental study harnesses, demonstrators, and technology transfer vehicles for the techniques for automation of high quality software development.
- *Foundations*: methods and languages should be built on appropriate mathematical foundations, and at the same time be practical.

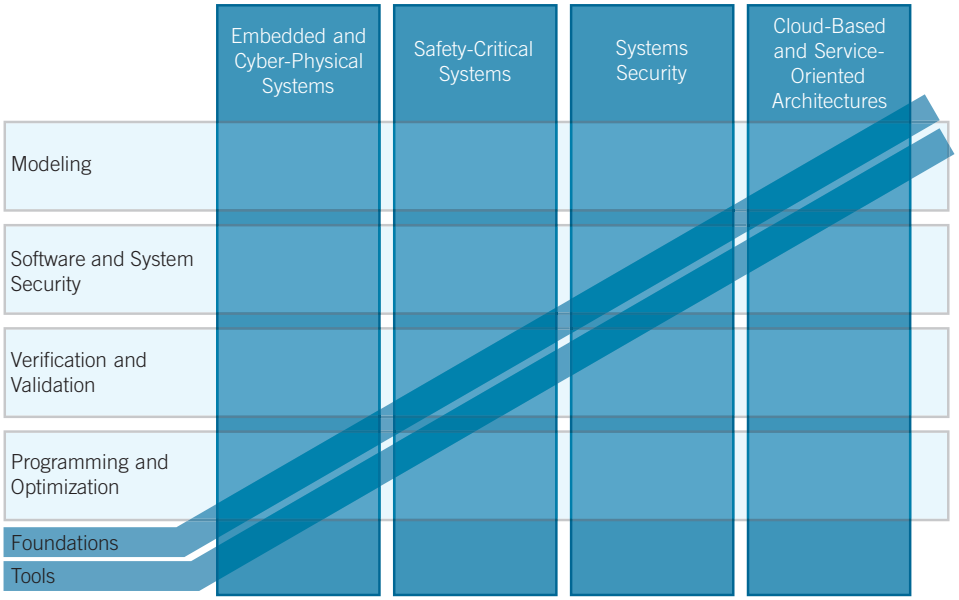

Figure 3.1: Main research lines, application areas, and cross-cutting issues.

3.1. Areas of Application

The following are some areas where the IMDEA Software Institute aims and expects to make an impact.

3.1.1. Embedded and Cyber-Physical Systems

One of the application areas of software where correctness is most critical is embedded systems. An embedded system is a computational artifact that is subject to physical constrains, and whose correct functioning cannot depend on human guidance. These systems are referred to as cyber-physical systems when there is a tight coupling between the networked computation and physical elements. Embedded and cyber-physical systems are involved in safety-critical applications (such as control systems of automobiles or aircraft) or systems for highly remote operation (satellite, space, etc.). Also, embedded systems are pervasive in areas of high economic impact, like mobile telephony or consumer electronics. These systems must be resource-aware and are often also real-time systems. This means that the

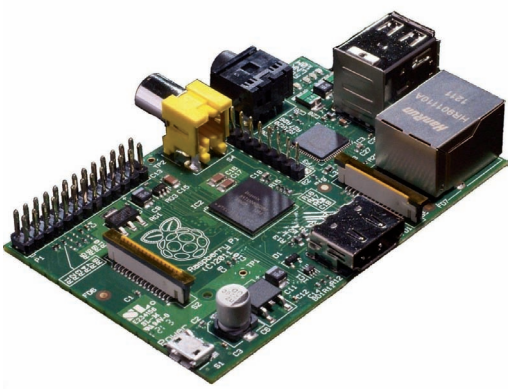

Resilience and low energy consumption are critical for autonomous functioning of embedded systems.

computation must be correctly performed within its time constraints, and also with an adequate use of resources. They also often need to take into account the behavior of physical elements. There is a clear need for rigorous techniques that can improve the quality of embedded and cyber-physical software, or the time to market of new devices or families of devices. Most of the research activities required by embedded and real-time systems and carried out at the IMDEA Software Institute are aligned with the Strategic Research Agenda of the European Technology Platform on Embedded Computing Systems, ARTEMIS.

3.1.2. Safety-Critical Systems

Software is becoming pervasive in areas such as transportation (avionics, automotive), health (diagnosis, therapy), and control (of nuclear plants, of railway signaling systems, of conflict detection systems), where a failure or malfunction may be extremely damaging, even in terms of human lives. The constraints for such safety-critical systems are extremely stringent: the systems must be able to function during extremely long period of times, in presence of human mistakes or hardware or software failures, and provide an acceptable level of services at all times. Thus, it is urgent to develop methods and tools that help support the development of fault-tolerant, resilient, and adaptable software and its (quantitative) evaluation against the aforementioned constraints. A particular challenge is to scale existing methods so that they become effective in the context of distributed and networking systems.

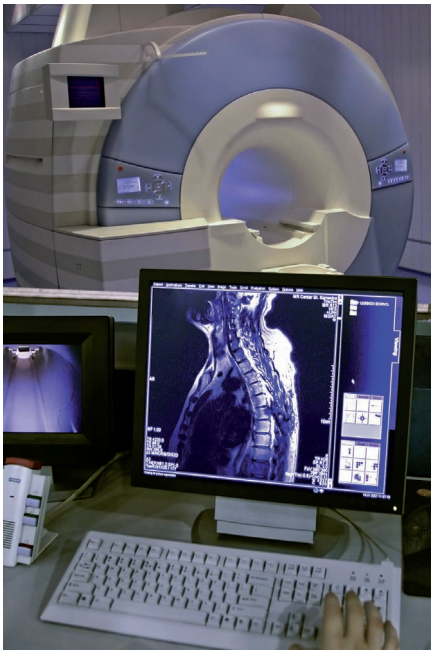

Modern devices, from cars to TAC scanners, completely depend on software working correctly. Failures or malfunctions can bring about from wrong diagnoses to fatal outcomes.

3.1.3. Systems Security

As our society increasingly relies on information technology, there is an urgent and unprecedented need to develop new security mechanisms for protecting infrastructures, data, and applications. Several concomitant factors aggravate the problems of information security. In order to face this challenge, one must provide scalable and rigorous techniques that can be integrated in prevailing software development processes to enforce security of applications. Since many attacks arise at the application level, it is particularly important to achieve security at the level of programming languages, drawing from methods developed in programming language research (design, analysis, and verification), and developing security solutions at a level of abstraction that matches the programming language. At the same time it is important to target a wide range of security- and privacy-related issues and scenarios, from whole systems to (big) data, and on the current plethora of platforms, ranging from apps to the cloud itself.

3.1.4. Cloud-Based and Service-Oriented Architectures

Cloud computing represents a novel, fast growing paradigm for providing massively distributed, flexible, and scalable environments for large-scale elastic provision of computation resources (such as virtual machines and storage), web content, dynamic computing environments (platforms), and application software. By providing a spectrum of computation resources as a service, cloud computing naturally extends the concept of Service-Oriented Computing (SOC) and facilitates its real-world implementation. One of the goals of SOC is to provide the necessary support for effectively programming, deploying, and maintaining services over highly-distributed networks. Cloud computing and SOC draw from many areas of computer science, including software engineering, concurrent and distributed systems, and modular and component-based programming. While these areas are well developed in isolation, there remain significant challenges to combine the methodologies that stem from each area in order to deliver cost-effective approaches that support the construction and deployment of electronic services. Cloud computing is a rapidly growing sector of IT worldwide, and has been identified as one of the strategic research and development priorities by the European Commission and international IT analysts.

Under the hood, cloud computing uses complex technologies to provide scalable access to computing resources and services for a potentially large number of users.

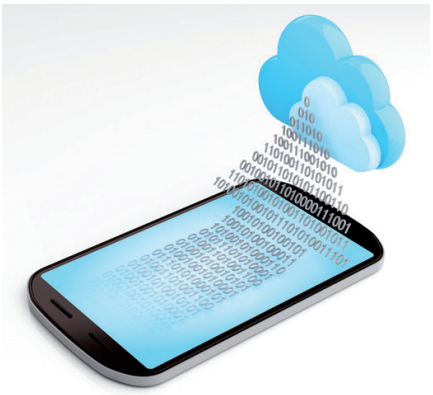

Cloud computing allows devices with restricted processing power, storage, or energy consumption to use virtualized processing power and resources.

3.2. Research Lines

3.2.1. Modeling

A model is an abstraction of some aspect of a system (like a blueprint in engineering), which is created to serve particular purposes, for example, to present a human-understandable description of some aspects of the system or to present information in a form that can be mechanically analyzed. The term Model-Driven Engineering (MDE) is used to describe software development approaches in which abstract models of software systems are created and systematically transformed to obtain concrete implementations or skeletons. Model-driven development holds the promise of reducing system development time and improving the quality of the resulting products.

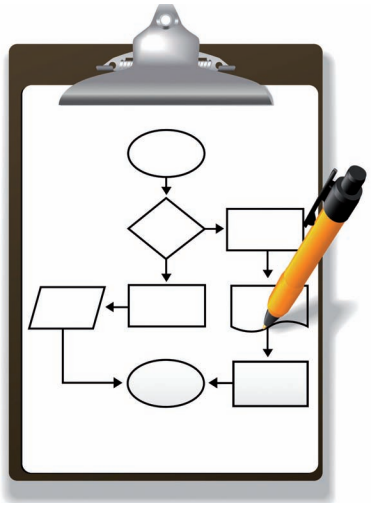

However, in mainstream MDE practice, models are usually informal, with no well-established semantics, and only used for documentation purposes. In fact, modeling has traditionally been a synonym for producing diagrams. Most models consist of a number of "bubbles and arrows" pictures and some accompanying text. The information conveyed by such models has a tendency to be incomplete, informal, imprecise, and sometimes even inconsistent.

In order to address the major challenges current MDE technologies are facing, we believe that the past and present work on formal methods is particularly relevant. Many of the

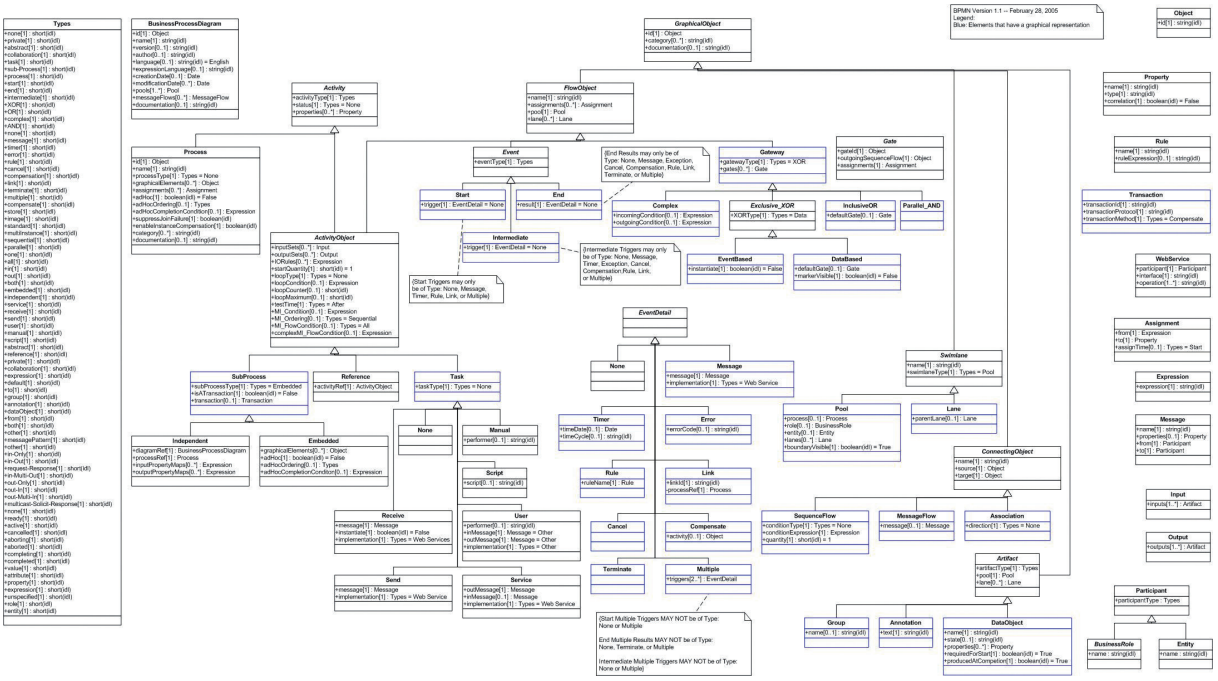

flaws in modeling are caused by the limitations of the diagrams being used. A diagram simply cannot express some of the essential information of a thorough specification. To specify software systems, formal languages offer some clear benefits over the use of diagrams. Formal languages are unambiguous, and cannot be interpreted differently by different people, for example, an analyst and a programmer. Formal languages make a model more precise and detailed, and can be manipulated by automated tools to ensure correctness and consistency with other elements of the model. On the other hand, a model completely written in a formal language is often not easily understood. In this sense, we believe that the interaction between the MDE and formal methods communities has a huge potential impact.

At the IMDEA Software Institute we are providing rigorous semantics for current MDE technologies (e.g., OCL, QVT) and we are developing tool-supported methodologies for applying these technologies for building *meaningful* models: i.e., models that have a clear and rich meaning, and that are therefore useful and valuable for developing quality software. At the same time, we are proposing new MDE technologies for specific areas of applications, including software and system security and graphical user interfaces, such as our ActionGUI technology.

3.2.2. Software and System Security

The goal of this line is to develop methods and tools that provide an accurate security analysis of systems and software, together with some countermeasures to defeat malicious agents.

While software security traditionally focuses on low-level protection mechanisms such as access control, the popularization of massively distributed systems dramatically increases the number and severity of vulnerabilities at the application level. These vulnerabilities may be exploited by malicious software such as viruses, Trojan horses, etc., but also (unintentionally) by buggy software, with disastrous effects.

Language-based security aims to achieve security at the level of the programming language, with the immediate benefit of countering application-level attacks at the same level at which such attacks arise. Language-based security is attractive to programmers because it allows them to express security policies and enforcement mechanisms within the programming language itself, using well-developed techniques that facilitate a rigorous specification and verification of security policies.

Language-based techniques can guarantee a wide range of policies including confidentiality, integrity, and availability, and their combination. However, their practical adoption has been hindered partly because known enforcement methods are confined to simple policies, such as non-interference for confidentiality. The most pressing challenges are defining unified enforcement mechanisms that support flexible and customizable policies, and developing methods for providing a quantitative assessment of security.

The IMDEA Software Institute is developing rich policy languages that capture precisely common instances of information release. Moreover, these policy languages are directly applicable to powerful abstraction mechanisms that pervade modern programming languages. These policy languages are supported by automated verification procedures, that allow users to detect fraudulent software.

We are also developing accurate methods for a quantitative evaluation of program security. These methods account for covert channels, including timing behavior and resource consumption, and for resistance to common attacks, such as viruses. The ultimate goal is to develop comprehensive adversarial models and effective protection strategies against covert channels.

Language-based methods have been studied primarily for mobile code and very few methods are known to scale to distributed systems. One main challenge is to ensure security of distributed applications, using a combination of cryptographic and language-based methods. Programming language techniques provide an attractive approach to

guarantee the security of distributed software, because they allow reasoning about programs and their cryptographic libraries in a unified framework. Moreover, rigorous programming language-based techniques can be used to demonstrate beyond reasonable doubt that standard cryptographic systems, some of which have a long history of flawed security proofs and hidden but effective attacks, are secure.

The IMDEA Software Institute is building tools that support the automated analysis of cryptographic systems and provide very strong guarantees of their correctness (cryptographic strength). The tools adopt the game-playing technique, that organizes the construction of cryptographic proofs as sequences of probabilistic games as a natural solution for taming the complexity of performing cryptographic proofs. The tools have been validated experimentally through the verification of widely deployed cryptographic standards.

Proving correctness of cryptographic protocols eliminates the possibility of compromising confidential information or forging false credentials.

3.2.3. Verification and Validation

Verification refers to the rigorous demonstration that software is correct; that is, it provides behavioral consistency according to a given specification of its intended behavior. By "intended behavior" we mean the properties that software is expected to satisfy when it is deployed. Software not possessing the properties might be defective: its execution might have unintended consequences. *Verified software* is software that is free of certain classes of defects because it has been rigorously proven that it satisfies its intended behavior. For these particular classes of defects, the verified software is termed zero-defect software. Such software does not require disclaimers that forgive developer error. Instead such software is guaranteed to be reliable — it behaves as intended.

How do we "rigorously prove" that software is correct? The basic principle is to represent properties as logical formulas so that verification of the properties is akin to proving a theorem using proof techniques from mathematical logic. However, modern software is very complex and typically composed of several components, where each component

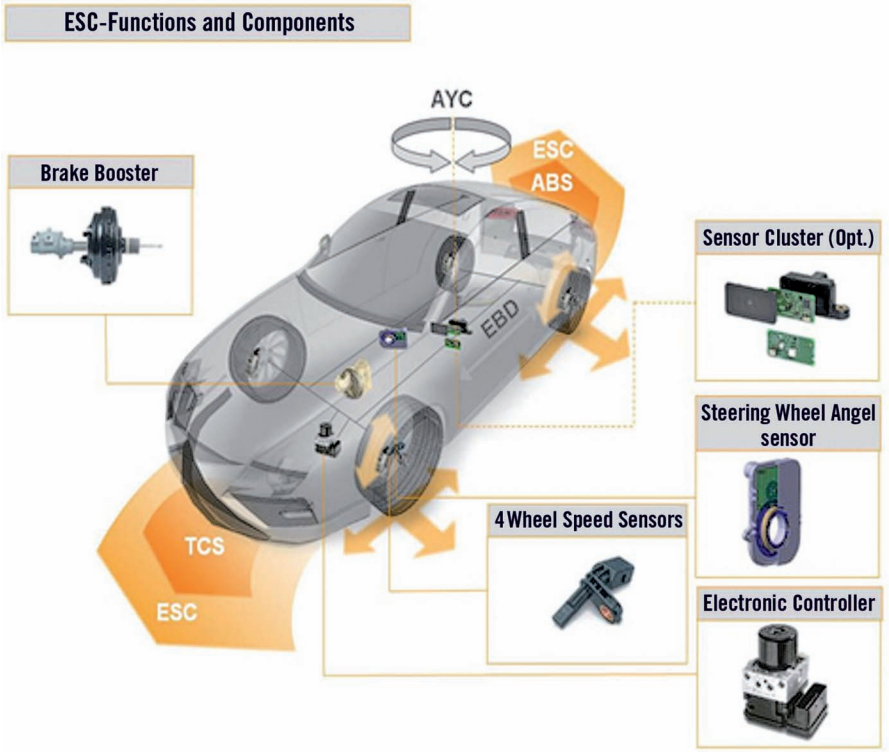

Modern cars and trucks contain as many 100 million lines of computer code. This software runs on more than 30 on-board computers and controls vital functions, including the brakes, engine, cruise control, and stability systems. It is under increasing scrutiny in the wake of recent problems with major manufacturers and it is currently impossible to fully test.

can be written in a different programming language. For such complex software, proving properties manually is very difficult. The question that arises naturally is this: can the logic-based proof techniques be made to scale so that software can be automatically verified as much as possible so that the manual verification burden is minimized? Apart from managing the complexity of proofs, the benefits of automatic verification are as follows. First, the verification can be repeated whenever necessary and with the same results, thus attesting to the accuracy of verification. Second, proofs can be mechanically checked for correctness. Third, verification results can be reused: once a program has been verified, its specifications can be repeatedly used in verification of a larger piece of software without re-verification.

Researchers at the IMDEA Software Institute are involved in various aspects of automatic software verification. They study expressive languages and logics for specification of properties of software, particularly of software written in modern programming lan-

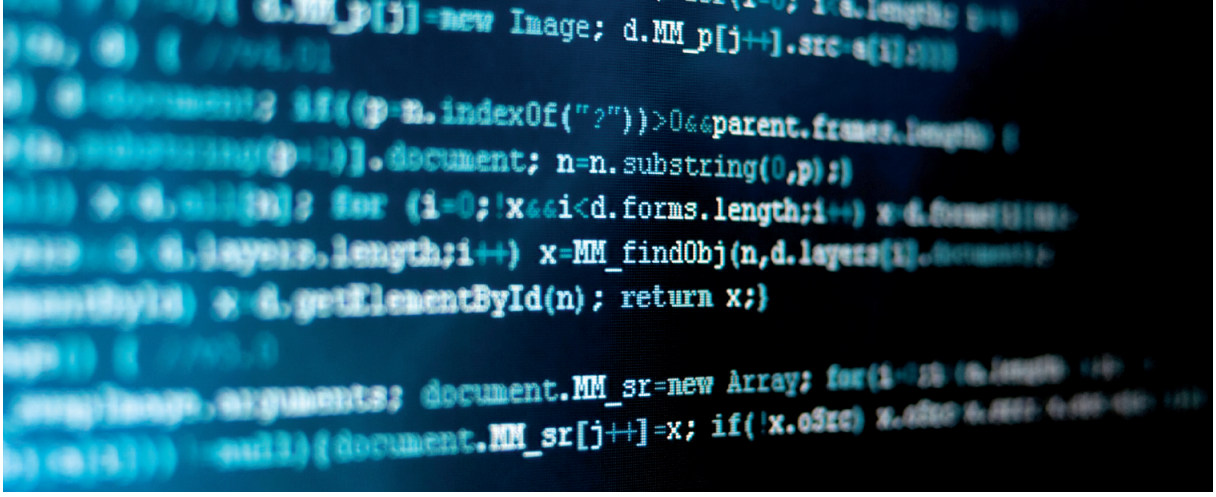

guages such as Java. Once a Java program is decorated with such specifications, off-the-shelf verifiers can be used to generate "verification conditions" which can then be discharged by theorem provers. Researchers are not only studying more efficient verification algorithms and decision procedures for improving theorem proving technology, but also are performing experiments on verifying realistic code such as Java libraries — whose programs are frequently used in building complex software — and design patterns, which provide generic solutions to common software problems. The automated proofs will be made publicly available in a repository linked to the Verified Software Repository of the international Verification Grand Challenge Project.

3.2.4. Advanced Programming and Optimization Tools

The goal of this line is to develop methods and tools that help programmers improve the quality and robustness of the programs they write, allow them to write better programs in a shorter time, and support efficient execution of code through highly optimizing compilers.

Programming is notoriously difficult and error prone. Current tools assist programmers in their task. However, much more sophisticated technology is still needed in order to reduce the time to market while actually increasing the degree of correctness of the delivered code beyond what is possible today.

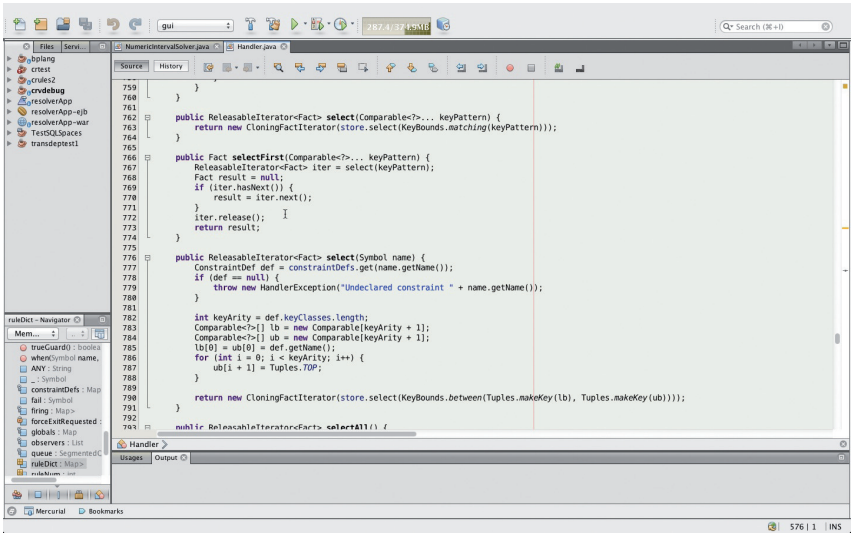

The aims of program correctness and robustness are similar to those in verification, but the focus here is on tools that find errors and verify programs *during the process of writing such programs*, rather than a posteriori. This focus requires efficient and fully automatic program analysis methods.

Abstraction-based techniques provide a unifying framework for this purpose. Their essence is abstract interpretation, a rigorous method which induces a dramatic reduction in the complexity of software analysis. It has been shown powerful enough to, for example, analyze automatically avionics software, a clear example of a large cyber-physical system, consisting of millions of lines of code, and subject to stringent conditions from the DO-178B standards. Researchers at the IMDEA Software Institute are developing tools that show that abstraction techniques can be embedded in development environments for routine use by programmers for on-line debugging, diagnosis, verification, and certificate generation, and that they combine naturally with (and reduce the need for) other techniques such as testing and run-time verification, which currently take more than 90% of overall development cost.

Abstraction-based techniques have also been shown particularly effective for high integrity and embedded software, where the properties of concerns are time and memory consumption, dynamic data sizes, energy consumption, termination, absence of errors or exceptions, etc. Researchers at the IMDEA Software Institute are developing advanced tools for debugging and verification of software with respect to these non-functional properties.

Another important way of improving the programming process, which allows programmers to write better programs in a shorter time, is by improving programming languages. Researchers at the IMDEA Software Institute are working on promising approaches such as extensible and multi-paradigm languages, support for domain-specific languages, support for multi-language applications, and service-oriented architectures.

Regarding the objective of supporting the efficient execution of code, abstraction-based techniques can also be used to ensure that programs are highly optimized before execu-

tion, i.e., that they run in the fastest and most resource-efficient way on the platforms and environmental conditions they are deployed on, while maintaining their observable behavior. Typical goals include saving on memory and processing time on sequential processors, adaptive task scheduling in parallel and distributed computers, self-reconfiguration, and automatic adaptation to environmental conditions.

A prominent form of such program optimization is automatic parallelization. As highly parallel processors are becoming an inexpensive and common facility in mainstream computing, there is an opportunity to build much faster, and eventually much better, software. Yet exploiting this enormous potential requires the development of new programming practices that reflect this profound change in the execution paradigm. Two common alternatives are to write parallel programs, using dedicated programming idioms and algorithms that help taming the complexity of parallel programs, or to automatically parallelize existing ones, using compilers for identifying parts of the application that are independent and can thus be run in parallel. Researchers at the IMDEA Software Institute are working on both approaches, developing languages and idioms more suited for parallelism and abstraction-based techniques and tools for allowing detection of common errors in parallel programs and for automatic parallelization of programs.

People

annual report 2012

The IMDEA Software Institute strives toward the level of excellence and competitiveness of the highest-ranked institutions worldwide. Success in this goal can only be achieved by recruiting highly-skilled personnel for the scientific teams and support staff. The Institute considers this one of the key critical factors and measures of its success.

Competition for talent in Computer Science is extremely high at the international level, since essentially all developed countries have identified the tremendous impact that IT has on the economy and the crucial competitive advantage that attracting truly first class researchers entails. To meet this challenge, the Institute offers a world-class working environment that is competitive with similar institutions in Europe and in the US and which combines the best aspects of a university department and a research laboratory.

Hiring at the Institute follows internationally-standard open dissemination procedures with public, international calls for applications launched periodically. These calls are advertised in the appropriate scientific journals and at conferences in the area, as well as in the Institute web page and mailing lists. Appointments at (or promotion to) the Assistant Professor, Associate Professor, and Full Professor levels (i.e., tenure-track hiring, tenure, and promotion decisions) are approved by the Scientific Advisory Board. In addition to staff positions, the Institute also has its own programs for visiting and staff researchers, post-docs, graduate scholarships, and internships. In all aspects related to human resources, the Institute follows the recommendations of the European Charter for Researchers and a Code of Conduct for their Recruitment (http://ec.europa.eu/), which it has duly signed.

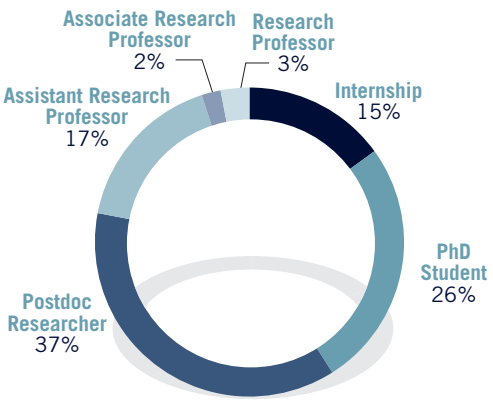

Figure 4.1. Type of position applied for.

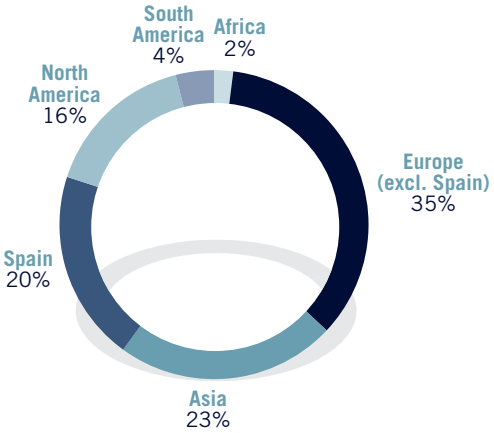

Figure 4.2. Location of previous institution for applicants at or above the postdoc level (by continent + Spain).

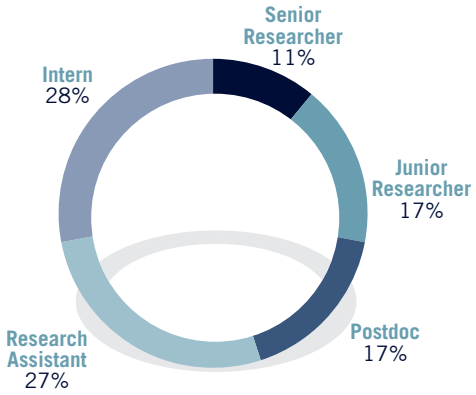

Figure 4.3. Type of position, all researchers.

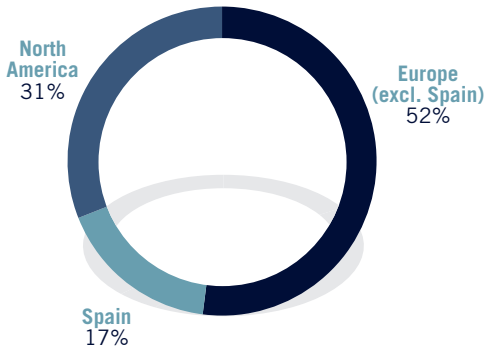

Figure 4.4. Where PhD was obtained (by continent + Spain).

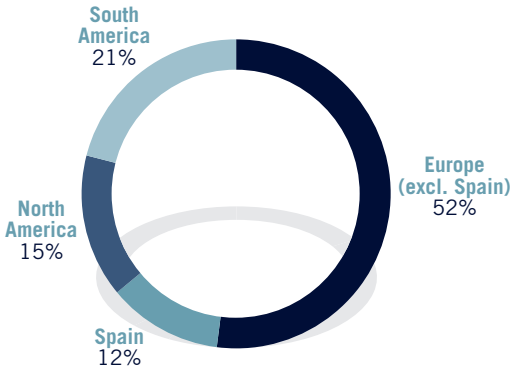

Figure 4.5. Location of previous institution, all (by continent + Spain).

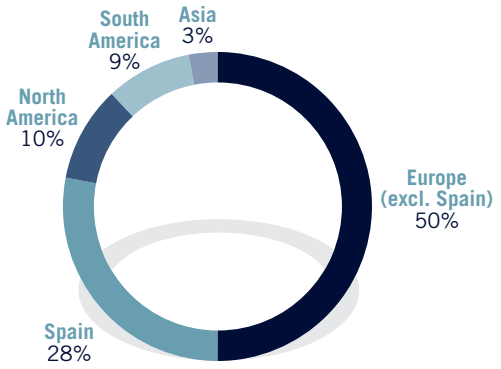

Figure 4.6. Nationality of researchers at or above postdoc level (by continent + Spain).

Applications

Figure 4.1 shows the proportions of applications received for each category during 2012: associate professors (senior researchers), assistant professors (junior researchers), postdoctoral researchers, research assistants, and interns. Figure 4.2 displays the location (by continents) of the institutions in which the applicants were at the time of application (for senior, junior, and postdoctoral positions). Spain is highlighted separately from the rest of Europe to provide a finer view of the data (level of internationalization).

Status

At the end of 2012, the scientific staff of the Institute was composed of seven senior faculty (full or associate professors, one part-time), eleven junior faculty (three non tenure-track), eleven postdoctoral researchers, and seventeen research assistants (PhD candidates). Fifteen interns spent a variable length of time (from one month to half a year) at the Institute collaborating with the faculty members. Additionally, a number of visitors have also been at the Institute during 2012, four of them for extended periods. Figure 4.3 shows the proportions of each category at the end of 2012 (where 28% are faculty members vs. 72% non-faculty). Figure 4.4 summarizes where these researchers obtained their PhD (by continents plus Spain), and Figure 4.5 shows the location where the Institute researchers were working previously to joining IMDEA. Finally, Figure 4.6 presents the nationalities of researchers at or above the postdoc level.

faculty

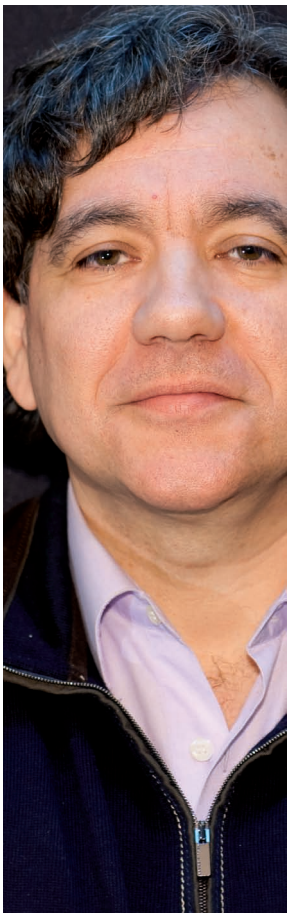

Manuel Hermenegildo

Professor and Scientific Director

Manuel Hermenegildo received his Ph.D. degree in Computer Science and Engineering from the University of Texas at Austin, USA, in 1986. Since January 1, 2007 he is Full Professor and Scientific Director of the IMDEA Software Institute. He is also a full Prof. of Computer Science at the Tech. U. of Madrid, UPM. Previously to joining the IMDEA Software Institute he held the P. of Asturias Endowed Chair in Information Science and Technology at the U. of New Mexico, USA. He has also been project leader at the MCC research center and Adjunct Assoc. Prof. at the CS Department of the U. of Texas, both in Austin, Texas, USA. He has received the Julio Rey Pastor Spanish National Prize in Mathematics and Information Science and Technology and the Aritmel National prize in Computer Science, and he is an elected member of the Academia Europaea. He is also one of the most cited Spanish authors in Computer Science. He has published more than 200 refereed scientific papers and monographs and has given numerous keynotes and invited talks in major conferences in these areas. He has also been coordinator and/or principal investigator of many national and international projects, area editor of several journals, and chair and PC member of a large number of conferences. He served as General Director for the research funding unit in Spain, as well as member of the European Union's high-level advisory group in information technology (ISTAG), and of the board of directors of the Spanish Scientific Research Council and the Center for Industrial and Technological Development, among other national and international duties.

Research Interests

His main areas of interest include programming language design and implementation; abstract interpretation-based program analysis, verification, debugging, and optimization; logic and constraint programming; parallelizing compilers; parallel and distributed processing.

Manuel Carro

Associate Professor and Deputy Director.

Manuel Carro received his Bachelor degree in Computer Science from the Technical University of Madrid (UPM), and his PhD degree from the same University in 2003. He is currently Associate Research Professor and Deputy Director at the IMDEA Software Institute, and an Associate Professor at the Technical University of Madrid. He has previously been representative of UPM at the NESSI and INES technological platforms, and is now representative of UPM at SpaRCIM and deputy representative of IMDEA Software at ERCIM and Informatics Europe. He has published over 70 papers in international conferences and journals, some of which merited the "Best Conference Paper" award. He has been organizer and PC member of many international conferences and workshops and participated in research projects at the regional, national, and European level. He was UPM's principal investigator for the S-Cube European Network of Excellence. He has completed the supervision of three PhD thesis and is actively supervising another one.

Research Interests

His interests span several topics, including the design and implementation of high-level (logic- and constraint-based) programming languages for improving the quality of software production, the analysis of service-based systems, and the effective usage of formal specifications in the process of teaching programming. He has long been interested in parallel programming and parallel implementations of declarative languages, and visualization of program execution.

Gilles Barthe

Professor

Gilles Barthe received a Ph.D. in Mathematics from the University of Manchester, UK, in 1993, and an *Habilitation à diriger les recherches* in Computer Science from the University of Nice, France, in 2004. He joined the IMDEA Software Institute in April 2008. Previously, he was head of the Everest team on formal methods and security at INRIA Sophia-Antipolis Méditerranée, France, and a member of the Microsoft Research-INRIA Joint Centre. He also held positions at the University of Minho, Portugal; Chalmers University, Sweden; CWI, Netherlands; University of Nijmegen, Netherlands. He has published more than 100 refereed scientific papers. He has been coordinator/principal investigator of many national and European projects, and served as the scientific coordinator of the FP6 FET integrated project "MOBIUS: Mobility, Ubiquity and Security" for enabling proof-carrying code for Java on mobile devices (2005-2009). In 2012, he was a PC member of several conferences (POST, MFCS, ASE-Tools, Tools Europe, and STM), and served as PC chair of ESSoS.

Research Interests

Gilles' research interests include program verification, programming languages, software and system security, cryptography, privacy, and foundations of mathematics and computer science. His recent work focuses on computer-aided cryptographic proofs.

Anindya Banerjee

Professor

Anindya Banerjee received his PhD from Kansas State University, USA, in 1995. After his PhD, Anindya was a postdoctoral researcher, first in the Laboratoire d'Informatique (LIX) of École Polytechnique, Paris and subsequently at the University of Aarhus. He joined the IMDEA Software Institute Institute in February 2009 as Full Professor. Immediately prior to this position, Anindya was Full Professor of Computing and Information Sciences at Kansas State University, USA. He was an Academic Visitor in the Advanced Programming Tools group, IBM T. J. Watson Research Center in 2007 and a Visiting Researcher in the Programming Languages and Methodology group at Microsoft Research in 2007–2008. He was a recipient of the Career Award of the US National Science Foundation in 2001. He is an associate editor of the journal Higher-Order and Symbolic Computation.

Research Interests

Anindya's research interests lie in language-based computer security, program analysis and verification, program logics, concurrency, programming language semantics, abstract interpretation and type systems. His primary research activities over the past couple of years have centered around automatic and interactive verification of properties of pointer-based programs and in verification of security properties of such programs.

Juan José Moreno-Navarro

Professor and Director for International and Industrial Relations

Juan José Moreno-Navarro received his Ph.D. degree in Computer Science from the Technical U. of Madrid (UPM), Spain in 1989. He developed a large part of his Ph.D. at RWTH Aachen (Germany). He has been Full Professor at the UPM Computer Science Department since 1996 and joined the IMDEA Software Institute upon its foundation. He served as Deputy Director up to 2008. Currently he is Director of International and Industrial Relations. He has published more than 100 papers in international conferences, books, and journals. He has organized, served in program committees, and given invited talks and tutorials in many conferences in the IST field.

He has been actively involved in research and university policy, serving as General Director for University Policies (Ministry of Education, 2009-2012), General Director for Technology Transfer and Enterprise Development (Ministry of Science and Innovation, 2009) and General Director for Planning and Coordination (Ministry of Science and Innovation, 2008-2009). During this period he has been chairman of CNEAI (Spain's Researcher Activity Evaluation Agency), has been involved in the setting up of several research infrastructures, secretary of the Spanish University Council and of the Spanish Science and Technology Council, and member of the steering board of several foundations (ANECA, UIMP, CENER-Ciemat, Universidad.es, CSIC, ISCII, IDAE, etc.)

He has been the founding director of SpaRCIM. He has been responsible for the Spanish ICT research program, in charge on International Relations for IST, FP6 IST Committee member, and Spanish National Contact Point for the Spanish Ministry of Education and Science. Currently he is chair of the Spanish Society of Software Engineering, general chair of the Spanish Conference of Informatics 2013, and coordinator of the Spanish Turing Year.

Research Interests

His research interests include all aspects related to declarative and rigorous software development technologies. This includes software development techniques, specially specification languages, automatic generation of code (programming from specifications), and software services description. His interests also include the integration of functional and logic programming. He has led the design and implementation of the language BABEL and took part in the activities of the international committee involved in the design of the language Curry. He is also currently involved in scientific policy analysis, bibliometry, and research impact evaluation and analysis.

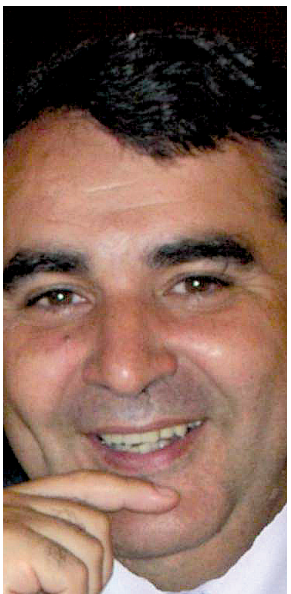

John Gallagher

Professor (part-time)

John Gallagher received the B.A. (Mathematics with Philosophy) and Ph.D. (Computer Science) degrees from Trinity College Dublin in 1976 and 1983 respectively. He held a research assistantship in Trinity College (1983-4), and post-doc appointments at the Weizmann Institute of Science, Israel (1987-1989) and Katholieke Universiteit Leuven, Belgium (1989). From 1984-1987 he was employed in research and development in a software company in Hamburg, Germany. Between 1990 and 2002 he was a lecturer and later senior lecturer at the University of Bristol, UK. Since 2002 he has been a professor at the University of Roskilde, Denmark, where he is leader of the research group Programming, Logic and Intelligent Systems and the Experience Lab as well as (part-time) Professor and holds a dual appointment at the IMDEA Software Institute since February 2007. He is a member of the executive committee of the Association of Logic Programming (2008-2011) and of the steering committee of the ACM SIGPLAN workshop series on Partial Evaluation and Program Manipulation (PEPM). He is an area editor for the journal Theory and Practice of Logic Programming. He has published approximately 50 peer-reviewed papers which have over 1200 citations.

Research Interests

His research interests focus on program transformation and generation, constraint logic programming, rewrite systems, temporal logics, semantics-based emulation of languages and systems, analysis and verification of energy consumption of programs and other properties, and has participated in and led a number of national and European research projects on these topics.

Manuel Clavel

Associate Professor

Manuel Clavel received his Bachelor's degree in Philosophy from the Universidad de Navarra in 1992, and his Ph.D. from the same university in 1998. Currently, he is an Associate Research Professor at the IMDEA Software Institute, as well as an Associate Professor at the Universidad Complutense de Madrid. He was Deputy Director from 2008 until April 2011. During his doctoral studies, he was an International Fellow at the Computer Science Laboratory of SRI International (1994 - 1997) and a Visiting Scholar at the Computer Science Department of Stanford University (1995 - 1997). His Ph.D. dissertation was published by the Center for the Study of Language and Information at Stanford University. Since then, he has published over 30 refereed scientific papers. He has also been involved in the supervision of 3 Ph.D. students (1 completed).

Research Interests

His research focuses on rigorous, tool-supported model-driven software development, including: modeling languages, model transformation, model quality assurance, and code-generation. Related interests include specification languages, automated deduction, and theorem proving.

César Sánchez

Assistant Professor

César Sánchez received his Ph.D. degree in Computer Science from Stanford University, USA, in 2007, studying formal methods for distributed algorithms. After a post-doc at the University of California at Santa Cruz, USA, César joined the IMDEA Software Institute in 2008, becoming a Scientific Researcher at the Spanish Council for Scientific Research (CSIC) in 2009. He holds a degree in Ingeniería de Telecomunicación (MSEE) from the Technical University of Madrid (UPM), Spain, in 1998. Funded by a Fellowship from La Caixa, he moved to Stanford University, USA, receiving a M.Sc. in Computer Science in 2001, specializing in Software Theory and Theoretical Computer Science. César is a recipient of the 2006 ACM Frank Anger Memorial Award. He keeps active collaborations with research groups in the USA and Europe.

Research Interests

César's research activities focus on formal methods for reactive systems with emphasis on the development and verification of concurrent, embedded and distributed systems. His foundational research includes the temporal verification of concurrent datatypes, runtime verification, and enhancements of linear temporal logics. In parallel, he is collaborating with industrial partners from the aerospace and embedded sectors to aid in the adoption of formal techniques for software development and validation. Current projects include the interactive formal generation of parallel software for satellite image processing, and the synthesis of advanced online debuggers for testing embedded software.

Pierre Ganty

Assistant Professor

Pierre joined the IMDEA Software Institute in the Fall 2009 after completing a nearly two year postdoc at the University of California, Los Angeles (UCLA). He holds a joint PhD degree in Computer Science from the University of Brussels (ULB), Belgium and from the University of Genova (Unige), Italy that he obtained late 2007. He is the principal investigator of the Spanish national project Paran'10 (2011-2013) on the verification of parameterized systems. He has supervised 6 internships at the IMDEA Software Institute. He served on the program committee of the international venues VMCAI (Verification, Model-Checking and Abstract Interpretation) in 2012 and RP (Reachability Problems) in 2013.

Research Interests

Pierre's research is focused on defining fully automated analysis techniques for systems with infinitely many states. Examples of such infinite state systems include sequential programs over unbounded data type, Internet of Things systems, communication protocols or event-based programs. In each of the previous example, there is an unbounded dimension: the data domain, the number of processes or the number of events; which is best modeled using an infinite state system.

Aleks Nanevski

Assistant Professor

Aleks received his Ph.D. degree in Computer Science from Carnegie Mellon University, USA in 2004. After holding postdoctoral positions at Harvard University (USA), and Microsoft Research, Cambridge (UK), Aleks joined the IMDEA Software Institute in September 2009. Prior to the PhD, Aleks finished his undergraduate studies in Computer Science at the University of Skopje, Macedonia in 1995.

Research Interests

Aleks' research is in the design and implementation of programming languages that facilitate verification of various program properties, ranging from type and memory safety, lack of memory leaks or information leaks, all the way to full functional correctness. His languages and systems unify programming and specification with automated and interactive theorem proving, via a common foundational framework of type theory. He is particularly interested in verifying programs that combine modern higher-order linguistic features such as higher-order functions, polymorphism, abstract types, objects and modules, with imperative ingredients such as pointer arithmetic, pointer aliasing, unstructured control flow, and concurrency.

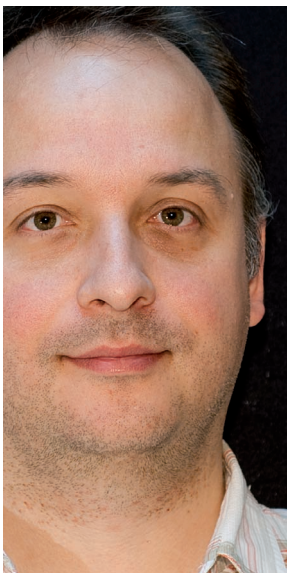

Alexey Gotsman

Assistant Professor

Alexey Gotsman received his Ph.D. degree in Computer Science from University of Cambridge, UK in 2009. During his Ph.D. studies, Alexey interned at Microsoft Research Cambridge, UK and Cadence Berkeley Labs, USA. He was a postdoctoral fellow at Cambridge before joining IMDEA in September 2010. Prior to his Ph.D., Alexey did his undergraduate and master studies in Applied Mathematics at Dnepropetrovsk National University, Ukraine, interning at the University of Trento, Italy in the process. He has received the prestigious Microsoft Research SEIF award to work on specifying and validating components on memory models of mobile platforms.

Research Interests

Alexey's research interests are in software verification, with particular focus on concurrent systems software. He is interested in developing both logics for reasoning about programs and automatic tools for verifying them.

Boris Köpf

Assistant Professor

Boris joined the IMDEA Software Institute in September 2011 after completing a post-doc at the Max Planck Institute for Software Systems (MPI-SWS). He received a Ph.D. degree from ETH Zurich in 2007, investigating formal methods for countering side-channel attacks. Before that, he studied mathematics at the Universidad de Chile, the Universidade Federal de Campinas, and the University of Konstanz, from which he received a M.Sc. degree. He is an alumnus of the German National Academic Foundation.

Research Interests

Boris' research focuses on the foundations of computer security. In particular, he is interested in quantitative notions of security, and in techniques for computing corresponding guarantees for real systems. He applies his research to the analysis of side-channel attacks (and countermeasures) and to privacy-preserving data publishing.

Juan Caballero

Assistant Professor

Juan Caballero joined the IMDEA Software Institute as an Assistant Research Professor in November 2011, after receiving his Ph.D degree in Electrical and Computer Engineering from Carnegie Mellon University, USA. Prior to joining the IMDEA Software Institute, Juan was a visiting graduate student researcher at University of California, Berkeley for two years. He was awarded the La Caixa fellowship for graduate studies in 2003. Juan also holds a M.Sc. in Electrical Engineering from the Royal Institute of Technology (KTH), Sweden, and a Telecommunications Engineer degree from the Technical University of Madrid (UPM), Spain.

Research Interests

Juan's research focuses on computer security, including security issues in systems, software, and networks. He enjoys designing program analysis techniques, specially techniques that work directly on program binaries. He applies those techniques for analyzing security properties of benign programs, as well as for malware analysis. In addition, he is interested in network security, the economic aspects of cybercrime, applying machine learning for security, and software engineering.

Pavithra Prabhakar

Assistant Professor

Pavithra Prabhakar obtained her doctorate in Computer Science from the University of Illinois at Urbana-Champaign in 2011, from where she also obtained a masters in Applied Mathematics. She has a masters degree in Computer Science from the Indian Institute of Science, Bangalore and a bachelors degree from the National Institute of Technology, Warangal, in India. She has been on the faculty of IMDEA Software since 2011, and spent the year between 2011-2012 at the California Institute of Technology as a CMI (Center for Mathematics of Information) fellow on leave of absence from IMDEA. She is the recipient of the Sohaib and Sara Abbasi fellowship from the University of Illinois and M.N.S Swamy medal from the Indian Institute of Science for the best master's thesis. Her paper at the ACM Hybrid Systems: Computation and Control Conference 2012 received a honorable mentions award.

Research Interests

Pavithra's main research interest is the formal analysis of cyber-physical systems. Her research is at the intersection of formal methods, hybrid systems and control theory with applications in robotics and aeronautics. Her research aims at building scalable analysis methods for systems consisting of mixed discrete continuous behaviors. To this end, she investigates foundational aspects such as decidability of verification problems and pre-orders for approximation; and develops algorithmic verification techniques and tools based on state-space reduction methods such as predicate abstraction and counter-example guided abstraction refinement. She has published widely in formal methods and hybrid systems conferences.

Pedro López-García

Researcher

Pedro López-García received a MS degree and a Ph.D. in Computer Science from the Technical University of Madrid (UPM), Spain in 1994 and 2000, respectively. In May 28, 2008 he obtained a Scientific Researcher position at the Spanish Council for Scientific Research (CSIC) and joined the IMDEA Software Institute. Immediately prior to this position, he held associate and assistant professor positions at UPM and was deputy director of the Artificial Intelligence unit at the Computer Science Department. He has published about 40 refereed scientific papers, many of them at conferences and journals of high or very high impact. He served as the scientific local coordinator of the european project ES_PASS "Embedded Software Product-based ASSurance," and is currently the principal investigator at the institute of the european FP7 FET project ENTRA "Whole-Systems Energy Transparency." He has also participated as a researcher in many other regional, national, and international projects.

Research Interests

His main areas of interest include energy-aware software engineering; automatic analysis and verification of non-functional program properties such as resource usage (user defined, energy, execution time, memory, etc.), non-failure and determinism; performance debugging; abstract interpretation; (automatic) granularity analysis/control for parallel and distributed computing; profiling; combined static/dynamic verification and unit-testing; type systems; tree automata; constraint and logic programming.

Mark Marron

Researcher

He joined the IMDEA Software Institute as a postdoctoral researcher in June 2008. Following four months as a Visiting Researcher at Microsoft Research in Redmond he returned to IMDEA as a Researcher. Recent research highlights include the release of a robust and scalable heap analysis toolkit (Jackalope analysis tools) for public use and the award of a prestigious Microsoft Innovation Award for work on heap analysis and memory use.

Research Interests

His research interests are on developing practical techniques for modeling program behavior and using this information to support error detection and optimization applications. His work to date has focused on the development of static analysis for the program heap which infers region, sharing, footprint and heap based data dependence information. More recent work has focused on using the information extracted by the analysis to support program parallelization, memory management, error detection, and software engineering applications.

Laurent Mauborgne

Researcher

Laurent Mauborgne received his Ph.D. in Computer Science from École Polytechnique, France, in 1999, and an Habilitation a diriger les recherches from University Paris-Dauphine (France) in 2007. He has been assistant professor at École normale supérieure, Paris, since 2000, and associate director of computer science studies there since 2006. He was also part-time professor at École Polytechnique. He was invited to spend a year at the IMDEA Software Institute in August 2009. He published 16 refereed papers in international conferences and 3 papers in journals. He gave courses in research summer schools and participated in the European projects DAEDALUS and ES_PASS. He was program committee member of the Static Analysis Symposium for 4 years. He is one of the authors of the Astrée analyzer, a tool that proved the absence of runtime errors in critical avionic codes.

Research Interests

The research of Laurent Mauborgne is focused on static analysis of programs and abstract interpretation. The goal is to develop theoretical as well as practical tools to analyze the behaviors of programs. This includes proving safety or temporal properties, optimizing compilation and computing resource usage. Among the recent subjects, he studied the cooperative combination of analyzers in different frameworks.

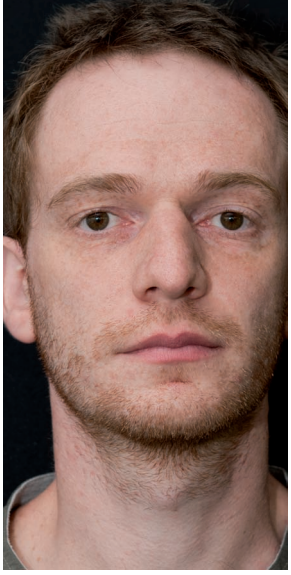

Pierre-Yves Strub

Researcher

Pierre-Yves Strub received his Ph.D. in Computer Science from École Polytechnique, France, in 2008. He joined the IMDEA Software Institute in 2013, after a post-doctoral position at the Microsoft-INRIA Joint Lab in Paris, France and at the LIAMA institute in Beijing, China.

Research Interests

Pierre-Yves research interests include formal proofs, proof assistants and their related type theory, certification of cryptographic algorithms and mathematical proofs, program verification via typing, and secure web programming. He is currently focused on EasyCrypt, a toolset for reasoning about relational properties of probabilistic computations with adversarial code, of which he is one of the main authors. He is also the main author of CoqMT, an extension of the Coq proof assistant.

postdoctoral researchers

César Kunz

Postdoctoral Researcher

César Kunz received a Computer Science degree from the National University of Córdoba (UNC), Argentina in 2004. He continued his studies at INRIA, France, funded by the FP6 FET integrated project "MOBIUS: Mobility, Ubiquity and Security", and received a Ph.D. from the École des Mines de Paris (ENSMP), France in February, 2009. He joined the IMDEA Software Institute as a postdoctoral researcher in February 2009.

Research Interests

His research interests lie around formal program analysis and verification, abstract interpretation, and program transformation. His primary research activities are centered on the certification of program correctness, the verification of compiler optimizations, and the transformation of verification results in the presence of program transformations.

Zoé Drey

Postdoctoral Researcher

Zoé Drey joined the IMDEA Software Institute as a postdoctoral researcher in October 2011. She received a Ph.D. degree in Computer Science from the University of Bordeaux 1, France, under the supervision of Charles Consel. Before joining IMDEA, she held a full-time teaching assistant position at the ENSEEIHT engineering school in Toulouse, France. Her Ph.D. studies were focused on making accessible the programming task in the field of networked entity orchestration, by providing adapted tools and methodologies to ease the development of reliable applications.

Research Interests

Her interests revolve around the design and implementation of domain-specific languages which reconcile language usability and reliability of the developed programs. In particular, she is interested in combining programming language semantics, logic and functional programming techniques, as well as software engineering methodologies to address this challenge. She is currently exploring the ways to make both the development process and the verification more usable to non-expert programmers, by adapting existing static analysis techniques to ease the instrumentation of domain-specific languages with debugging and verification tools (e.g., by automatically specializing existing debuggers/analyzers for user-friendly error reporting, and/or by providing high-level interfaces to existing specification languages)

Noam Zeilberger

Postdoctoral Researcher

Noam Zeilberger joined IMDEA Software in October 2011 as a postdoctoral researcher. Previously, he held a two-year postdoctoral fellowship of the *Fondation Sciences Mathématiques de Paris*, working at Université Paris 7. He obtained his Ph.D. in May 2009 from the Computer Science Department of Carnegie Mellon University, under the supervision of Peter Lee and Frank Pfenning.

Research Interests

Noam is interested broadly in the connections between logic and language and computation, and is excited by the potential of type theory (and its twin sister category theory) as a common foundation for (and a means of facilitating communication between) different areas of computer science. His work has focused on the Curry-Howard correspondence in general, and more specifically on: the problem of side-effects; continuations and computational duality; linear logic and focalisation; refinement types and dependent types. Since joining IMDEA and working with Gilles Barthe, he has also become interested in the notion of zero-knowledge from cryptography/complexity theory, and how it relates to notions of knowledge from proof theory.

Alexander Malkis

Postdoctoral researcher

Alexander has obtained his Diploma degree from the University of Saarland, Germany, in 2004-2005, for a work on polyforms (in other terminology, bond animals) under the guidance of Prof. Dr. Raimund Seidel; during his studies Alexander was financed by the prominent foundation "Studienstiftung des deutschen Volkes". He continued his studies in Saarbruecken and Freiburg, funded by the Max-Planck society and the DFG (German science foundation), obtaining his PhD thesis in 2011 at the University of Freiburg for a work on verification of multithreaded programs under guidance of Prof. Dr. Andreas Podelski. In April 2011, he joined the IMDEA Software Institute.

Research Interests

There is a range of topics in which Alexander is interested in, among them: polynomial verification of large program classes; emptiness of language intersection (complexity and algorithms); thread simulations, liveness, procedure abstractions under concurrency; a working verifier for multithreaded C; verifying multithreaded programs with rich structure and semantics, e.g. with heap, probabilism, recursion, for multicore systems; modeling biological and social systems; and synthesis of multithreaded embedded software.

José Francisco Morales

Postdoctoral researcher

Jose F. Morales joined the IMDEA Software Institute as a postdoctoral researcher in November 2011, after receiving his Ph.D degree in Computer Science from the Technical University of Madrid (UPM), Spain. Previously, he held a teaching assistant position at the Universidad Complutense de Madrid, starting in 2005.

Jose's work to date has focused on mechanisms for the efficient execution of logic programs: inference of program properties by abstract interpretation, highly-optimizing translation to low-level code using those properties, and the development of abstractions for the specification and automated construction of abstract machines.

Research Interests

His current research interests include the design of multiparadigm languages (combining imperative, logic, functional, and object-oriented programming), assertion languages and type systems, abstract interpretation, abstract machines, compiler optimizations, and native code generation.

Santiago Zanella Béguelin

Postdoctoral researcher

Santiago Zanella Béguelin obtained his degree in Computer Science from Universidad Nacional de Rosario (UNR), Argentina in 2006. He received his Ph.D. degree from École Nationale Supérieure des Mines de Paris in 2011 under the supervision of Gilles Barthe. From 2006 to 2011 he was a member of the *Secure Distributed Computations and their Proofs* team at the Microsoft Research-INRIA Joint Centre, Paris. He joined IMDEA in November 2009.

Research Interests

His main areas of interest include program specification and verification, quantitative analysis of programs, security proofs of cryptographic systems, language-based security, and proof assistants.

Santiago has devised novel program logics and programming language techniques that can be used to establish the security of cryptographic systems with an unprecedented level of assurance, making a jump from qualitative to quantitative guarantees, and from informal arguments to fully formalized, independently verifiable proofs. These ideas have been realized in the CertiCrypt framework, and applied to obtain certified security proofs of prominent and practically-relevant cryptographic systems, such as the Optimal Asymmetric Encryption Padding (OAEP) scheme.

He is currently working on developing automated tools to bring verification of security of cryptographic systems to practice, using off-the-shelf SMT solvers and automated theorem provers.

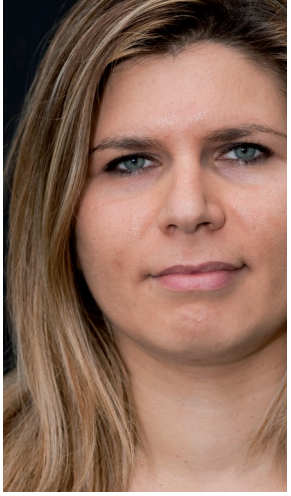

Zorana Banković

Postdoctoral researcher

Zorana Banković obtained her Electrical Engineer degree from the Faculty of Electrical Engineering at the University of Belgrade (Serbia) in 2005 and her Ph.D. degree from the Universidad Politécnica de Madrid (UPM) in 2011. Her dissertation was given the UPM special award as one of the four best theses of Telecommunications School that year. Before joining IMDEA Software, she was a researcher at the Department of Electronic Engineering at UPM. She has participated in 11 research and development projects, and authored 10 journal publications. During that time her main research interests included energy-efficient security solutions for wireless sensor networks, anomaly detection and thermal–aware optimizations in data centers, such as floorplanning, dynamic resource scheduling and allocation, as well as the design of a reputation system, that allows applying optimization techniques to each state of a data center.

After joining IMDEA Software in October 2012, her research has mainly been related to ENTRA research project, funded by the EU 7th Framework Program Future and Emerging Technologies (FET)

Research Interests

Her current research interests are in "energy-aware" software development using advanced program analysis and modeling of energy consumption in computer systems, aimed at making predictions of energy consumption early in the software design phase, and therefore enabling the development of greener IT through energy-efficient usage of hardware resources. Zorana's work includes research and development of energy optimization techniques at all software levels (compiler, OS, algorithms), as well as identification of static analyses that provide necessary input to the optimization stages which aim at improving resource consumption.

Ruy Ley Wild

Postdoctoral researcher

Ruy Ley-Wild joined the IMDEA Software Institute as a postdoctoral researcher in December 2011. He received his Ph.D. degree in Computer Science from Carnegie Mellon University under the supervision of Guy Blelloch. During his Ph.D. studies, he was funded by a Bell Labs Graduate Research Fellowship and interned at Bell Labs, Toyota Technological Institute at Chicago, and Microsoft Research Cambridge.

Research Interests

Ruy is broadly interested in the design and implementation of programming languages that express computation at a suitable level of abstraction and logics that enable high-level reasoning about the correctness and complexity of such programs. In particular, he has worked on compilation, cost semantics, and high-level dependence-tracking for self-adjusting computation. He is currently working with Aleks Nanevski on a type-theoretic approach to semantics and logics for a higher-order, stateful, concurrent language.

François Dupressoir

Postdoctoral researcher

François Dupressoir joined the IMDEA Software Institue as a postdoctoral researcher in February 2013. He successfully defended his Ph.D. in Computer Science at the Open University (U.K.) under the supervision of Andy Gordon, Jan Jürjens, and Bashar Nuseibeh. His Ph.D. studies were partially funded by a Microsoft Research Ph.D. scholarship, and led him to internships at the European Microsoft Innovation Center, and at Microsoft Research in Redmond and Cambridge.

Research Interests

François is broadly interested in program verification, theorem proving and cryptography. He is currently working with Gilles Barthe on methods for formally reasoning about cryptographic security properties of real-world systems, especially focusing on obtaining strong correctness and security results on low-level implementations of schemes and protocols, and studying how such properties can be preserved through compilation.

Benedikt Schmidt

Postdoctoral researcher

Benedikt Schmidt joined the IMDEA Software Institute as a postdoctoral researcher in February 2013. He received his Ph.D. degree in Computer Science from ETH Zurich, under the supervision of David Basin.

Research Interests

Benedikt is broadly interested in the areas of theorem proving, program verification, and rewriting and in their application to analyzing the security of systems. So far, his work has focused on the symbolic analysis of security protocols including interactive machine-checked approaches and fully automated approaches. Currently, he is also interested in the verification of cryptographic primitives and protocols against the computational model of attacks.

Ilya Sergey

Postdoctoral researcher

Ilya Sergey joined the IMDEA Software Institute as a postdoctoral researcher in December 2012. He received his Ph.D. degree in Computer Science from KU Leuven (Belgium) under the supervision of Dave Clarke in November 2012. During his doctoral studies he was a visiting Ph.D. fellow at the Department of Computer Science of Aarhus University, hosted by Olivier Danvy, and a research intern in the Programming Principles and Tools group at Microsoft Research Cambridge, supervised by Simon Peyton Jones.

Research Interests

Ilya's research interests dwell in the area of the design and implementation of programming languages, including but not limited to program semantics, certified programming, program transformations and refactoring techniques. He is particularly interested in developing methods of systematic derivation of correct-by-construction static analyses for higher-order languages by means of abstract interpretation, as well as their efficient implementations. Since joining IMDEA and working with Aleks Nanevski and Anindya Banerjee, he also became passionate about verification of multithreaded programs. He is currently working on a type-theoretic approach to specification and checking of properties of higher-order concurrent programs.




visiting faculty

Neil Jones
Visiting Professor

DIKU University of Copenhagen, Denmark
Visiting during Feb. 2012–Mar. 2012

Diego Garbervetsky
Visiting Professor

University of Buenos Aires, Argentina
Visiting during Jun. 2012–Jul. 2012

Peter Stuckey
Visiting Professor

University of Melbourne, Australia
Visiting during Jul. 2012–Dec. 2012

María García de la Banda
Visiting Professor

Monash University, Australia
Visiting during Jul. 2012–Dec. 2012

research assistants

ph.d. students

Research Assistants hold full scholarships or contracts at the Institute, performing research within one of the Institute's research lines under the supervision of a junior or senior researcher, and they undertake their Ph.D. at one of the Universities that the Institute has agreements with. Many of our Research Assistants obtain their degrees from the Technical University of Madrid (UPM), with whom the Institute collaborates in the development of graduate programs, and also at Universidad Complutense de Madrid (UCM).

Álvaro García

Research Assistant

Degree: Technical University of Madrid (UPM), Spain

Research: Efficient implementations of functional programming languages: theories and models for higher-order languages and lambda calculus, inter-derivation of program semantics, and abstract machines.

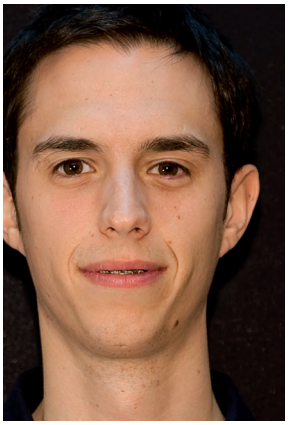

Miguel Angel García de Dios

Research Assistant

Degree: Universidad Complutense de Madrid (UCM), Spain

Research: Formal specification and verification, and rigorous tool supported modeling and validation of software systems.

Julian Samborski-Forlese

Research Assistant

Degree: Universidad Nacional de Rosario (UNR), Argentina

Research: Applications of formal methods and abstract interpretation to program verification; quantum computing; functional programming languages; semantics.




Juan Manuel Crespo
Research Assistant

Degree: Universidad Nacional de Rosario (UNR), Argentina

Research: Relational logics, formal verification of cryptographic primitives and protocols, programming languages.

Federico Olmedo
Research Assistant

Degree: Universidad Nacional de Rosario (UNR), Argentina

Research: Verification of cryptographic systems and semantics of programming languages.

Alejandro Sánchez
Research Assistant

Degree: Universidad Nacional de Córdoba (UNC), Argentina

Research: Concurrent systems, decision procedures, dynamic memory analysis, program verification.

Carolina Inés Dania
Research Assistant

Degree: Universidad Nacional de Córdoba (UNC), Argentina

Research: Tool-supported model-driven software development. Oriented on formal specification languages, security models, transformation and code generation.

Javier Valdazo Parnisari
Research Assistant

Degree: Universidad Nacional de Córdoba (UNC), Argentina

Research: Formal specification and verification. Rigorous tool supported modeling and validation of software systems. Model driven software engineering. Model transformations. Security models, transformation and enforcement.

Germán Andrés Delbianco
Research Assistant

Degree: Universidad Nacional de Rosario (UNR), Argentina

Research: Program verification with dependent types. In particular, developing new logics for reasoning about higher-order programs with imperative features, e.g., dynamic mutable state, continuations and concurrency, from a computational effects perspective.

Umer Liqat
Research Assistant

Degree: Dresden University of Technology (TUD), Germany

Research: Program analysis and verification, Automatic analysis and verification of (user definable) resource usage. In particular energy consumption analysis and optimization. Constraint and Logic programming.

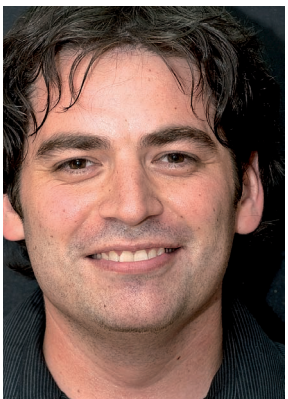

Gonzalo Ortiz
Research Assistant

Degree: Universidad Complutense de Madrid, Spain

Research: Languages and libraries to develop secure data-centric applications; model driven security; model-view-controller architecture; GUI development and libraries.

Antonio Nappa
Research Assistant

Degree: Università degli Studi di Milano, Italy

Research: Computer Security, Malware Analysis and Cybercrime

Pablo Chico de Guzmán
Research Assistant

Degree: Technical University of Madrid (UPM), Spain

Research: Advanced compilation techniques for Declarative Languages in order to improve their performance and semantics.

Artem Khyzha
Research Assistant

Degree: Dnipropetrovsk National University, Ukraine

Research: Developing compositional reasoning techniques for concurrent software. Application of separation logic to software verification.

Alejandro Serrano
Research Assistant

Degree: Autonomous University of Madrid (UAM), Spain

Research: Static resource analysis based on abstract interpretation, and its application to energy transparency and optimization in embedded systems.

Miriam García
Research Assistant

Degree: MSc in Mathematical Modelling in Engineering, University of L'Aquila and University of Hamburg

Research: Stability analysis based on model-checking techniques . Hybrid systems. Applied mathematics (PDEs, dynamical systems).

Goran Doychev
Research Assistant

Degree: Saarland University, Germany

Research: Obtaining quantitative security guarantees for computer systems and networks.

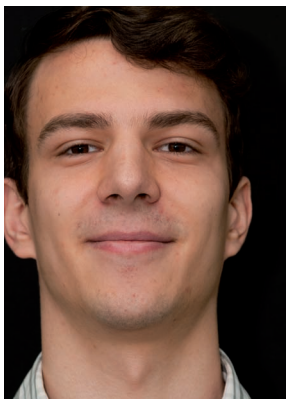

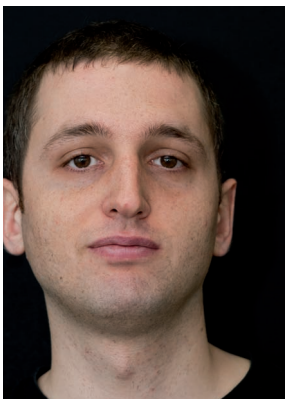

interns

Intern	Period	Nationality
Guido Genzone	Nov. 2011 - Apr. 2012	Argentina
Agustin Romano	Feb. 2012 - Jul. 2012	Italy
Shiva Shabaninejad	Feb. 2012 - Jul. 2013	Iran
Lucio Nardelli	Mar. 2012 - Aug. 2012	Argentina
Eugenia Simich	Apr. 2012 - Oct. 2012	Argentina
Ignacio Echeverria	Apr. 2012 - Sep. 2012	Argentina
M. Zubair Rafique	Apr. 2012 - Apr. 2013	Pakistan
Aditya Desai	May 2012 - Jul. 2012	India
Piotr Mardziel	May 2012 - Jul. 2012	Poland
Dominik Feld	May 2012 - Jul. 2012	Germany
Prasoon Dadich	Jun. 2012 - Aug. 2012	India
Arbob Ahmad	Jun. 2012 - Aug. 2012	USA
Alejandro Serrano	Jun. 2012 - Aug. 2012	Spain
Sachar Itzhaky	Jun. 2012 - Sep. 2012	Israel
Peerachai Kaowichakorn	Aug. 2012 - Jan. 2013	Thailand
Önder Babur	Nov. 2012 - Jul. 2013	Turkey
Edgardo Zoppi	Nov. 2012 - Mar. 2013	Argentina

project staff

Project Staff provide additional support for the development of projects and contracts being carried out at the Institute. They are typically co-funded by such projects.

Guillermo Jiménez
Technical Project Staff

Degree: Bach. Computer Science

research support
technical staff

Research Support Technical Staff provide support for the research infrastructures provided to the researchers at the Institute. These include virtualization environments, version control systems, research collaboration tools, continuous integration platforms, experimentation harnesses, computing farms, communications, etc. They are currently co-funded by different projects and agreements.

Roberto Lumbreras
Communication Infrastructures

Degree: MSc. Elec. & Computer Eng.

Juan Céspedes
Network and Systems Engineer (part-time)

Degree: MSc. Elec. & Computer Eng.

Inés Huertas
Network Technician (part-time)

Degree: Bach. Telematics

management & administration

María Alcaraz
General Manager

Degree: MBA, MSc. Economics

Marta Sedano
Technology Manager

Degree: BA (Hons) Business

Paola Huerta
Human Resources Assistant

Degree: MSc. History

Tania Rodríguez
Administrative Assistant (part-time)

Degree: MSc. Economics

Research Projects and Contracts

annual report 2012

Projects awarded by national or international agencies and contracts with industry are an important source of funding and opportunities for technology transfer for the Institute. During 2012 the Institute participated in 19 funded research projects and contracts, 13 of which (68% of the total) involve collaboration with industry. Of the 19 projects, 10 are international (8 funded by the European Union, 1 by the US ONR and Stanford University, and 1 by the Danish Research Council), 5 of them are direct industrial funding, and the rest are funded by national (3) and regional (1) agencies. Figure 5.1 shows the origin of project funding. The Institute has also obtained 16 fellowships.

The trend in external funding for the period 2008-2012, including the forecast for 2013, is shown in Figure 5.2. Forecast external funding for 2013 shows an increase of 45.17% over 2012. In 2012 33.69% of the external funding came from projects and contracts involving industry, and the forecast for 2013 implies a 118% increase. The percentage of external funding with respect to the total budget of the Institute was already around 30% in 2012, and also shows an increasing trend for 2013.

The rest of the chapter summarizes the main projects, contracts, grants, and fellowships awarded to the Institute in 2012, and briefly describes some interesting recently granted projects.

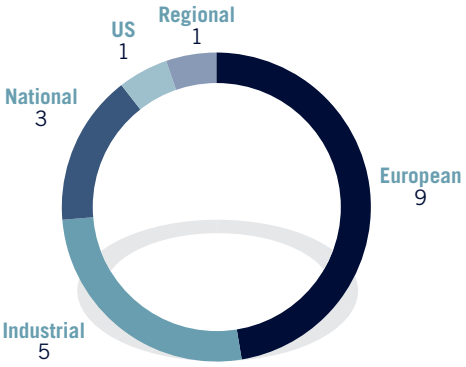

Figure 5.1. Projects by origin of funding.

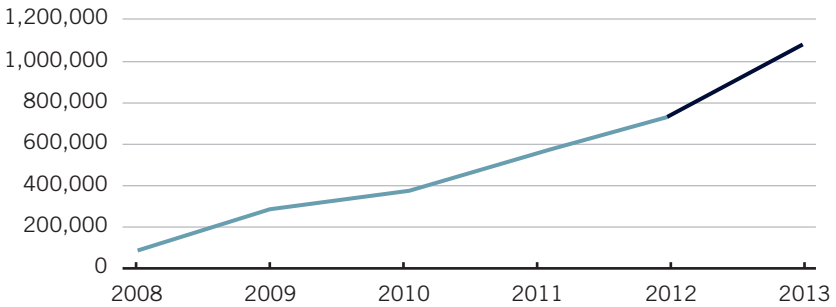

Figure 5.2. Evolution in external funding since 2008.

5.1. Ongoing Projects

ENTRA

Whole-systems energy transparency

Funding: European Union - 7th Framework Program - FET proactive MINECC call
Duration: 2012-2015
Project Coordinator: Prof. John Gallagher

ENTRA is an FP7 "Future and Emerging Technologies" project under the proactive "MINECC" objective - "Minimizing Energy Consumption of Computing to the Limit". The ENTRA project proposes radical advances in energy-aware software design and management with the objective of providing an important key to the pervasive realization of energy-aware computing. Though huge advances have been made in power-efficient hardware, most of the potential energy savings are wasted by software that does not exploit energy-saving features of hardware, and by poor dynamic management of tasks and resources. The budget of the project is approximately 2.7M Euros.

The project is built around the central concept of *energy transparency* at every stage of the software lifecycle. The project will develop novel *program analysis* and *energy modeling* techniques, making energy usage transparent through the system layers. This will enable *energy optimizations* both during code development and at run-time, and promote energy efficiency to a first-class software design objective.

AutoCrypt

Funding: US Office of Naval Research (ONR), through Stanford University
Duration: 2012-2015
Project Coordinator: Prof. Gilles Barthe

AutoCrypt is a joint project with Stanford University, University of Pennsylvania, and SRI, funded by ONR and which will run from July 2012 until July 2015. It has an overall budget of 2 Million Euros. AutoCrypt aims to use computer technology to provide mathematical guarantees that a cryptographic algorithm is secure, and that it is adequate for a given product, process, or service.

Within the project, the IMDEA Software team will use their EasyCrypt tool (http://www.easycrypt.info) to develop a systematic classification of cryptographic algorithms and to create a cryptographic atlas that will be used by researchers and companies to choose the most suitable algorithm for their needs.

HATS

Highly Adaptable and Trustworthy Software using Formal Models.

Funding: European Union – 7th Framework Program – FET proactive *Forever Yours* call
Duration: 2009-2013
Principal Investigator: Prof. Gilles Barthe

HATS is an Integrated Project funded by the European Union within the 7th Framework Program, "Future and Emerging Technologies" under the proactive "Forever Yours" objective. The main outcome envisaged by this project is an integrated architectural framework and a methodology for rigorous development of highly adaptable and trustworthy software. The IMDEA Software Institute is one of the research centers in a consortium of 8 academic partners, 2 industrial research centers, and 1 SML, from 7 countries. The budget for the project is approximately 6M Euros.

Specifically, HATS strives to turn software product family (SWPF) development into a rigorous approach. The technical core of the project is an Abstract Behavioral Specification language which allows precise description of SWPF features and components and their instances. The main project outcome is a methodological and tool framework achieving not merely far-reaching automation in maintaining dynamically evolving software, but an unprecedented level of trust while informal processes are replaced with rigorous analysis based on formal semantics.

The IMDEA Software Institute is responsible for the development of a highly adaptable architecture that allows cost-effective verification of the executable programs that will be automatically generated from Abstract Behavioral Specifications. The security architecture is specifically directed towards security policies expressed using information flow and functional correctness policies.

NESSoS

Network of Excellence on Engineering Secure Future Internet Software Services and Systems

Funding: European Union, Cooperation Program (NoE) – 7th Framework Program
Duration: 2011-2013
Principal Investigator: Prof. Manuel Clavel

The Network of Excellence on Engineering Secure Future Internet Software Services and Systems (NESSoS) aims at constituting and integrating a long lasting research community on engineering secure software-based services and systems. The NESSoS consortium involves 12 partners, including 2 companies (namely, Siemens and ATOS), from 7 countries. The budget for the project is approximately 3.5 M Euros.

The domain of Engineering Secure Software Services covers a collection of engineering activities that aim at the creation of software services —i.e. ICT services delivered through the deployment of software systems— that are both behaviorally correct (typically guided by software engineering principles) as well as secure (typically guided by security engineering principles). The approach of engineering secure software services is based on the principle of addressing security issues from the very beginning in system design and analysis, thus contributing to reducing system and service vulnerabilities, improving the necessary assurance level, thereby considering risk and cost issues during development in order to prioritize investments.

IMDEA Software plays a prominent role in three research workpackages: secure service architectures and design; programming environments for secure and composable services; and security assurance for services. Also, IMDEA Software leads the researcher mobility program within the consortium. This program is a mechanism that supports the integration of activities across the various sites: it brings together researchers working on related topics; it drives knowledge exchange and knowledge generation through union and diversity; and, finally, it increases the capability of joint cooperation among researchers.

VARIES

Variability in safety critical embedded systems

Funding: ARTEMIS- European Union - 7th Framework Program
Duration: 2012-2015
Principal Investigator: Laurent Mauborgne

VARIES is an ARTEMIS Joint Undertaking project granted under the FP7 ARTEMIS-2011-1 Call. The 26 partners-strong international consortium includes the participation of national partners Hi-Iberia, Integrasys, and Tecnalia. The main goal of the VARIES project is to help Embedded Systems (ES) developers to maximize the full potential of variability in safety-critical ES. The objectives of this project will be therefore (i) to enable companies to make informed decisions on variability use in safety-critical ES; (ii) to provide effective variability architectures and approaches for safety-critical ES; and (iii) to offer consistent, integrated and, continuous variability management over the entire product life cycle.

The VARIES project will deliver the VARIES Platform: a complete, cross-domain, multi-concern, state-of-the-art reference platform for managing variability in safety-critical ES. Special attention will be given to aspects specific to safety-critical ES, in particular the impact of reuse and composition on certification.

In addition to this ambitious goal, the VARIES project will create a Center of Innovation Excellence (CoIE) for managing variability in ES. The VARIES CoIE will support the European ES industry on the 3 aforementioned objectives.

DESAFIOS-10

High-Quality, Reliable, Distributed, and Secure Software Development

Funding: Spanish Ministry of Science and Innovation
Duration: 2011-2013
Principal Investigator: Prof. Gilles Barthe

The overall goal of the DESAFIOS-10 project is to contribute both foundations and technologies for the development of software systems with certified quality and reliability, based on formal methods and declarative programming. The consortium involves groups from three different Institutions (Universidad Complutense de Madrid, Universidad Politécnica de Madrid, and IMDEA Software) and a number of industrial users.

This project arises as a natural evolution of the previous coordinated project DESAFIOS, which involved only the research groups from Universidad Complutense de Madrid and

Universidad Politécnica de Madrid. In contrast, DESAFIOS-10 emphasizes the security and reliability aspects of this research, which is precisely the workpackage led by IMDEA Software.

PROMETIDOS

Methods for Rigorous Software Development

Funding: Regional Government of Madrid
Duration: 2011-2013
Principal Investigator: Prof. Gilles Barthe

The PROMETIDOS-CM research program is focused on four main areas: specification and validation, to provide a solid foundation for the description and analysis of services; reliability and security, to guarantee robust solutions from start to end; declarative programming, to develop the next generation of languages for services; and efficiency, to optimize quality of service with respect to performance. A common goal for all these research lines is the development of tools that will rigorously support their scientific results and that can be eventually transferred to industry.

PROMETIDOS-CM is a consortium involving groups from Universidad Complutense de Madrid, Universidad Politécnica de Madrid, and the IMDEA Software Institute, which is the project coordinator.

PARAN-10

Parametrized Verification of Computing Systems

Funding: Spanish Ministry of Science and Innovation
Duration: 2011-2013
Principal Investigator: Prof. Pierre Ganty

This project aims at developing novel techniques for production, verification and certification of computing systems where *parameters* play an essential role. Parameters either at the level of the system specification or at the level of the verification technique make it possible to address scalability and undecidability issues. However, specification and verification in the presence of parameters are highly non-trivial, and pose problems for automated verification methods (such as model checking) as well as interactive approaches to computing systems verification (such as theorem-proving), both of which are relevant in practice.

The project is organized along three research lines: model-checking of parametrized systems, parametric model-checking, and programming languages and logics for parametrization. In these three lines the project aims at making fundamental contributions to advance the state of the art as well as develop prototype implementations in order to explore and demonstrate the practical relevance of the proposed approaches.

RMT

Rich-Model Toolkit – An Infrastructure for Reliable Computer Systems (COST Action IC0901)

Funding: European Union, Cost action
Duration: 2009 - 2013
Principal Investigator: Prof. César Sánchez

This initiative explores directions and techniques for making automated reasoning (including analysis and synthesis) applicable to a wider range of problems, as well as making them easier to use by researchers, software developers, hardware designers, and information system users and developers. It includes participants from over 20 countries. A selection of the topics of interest is:

Standardization of expressive languages: Definitions of formats to represent systems, formulas, proofs, counterexamples. A framework to specify translations between specification languages, as well as benchmarks and competitions for automated reasoning, verification, analysis, and synthesis.

Decision procedures: Creation of decision procedures for new classes of constraints, including implementation of SAT and SMT and their certification. This requires the encoding of synthesis and analysis problems into SMT. The encoding of description logics (widely used in the Semantic Web) and the problem of scalable reasoning about knowledge bases are also addressed.

Transition system analysis: One key topic of study is abstraction-based approaches and refinement for verification of infinite-state systems. The application of constraint-based program analysis is also being analyzed, as well as data-flow analysis for complex domains. The application of TSA to programming languages and bytecode is being explored by extracting transition systems from them.

High-level synthesis: The project is devising new algorithms for synthesis from high-level specifications, and decision procedures are being extended to perform synthesis tasks. A relevant point being explored is the connection between invariant generation and code synthesis.

AMAROUT Europe

Funding: European Union, Marie Curie Action (PEOPLE-COFUND) - 7th Framework Program
Duration: 2009-2013
General Coordinator: Prof. Manuel Hermenegildo

AMAROUT Europe is a Marie Curie Action (PEOPLE-COFUND) to foster and consolidate the European Research Area by attracting to Europe and, in particular, to the region of Madrid, top research talent. AMAROUT helps the IMDEA network contribute to the goal of turning Madrid into one of the top knowledge generation regions in Europe. To accomplish this, the AMAROUT program finances up to 132 researchers to join the IMDEA network of research institutes for one year (renewable up to twice). The total budget for the program is around 11 M Euros of which the European Union cofinances 40%.

Both “experienced” and “very experienced” researchers from any country worldwide can apply for AMAROUT fellowships at any of the seven IMDEA Institutes participating in the program (Energy, Food, Materials, Nanoscience, Networks, Software, and Water). The AMAROUT Selection Committee consists of seven Evaluation Panels, one for each of the participating IMDEA Institutes. Each Evaluation Panel is formed by the Director of the Institute, three members of its Scientific Advisory Board, and two external, independent peer reviewers. The main AMAROUT selection criteria is the candidate’s demonstrated ability and commitment to research, as well as the match of experience and interests with the research theme and lines of the IMDEA Institute chosen by the candidate.

The AMAROUT Program is a joint initiative from the seven IMDEA research institutes. The IMDEA Software Institute was in charge of writing and submitting the proposal and is the beneficiary, acting as the administrator of the program for the other institutes. As such, it is in charge of the project management and its structure: Scientific Committee (SC); Fellowships Management Unit (FMU); Secretary; and Local Board of Prospective (BP).

NUSA

Numeric and Symbolic Abstractions for Software Model Checking

Funding: The Danish Council for Independent Research - Natural Sciences
Duration: 2011-2013
Principal Investigator: Prof. John Gallagher

Abstract interpretation and model checking are two approaches to verifying or deriving properties of software and hardware systems. While model checking is applied to finite-

state systems (typically hardware), abstract interpretation is usually aimed at infinite-state software systems. Indeed, the very notion of verification by abstraction starts from the assumption that the system under consideration is infinite or very large. Both abstract interpretation and model checking are the subject of major research efforts, both in academic and industrial laboratories, since they hold out the promise of an automatic, push-button approach to obtaining guarantees of system behavior. This proposal lies in the intersection of abstract interpretation and model checking. The main question for investigation in this project is how the framework and accumulated experience of abstract interpretation can be applied to model checking infinite state systems - in short, to define abstract model checking methods that exploit the generality and power of the framework of abstract interpretation.

AMAROUT II Europe

Funding: European Union, Marie Curie Action (PEOPLE-COFUND) - 7th Framework Program
Duration: 2012-2016
General Coordinator: Prof. Manuel Hermenegildo

AMAROUT-II Europe is a Marie Curie Action (PEOPLE-COFUND) which continues the AMAROUT action sharing with it the objectives of fostering and consolidating the European Research Area by attracting to Europe and, in particular, to the region of Madrid top research talent. As in the previous AMAROUT program “experienced” and “very experienced” researchers from any country (worldwide) can apply for AMAROUT II fellowships at any of the seven IMDEA Institutes participating in the program (Energy, Food, Materials, Nanoscience, Networks, Software, and Water). The program is seeking to attract, over 4 years, 152 experienced researchers to carry out research projects within the IMDEA network. The program keeps a call open permanently until months 36. Applications are evaluated by batches, according to quarterly cut-off dates. To promote the program and its calls, both nationally and abroad, best practices developed during the previous AMAROUT program are being applied. IMDEA Software is the mono-beneficiary of the AMAROUT-II programme, the same role that it is currently performing for the previous AMAROUT programme.

As in AMAROUT, the AMAROUT-II Program is a joint initiative from the seven IMDEA research institutes. The IMDEA Software Institute was in charge of writing and submitting the proposal and is the beneficiary, acting as the administrator of the program for the other institutes.

Microsoft Research Software Engineering Innovation Foundation Awards

The *SEIF (Software Engineering Innovation Foundation)* awards are given by Microsoft to support research in software engineering technologies, tools, practices, and teaching methods. These awards are given to project proposals which can be related to any of the core areas of interest in software engineering and have been given in 2012 for the third time. More than 100 project proposals were received this year, among which 10 projects were selected to receive the prize and the associated grant. Out of these 10 selected projects, two were granted to the IMDEA Software Institute for the following projects:

- **Alexey Gotsman** with the project *Specifying and Validating Components on Memory Models of Mobile Platforms*.
- **Mark Marron** with the project *MemAlyzer: Finding and Fixing Memory Usage Problems*.

The rest of the selected applications were from centers in the US (6), Canada (1 – U. of Calgary) and Switzerland (1 – ETH Zurich).

AbsInt GmbH

This project is a contract with AbsInt Angewandte Informatik GmbH to collaborate in the development of static analyzers by abstract interpretation. It is coordinated by Laurent Mauborgne. The goal of this contract is to develop advanced abstract interpretation techniques allowing fine tuning and increasing the precision and efficiency of the Astrée static analyzer sold and maintained by AbsInt. IMDEA Software brings its expertise and advice on sound abstractions of the memory model of the C language and on adaptive relational abstract domain tuning.

Telefónica Digital

In 2012, IMDEA has started a cooperation project with *Telefónica Digital*, Spain, aimed at research and development in components for automatic management of cloud scalability towards their integration into *Claudia*, a product developed within the European FI-WARE initiative. *Claudia* facilitates the definition and automatic deployment and management of virtual machines, storage, and connectivity resources that comprise the virtual infrastructure on which cloud applications are run.

The Institute is in charge of providing advice on the software architecture and high-level design of the software components, within the FI-WARE requirements, and partic-

ipates in their development and testing. The component integration is based on the OpenStack cloud architecture.

As mentioned before, Telefonica Digital and the Institute have also been working during 2012 towards the establishment of a *Joint Research Unit* (JRU) within their more global strategic partnership.

Boeing Research & Technology Europe

IMDEA has also been contracted in 2012 by *Boeing Research & Technology Europe* (which is located in Spain), to work jointly in research and development in the fields of Big Data and Social Network Analytics. In particular, the Institute and Boeing are jointly designing and implementing a framework for data mining in social media. The framework includes a declarative embedded language designed by IMDEA Software. This language supports the description of workflows that integrate map-reduce jobs and native applications. The implementation avoids costly recomputations increasing the efficiency of social media processing, with applications in rich Web interfaces that rely on live collection of social network information from Twitter streams and other sources.

5.2. Projects with Associated Groups

Part of the research of the Institute is performed in collaboration with research groups at associated institutions. This is exemplified by the existence of research projects led by these institutions but in which IMDEA personnel take part (and the resulting joint publications and results). We provide a summary list of the most relevant such projects which were active during 2012.

Project	Duration	Description	Funding Agency
S-CUBE	2008-2012	The European network of excellence in software and services	European Union – Network of Excellence
DOVES	2009-2014	Development of verifiable and efficient software	MINECO
SpaRCIM	2003-2014	Spanish Research Consortium for Informatics and Mathematics	European Union / MINECO

5.3. Recently Granted Projects (not started in 2012)

ADVENT

IMDEA Software is the main partner and coordinator of the ADVENT research project. The project was awarded during the year 2012 and will run from April 2013 to 2016. It is funded by the very competitive EU 7th Framework Programme, Future and Emerging Technologies (FET) *Young Explorers Initiative*, and has an overall budget of 1 million Euro. In addition to IMDEA Software, the consortium includes as partners Tel Aviv Univesity (Israel), The Max Planck Institute (Germany), and Katholieke Universiteit Leuven (Belgium).

The ADVENT project (http://advent-project.eu) will develop innovative methods and tools for cost-effective verification of real-world systems software, making it possible to guarantee an unprecedented level of reliability. ADVENT will achieve this by exploiting a trend among programmers to use informally described patterns, idioms, abstractions, and other forms of structure contained in their software, which are together called its architecture.

Building on the emerging technology of separation logic, ADVENT will formalize such software engineering concepts used by systems programmers to reason about their software informally, and will use the results to drive the design of verification techniques. This is a radically novel approach to the problem of verifying complex software: it departs from the common practice of building generic verification tools that, not being able to take advantage of programmers' knowledge and intuition, do not scale to big and complicated systems.

The architecture-driven verification techniques resulting from the project have the potential to yield a dramatic leap in the cost-benefit ratio of the verification technology. This will allow verification to scale to systems of real-world size and complexity that so far have been beyond the reach of quality assurance methods guaranteeing correctness.

The IMDEA Software team will be led by Alexey Gotsman.

4CaaSt

The IMDEA Software Institute has joined the EU 7th FP project *4CaaSt* as a partner in January 2013. The goal of the project is to create an advanced PaaS Cloud platform which supports the optimized and elastic hosting of Internet-scale multi-tier applications. 4CaaSt embeds all the necessary features, easing programming of rich applications and enabling the creation of a true business ecosystem where applications coming from different providers can be tailored to different users, mashed up and traded together.

IMDEA's participation focuses on the cloud application lifecycle engineering, management and experimentation, with a special focus on the process of resolution of application and component specifications (*blueprints*) to produce deployable configurations of cloud components that can be packaged into products that can be commercialized and marketed, thus promoting fast value uptake for small and medium enterprises using the cloud technologies. Manuel Carro is the PI of 4Caast at IMDEA Software.

5.4. Fellowships

1. *Microsoft Research PhD Scholarship funds*, awarded in 2011, active in 2012-2015 (**Alexey Gotsman**).
2. *Juan de la Cierva Postdoc grant*, Spanish Ministry of Science and Innovation, awarded in 2010 and ending in 2014 (**César Kunz**, through UPM).
3. *Juan de la Cierva Postdoc grant*, Spanish Ministry of Science and Innovation, awarded in 2011 and ending in 2015 (**Juan Caballero**).
4. *Ramón y Cajal grant*, Spanish Ministry of Science and Innovation, awarded in 2010 and ending in 2015 (**Aleksandar Nanevski**).
5. *Marie Curie AMAROUT Incoming Fellowships (3)*, European Union – Framework Program, awarded in 2009 and ending in 2012 (**Aleksandar Nanevski**, **Pierre Ganty**, and **Laurent Mauborgne**).
6. *Marie Curie AMAROUT Reintegration Fellowships* awarded in 2010 and ending in 2013 (**Juan Caballero**).
7. *Marie Curie AMAROUT Incoming Fellowships (4)* awarded in 2010 and active in 2012 (**Ruy Ley Wild**, **Boris Köpf**, **Alexey Gotsman** and **Alexander Malkis**).
8. *ERCIM / Marie-Curie Grants 2011*, European Union, Framework Program (**Zoé Drey**, September 2011 – August 2012).
9. *Predoctoral Grants*, Madrid Regional Government, awarded in 2009 and ending in 2013 (**Álvaro García**).
10. *FPI Doctoral Grant*, Spanish Ministry of Science and Innovation, awarded in 2010 and continuing until 2014 (**Juan Manuel Crespo**).
11. *FPU Doctoral Grant*, Spanish Ministry of Education, Culture and Sports, awarded in 2012 and ending in 2016 (**Julian Samborski**).

Dissemination of Results

6.1. Publications

6.1.1. Refereed Publications

1. Stan Rosenberg, *Anindya Banerjee*, David A. Naumann. *Decision Procedures for Region Logic*. VMCAI, pages 379–395, 2012.

2. *Gilles Barthe*, Jorge Cuéllar, Javier Lopez, Alexander Pretschner. *Preface*. Journal of Computer Security, Vol. 20, Num. 4, pages 307–308, 2012.

3. José Bacelar Almeida, Manuel Barbosa, Endre Bangerter, *Gilles Barthe*, Stephan Krenn, *Santiago Zanella Béguelin. Full Proof Cryptography: Verifiable Compilation of Efficient Zero-Knowledge Protocols*. ACM Conference on Computer and Communications Security, pages 488–500, 2012.

4. *Gilles Barthe*, David Pointcheval, *Santiago Zanella Béguelin. Verified Security of Redundancy-Free Encryption from Rabin and RSA*. ACM Conference on Computer and Communications Security, pages 724–735, 2012.

5. *Gilles Barthe*, Benjamin Grégoire, *César Kunz*, Yassine Lakhnech, *Santiago Zanella Béguelin. Automation in Computer-Aided Cryptography: Proofs, Attacks and Designs*. Certified Programs and Proofs - Second International Conference, CPP 2012, Lecture Notes in Computer Science, Vol. 7679, pages 7–8, Springer, 2012.

6. *Gilles Barthe*, Gustavo Betarte, Juan Diego Campo, Carlos Luna. *Cache-Leakage Resilient OS Isolation in an Idealized Model of Virtualization*. 25th IEEE Computer Security Foundations Symposium, CSF 2012, pages 186–197, IEEE, 2012.

7. Michael Backes, *Gilles Barthe*, Matthias Berg, Benjamin Grégoire, *César Kunz*, Malte Skoruppa, *Santiago Zanella Béguelin. Verified Security of Merkle-Damgrd*. 25th IEEE Computer Security Foundations Symposium, CSF 2012, pages 354–368, IEEE, 2012.

8. *Gilles Barthe*, *Juan Manuel Crespo*, Benjamin Grégoire, *César Kunz*, *Santiago Zanella Béguelin. Computer-Aided Cryptographic Proofs*. Interactive Theorem Proving - Third International Conference, ITP 2012, Lecture Notes in Computer Science, Vol. 7406, pages 11–27, Springer, 2012.

9. *Gilles Barthe*, *Juan Manuel Crespo*, Dominique Devriese, Frank Piessens, Exequiel Rivas. *Secure Multi-Execution through Static Program Transformation*. Formal Techniques for Distributed Systems - Joint 14th IFIP WG 6.1 International Conference, FMOODS 2012 and 32nd IFIP WG 6.1 International Conference, FORTE 2012, Lecture Notes in Computer Science, Vol. 7273, pages 186–202, Springer, 2012.

10. *Gilles Barthe*, Benjamin Grégoire, *Santiago Zanella Béguelin. Probabilistic Relational Hoare Logics for Computer-Aided Security Proofs*. Mathematics of Program Construction - 11th International Conference, MPC 2012, Lecture Notes in Computer Science, Vol. 7342, pages 1–6, Springer, 2012.

11. *Gilles Barthe*, Benjamin Grégoire, *Santiago Zanella Béguelin. Computer-Aided Cryptographic Proofs*. Static Analysis - 19th International Symposium, SAS 2012, Lecture Notes in Computer Science, Vol. 7460, pages 1–2, Springer, 2012.

12. *Gilles Barthe*, Delphine Demange, David Pichardie. *A Formally Verified SSA-Based Middle-End - Static Single Assignment Meets CompCert*. 21st European Symposium on Programming, ESOP 2012, Lecture Notes in Computer Science, Vol. 7211, pages 47–66, Springer, 2012.

13. *Gilles Barthe*, *Boris Köpf*, *Federico Olmedo*, *Santiago Zanella Béguelin. Probabilistic relational reasoning for differential privacy*. Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2012, pages 97–110, ACM, 2012.

14. *Gilles Barthe*, Benjamin Grégoire, Sylvain Heraud, *Federico Olmedo*, *Santiago Zanella Béguelin*. *Verified Indifferentiable Hashing into Elliptic Curves*. 1st Conference on Principles of Security and Trust, POST 2012, Lecture Notes in Computer Science, Vol. 7215, Springer, 2012.

15. *Juan Caballero*. *Understanding the Role of Malware in Cybercrime*. ERCIM News, Vol. 2012, Num. 90, pages 15–16, July 2012.

16. Chris Grier, Lucas Ballard, *Juan Caballero*, Neha Chachra, Christian J. Dietrich, Kirill Levchenko, Panayiotis Mavrommatis, Damon McCoy, *Antonio Nappa*, Andreas Pitsillidis, Niels Provos, M. Zubair Rafique, Moheeb Abu Rajab, Christian Rossow, Kurt Thomas, Vern Paxson, Stefan Savage, Geoffrey M. Voelker. *Manufacturing Compromise: The Emergence of Exploit-as-a-Service*. Proceedings of the 19th ACM Conference on Computer and Communication Security, October 2012.

17. *Juan Caballero*, Gustavo Grieco, *Mark Marron*, *Antonio Nappa*. *Early Detection of Dangling Pointers in Use-After-Free and Double-Free Vulnerabilities*. Proceedings of the 2012 International Symposium on Software Testing and Analysis, July 2012.

18. *M. V. Hermenegildo*, F. Bueno, *M. Carro*, *P. López*, E. Mera, *J.F. Morales*, G. Puebla. *An Overview of Ciao and its Design Philosophy*. Theory and Practice of Logic Programming, Vol. 12, Num. 1–2, pages 219–252, Cambridge University Press, January 2012. http://arxiv.org/abs/1102.5497.

19. *J. F. Morales*, R. Haemmerlé, *M. Carro*, *M. V. Hermenegildo*. *Lightweight compilation of (C)LP to JavaScript*. Theory and Practice of Logic Programming, 28th Int'l. Conference on Logic Programming (ICLP'12) Special Issue, Vol. 12, Num. 4-5, pages 755–773, Cambridge U. Press, 2012.

20. Dragan Ivanovic, *Manuel Carro*, *Manuel Hermenegildo*. *A Constraint-Based Approach to Quality Assurance in Service Choreographies*. 10th International Conference on Service Oriented Computing, ICSOC'12, LNCS, Num. 7637, Springer Verlag, November 2012.

21. George Bariannys, *Manuel Carro*, Dimitris Plexousakis. *Deriving Specifications for Composite Web Services*. IEEE Signature Conference on Computers, Software, and Applications, IEEE Computer Society, IEEE, July 2012.

22. Dragan Ivanovic, *Manuel Carro*, *Manuel Hermenegildo*. *Analyzing Service-Oriented Systems Using Their Data and Structure*. European Software Services and Systems Research – Results and Challenges (ICSE Workshop), July 2012.

23. Dimka Karastoyanova, Zsolt Nemeth, *Manuel Carro*, Dragan Ivanovic, Cesare Pautasso, Claudia Di Napoli, and Maurizio Giordano. *Research Challenges on Service Technology Foundations*. European Software Services and Systems Research – Results and Challenges (ICSE Workshop), July 2012.

24. D. Ivanovic, *M. Carro*, *M. Hermenegildo*. *Exploring the Impact of Inaccuracy and Imprecision of QoS Assumptions on Proactive Constraint-Based QoS Prediction for Service Orchestrations*. Proceedings of the 4th International Workshop on Principles of Engineering Service-Oriented Systems, PESOS 2012, pages 931–937, IEEE Press, June 2012.

25. D. Ivanovic, *M. Carro*, *M. Hermenegildo*. *Constraint-Based Runtime Prediction of SLA Violations in Service Orchestrations*. XII Jornadas sobre Programación y Lenguajes (PROLE), Universidad de Almería, 2012.

26. *P. Chico de Guzmán*, *M. Carro*, *M. Hermenegildo*, *P. Stuckey*. *A General Implementation Framework for Tabled CLP*. FLOPS'12, LNCS, Num. 7294, pages 104–119, Springer Verlag, May 2012.

27. *P. Chico de Guzmán*, A. Casas, *M. Carro*, *M. Hermenegildo*. *A Segment-Swapping Approach for Executing Trapped Computations*. PADL'12, LNCS, Vol. 7149, pages 138–152, Springer Verlag, January 2012.

28. *P. Chico de Guzmán*, *M. Carro*, *M. Hermenegildo*, *P. Stuckey*. *A General Implementation Framework for Tabled CLP*. XII Jornadas sobre Programación y Lenguajes (PROLE), Universidad de Almería, 2012.

29. *Carolina Dania*. *Modeling Social Networking Privacy*. ESSoS-DS 2012. Doctoral Symposium, Workshop Proceedings, Vol. 834, pages 49–54, CEUR, 2012.

30. *Z. Drey*, *J.F. Morales*, *M. V. Hermenegildo*. *Reversible Language Extensions and their Application in Debugging*. 12th International Colloquium on Implementation of Constraint and Logic Programming Systems (CICLOPS 2012), 15 pages, September 2012.

31. Javier Esparza, *Pierre Ganty*, Rupak Majumdar. *A Perfect Model for Bounded Verification*. LICS '12: Proc. 27th Annual ACM/IEEE Symp. on Logic in Computer Science, 2012.

32. Mai Ajspur, *John P. Gallagher*. *Towards Abstract Interpretation of Epistemic Logic*. 8th Scandinavian Logic Symposium, 20-21 August 2012, Roskilde University, Denmark, Roskilde University, 2012. Extended Abstract.

33. E.J. Gallego Arias, R. Haemmerlé, *M. V. Hermenegildo*, *J.F. Morales* . *The Ciao CLP(FD) Library: A Modular CLP Extension for Prolog*. 12th International Colloquium on Implementation of Constraint and LOgic Programming Systems (CICLOPS 2012), 15 pages, September 2012.

34. *Pierre Ganty*, Rupak Majumdar. *Algorithmic Verification of Asynchronous Programs*. ACM Transactions on Programming Languages and Systems, Vol. 34, Num. 1, pages 6:1–6:48, ACM, 2012.

35. *Pierre Ganty*, Rupak Majumdar, Benjamin Monmege. *Bounded Underapproximations*. Formal Methods in System Design, Vol. 40, Num. 2, pages 206–231, 2012.

36. *Alexey Gotsman*, Madanlal Musuvathi, Hongseok Yang. *Show no weakness: sequentially consistent specifications of TSO libraries*. Proceedings of the International Symposium on Distributed Computing (DISC'12), Salvador, Bahia, Brazil, LNCS, Vol. 7611, pages 31–45, Springer, 2012.

37. *Alexey Gotsman*, Hongseok Yang. *Linearizability with ownership transfer*. Proceedings of the 23rd International Conference on Concurrency Theory (CONCUR'12), Newcastle upon Tyne, UK, LNCS, Vol. 7454, pages 256–271, Springer, 2012. **Best paper award**.

38. Sebastian Burckhardt, *Alexey Gotsman*, Madanlal Musuvathi, Hongseok Yang. *Concurrent library correctness on the TSO memory model*. Proceedings of the 21st European Symposium on Programming (ESOP'12), Tallinn, Estonia, LNCS, Vol. 7211, pages 87–107, Springer, 2012.

39. E. Albert, P. Arenas, G. Puebla, *M. Hermenegildo*. *Certificate Size Reduction in Abstraction-Carrying Code*. Theory and Practice of Logic Programming, Vol. 12, Num. 3, pages 283–318, 2012.

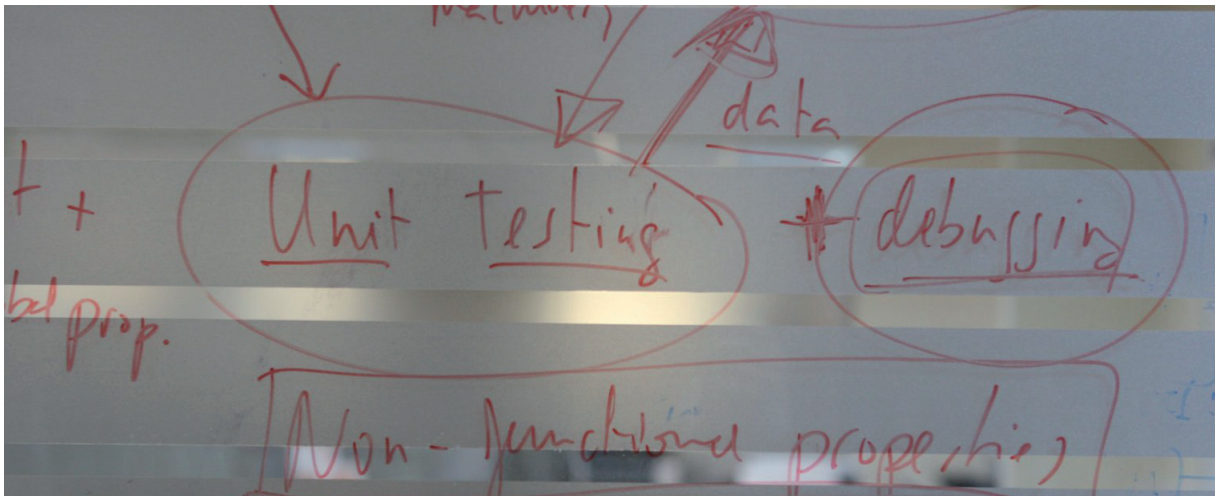

40. *Boris Köpf*, *Laurent Mauborgne*, Martín Ochoa. *Automatic Quantification of Cache Side-Channels*. 24th International Conference on Computer Aided Verification (CAV '12), Lecture Notes in Computer Science, Vol. 7358, pages 564–580, Springer-Verlag, 2012.

41. *Ruy Ley-Wild*, Umut A. Acar, Guy Blelloch. *Non-Monotonic Self-Adjusting Computation*. European Symposium on Programming, ESOP 2012, Lecture Notes in Computer Science, Vol. 7211, pages 476–496, Springer, 2012.

42. *Alexander Malkis*, *Anindya Banerjee*. *Verification of software barriers*. 17th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, pages 313–314, 2012.

43. *Mark Marron*, Ondrej Lhoták, *Anindya Banerjee*. *Programming Paradigm Driven Heap Analysis*. 21st International Conference on Compiler Construction (CC '12), Lecture Notes in Computer Science, Vol. 7210, pages 41–60, 2012.

44. *Mark Marron*, *César Sánchez*, Zhendong Su, Manuel Fähndrich. *Abstracting Runtime Heaps for Program Understanding*. IEEE Transactions on Software Engineering, Vol. 99, IEEE Computer Society, 2012.

45. Patrick Cousot, Radhia Cousot, *Laurent Mauborgne*. *Theories, Solvers and Static Analysis by Abstract Interpretation*. Journal of the ACM, Vol. 59, Issue 3, ACM, December 2012.

46. *J. F. Morales*, *M. V. Hermenegildo*, R. Haemmerlé. *Modular Extensions for Modular (Logic) Languages*. Proceeding of the 21th International Symposium on Logic-Based Program Synthesis and Transformation (LOPSTR'11), LNCS, Vol. 7225, pages 139–154, Springer, July 2012.

47. *J. F. Morales*, R. Haemmerlé, *M. Carro*, *M. V. Hermenegildo*. *Lightweight compilation of (C)LP to JavaScript*. XII Jornadas sobre Programación y Lenguajes (PROLE), Universidad de Almería, 2012.

48. *Alejandro Sánchez*, Sriram Sankaranarayanan, *César Sánchez*, Bor-Yuh Evan Chang. *Invariant Generation for Parametrized Systems Using Self-reflection*. Proc. of the 19th International Symposium on Static Analysis (SAS'12), LNCS, Vol. 7460, pages 146–163, Springer, 2012.

49. Marina Zapater, *César Sánchez*, José L. Ayala, José M. Moya, José L. Risco-Martín. *Ubiquitous Green Computing Techniques for High Demand Applications in Smart Environments*. Sensors, Vol. 12, Num. 8, pages 10659–10677, August 2012.

50. *César Sánchez*, *Julián Samborski-Forlese*. *Efficient Regular Linear Temporal Logic Using Dualization and Stratification*. Proceedings of the 19th International Symposium on Temporal Representation and Reasoning (TIME'12), pages 13–20, IEEE Computer Society, 2012.

51. *César Sánchez*, *Julián Samborski-Forlese*. *How to Translate Efficiently Extensions of Temporal Logics into Alternating Automata*. Proceedings of The 9th International Colloquium on Theoretical Aspects of Computing (ICTAC'12), LNCS, Vol. 7521, pages 30–45, Springer, 2012.

52. Vladimeros Vladimerou, *Pavithra Prabhakar*, Mahesh Viswanathan, Geir E. Dullerud. *Verification of Bounded Discrete Horizon Hybrid Automata*. IEEE Transactions on Automatic Control, Vol. 57, Num. 6, pages 1445–1455, 2012.

53. *Pavithra Prabhakar*. *Foundations for Approximation Based Analysis of Stability Properties of Hybrid Systems*. 50th Annual Allerton Conference on Communication, Control and Computing, 2012.

54. *Pavithra Prabhakar*, Mahesh Viswanathan. *Conformance Testing of Boolean Programs with Multiple Faults*. Formal Techniques for Distributed Systems - Joint 14th IFIP WG 6.1 International Conference, FMOODS 2012 and 32nd IFIP WG 6.1 International Conference, FORTE 2012, Stockholm, Sweden, June 13-16, 2012. Proceedings, Lecture Notes in Computer Science, Vol. 7273, pages 101–117, Springer, 2012.

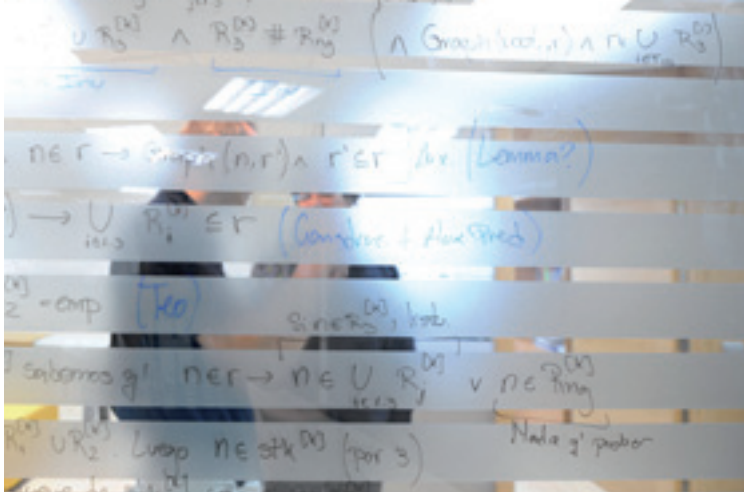

55. *Pavithra Prabhakar*, Geir E. Dullerud, Mahesh Viswanathan. *Pre-orders for reasoning about stability*. Hybrid Systems: Computation and Control (part of CPS Week 2012), HSCC'12, Beijing, China, April 17-19, 2012, pages 197–206, ACM, 2012. **Honorable mention award.**

56. *Martin Wirsing*, Jonas Eckhardt, Tobias Mühlbauer, Jose Meseguer. *Design and Analysis of Cloud-Based Architectures with KLAIM and Maude*. Rewriting Logic and Its Applications. Lecture Notes in Computer Science, Vol. 7571, pages 54–82, Springer, 2012.

57. Jonas Eckhardt, Tobias Mühlbauer, Musab AlTurki, José Meseguer, *Martin Wirsing*. *Stable Availability under Denial of Service Attacks through Formal Patterns*. Fundamental Approaches to Software Engineering. Lecture Notes in Computer Science, Vol. 7212, pages 78–93, Springer 2012.

6.1.2. Edited Volumes

1. *Gilles Barthe*, Benjamin Livshits, Riccardo Scandariato (Eds.): *Engineering Secure Software and Systems - 4th International Symposium*, ESSoS 2012, Eindhoven, The Netherlands, February, 16-17, 2012. Proceedings. Lecture Notes in Computer Science 7159, Springer 2012.

2. *Gilles Barthe*, Anupam Datta, Sandro Etalle (Eds.): Formal Aspects of Security and Trust - 8th International Workshop, FAST 2011, Leuven, Belgium, September 12-14, 2011. Revised Selected Papers. Lecture Notes in Computer Science 7140, Springer 2012.

6.1.3. Doctoral and Master Theses

1. Pablo Chico de Guzmán. *Advanced Evaluation Strategies for Tabling and Parallelism in Logic Programs.* Ph.D. Thesis, Universidad Politécnica de Madrid (UPM), November 2012. Advisers: Manuel Hermenegildo and Manuel Carro (IMDEA Software Institute).

2. Shiva Shabaninejad. *Valuation: Developer Support For By-Reference To By-Value Type Conversion.* M.Sc. Thesis, Technical University of Madrid (UPM). July 2012. Advisers: Manuel Carro and Mark Marron (IMDEA Software Institute).

3. Peerechai Kaowichakorn. *Probabilistic Analysis of Quality of Service.* M.Sc. Thesis, Universidad Politécnica de Madrid (UPM). October 2012. Advisors: Manuel Carro (IMDEA Software Institute) and Stefanie Betz (Blekinge Institute of Technology, Karlsruhe, Germany).

4. Goran Doychev. *Analysis and Mitigation of Information Leaks in Web Browsing Traffic.* M.Sc. Thesis, Saarland University, June 2012. Advisors: Boris Köpf (IMDEA Software Institute) and Michael Backes (Saarland University and MPI SWS, Germany).

5. Javier Valdazo. *Developing Secure Business Applications From Secure BPMN Models.* M.Sc. Thesis, Universidad Complutense de Madrid, 2012. Advisor: Manuel Clavel (IMDEA Software Institute and UCM).

6. Gustavo Grieco. *Inferring the Type Graph of a Binary Program*, University of Rosario, Argentina, August 2012. Advisor: Juan Caballero (IMDEA Software Institute).

6.2. Invited Talks

6.2.1. Invited and Plenary Talks by IMDEA Scientists

Anindya Banerjee:

1. Experiments in Mechanized Verification. *City University of New York*, NY, USA, May, 2012.

2. Modular Reasoning about Object-based Programs. *Laboratory for Foundations of Computer Science* (LFCS) Seminar, University of Edinburgh, UK, September 2012.

3. Flexible Confidentiality Policies and their Modular Enforcement. *IBM T. J. Watson Research Center*, USA, November 2012.

Gilles Barthe:

4. Computer-Aided Cryptographic Proofs and Designs. *17th European Symposium on Research in Computer Security* (ESORICS 2012), Consiglo Nazionale delle Ricerche, Pisa, Italy, September 2012.

5. Computer-Aided Cryptographic Proofs. *19th International Static Analysis Symposium* (SAS 2012), École Normale Supérieure, Deauville, France, September 2012.

6. Computer-Aided Cryptographic Proofs. *Interactive Theorem Proving* (ITP 2012), Princeton University, NJ, USA, August 2012.

7. Automation in Computer-Aided Cryptography: Proofs, Attacks and Designs. *Second International Conference on Certified Programs and Proofs* (CPP 2012), Kyoto, Japan, December 2012.

8. Computer-Aided Cryptographic Proofs. *ACM SIGPLAN 7th Workshop on Programming Languages and Analysis for Security* (PLAS 2012), Beijing, China, June 2012.

9. Probabilistic Relational Hoare Logics For Computer-Aided Security Proofs. *Mathematics of Program Construction* (MPC 2012), Madrid, Spain, June 2012.

10. Computer-Aided Cryptographic Proofs and Designs. *10th School For Young Researchers About Modelling And Verifying Parallel Processes* (MoVeP 2012), Marseille, December 2012.

Manuel Carro:

11. On Constraint Tabled Programming. *Universidade do Porto*, Portugal, June, 2012.

12. Constraint-Based Runtime Prediction of SLA Violations in Service Orchestrations. *22nd Workshop on Logic-based methods in Programming Environments* (WLPE 2012), Budapest, Hungary, September 2012.

Boris Köpf:

13. Quantitative Information Flow – Fundamental Techniques and Applications to Side Channel Analysis. *International Workshop on Quantitative Aspects in Security Assurance* (QASA 2012), Pisa, Italy, September 2012.

14. Quantifying Side Channels in RSA and AES. *10th Workshop on Quantitative Aspects of Programming Languages* (QAPL 2012), Talinn, Estonia, April 2012.

Juan José Moreno:

15. Oportunidades laborales en diferentes sectores de aplicación de las TICS: El sector académico y de investigación. *eSkill Week*, Madrid, June 2012.

16. Estratégias para la internacionalización de la investigación. *Universidad de Santiago de Compostela*, Spain, May 2012.

17. Higher Education Environment, MDP for International Centers of Excellence, EOI, January 2012.

Aleksandar Nanevski:

18. How to Make Ad Hoc Proof Automation Less Ad Hoc. *ICT Innovations 2012 - Secure and Intelligent Systems* (ICTI 2012), Ohrid, Macedonia, September 2012.

Pavithra Prabhakar:

19. Foundations for Approximation based Analysis of Stability Properties of Hybrid Systems. Session on Verification of Cyberphysical Systems: Tools and Algorithms, *50th Annual Allerton Conference*, Allerton House, Monticello, IL, USA, October 2012.

20. Pre-orders for reasoning about Stability. *Department of Computer Science and Automation, Indian Institute of Science*, Bangalore, India, April 2012.

21. Approximation-based Verification of Hybrid Systems. *MIT Electrical Engineering and Computer Science*, Boston, MT, USA, February 2012.

22. Approximation-based Verification of Hybrid Systems. *NASA Jet Propulsion Lab*, Pasadena, CA, USA, March 2012.

6.2.2. Invited Seminars and Lectures by IMDEA Scientists

1. *Anindya Banerjee*. Flexible Confidentiality Policies and their Modular Enforcement. *Florida International University Distinguished Lecture,* Miami, FL, USA, December, 2012.

2. *Pierre Ganty*. Precise and Automated Analysis for Systems with Infinitely Many States. *Microsoft Research*, Redmond, USA, December 2012.

3. *Pierre Ganty*. A Perfect Model for Bounded Verification. *Verimag*, Grenoble, France, May 2012.

4. *Alexey Gotsman*. Library abstraction for C/C++ concurrency. *Microsoft Research*, Redmond, WA, USA, August 2012.

5. *Alexey Gotsman*. Compositional reasoning on weak memory models. *Microsoft Research*, Cambridge, UK, March 2012.

6. *Alexey Gotsman*. Library abstraction for C/C++ concurrency. *HP Labs*, Palo Alto, CA, USA, September 2012.

7. *Alexey Gotsman*. Concurrent library correctness on weak memory models. *University of Newcastle*, UK, April 2012.

8. *Alexey Gotsman*. Semantics of eventual consistency *University of Nantes*, France, November 2012.

9. *Boris Köpf*. Quantitative Analysis of Side Channels in RSA and AES. *Cornell University*, USA, January 2012.

10. *Boris Köpf*. Quantitative Analysis of Side Channels in RSA and AES. *Queen Mary University of London*, UK, April 2012.

11. *Boris Köpf*. Quantitative Analysis of Side Channels in RSA and AES. *LMU Munich*, Germany, July 2012.

12. *Boris Köpf*. Automatic Quantification of Cache Side Channels. *Dagstuhl Seminar 12481*, Germany, November 2012.

13. *Pavithra Prabhakar*. Approximation-based Verification of Hybrid Systems. *Center for Control, Dynamical Systems, and Computation* (CCDC) Seminar, University of California, Santa Barbara, CA, USA, June 2012.

14. *Pavithra Prabhakar*. Approximation-based Verification of Hybrid Systems. *PRECISE Seminar Series*, University of Pennsylvania, Philadelphia, PA, USA, April 2012.

15. *Pavithra Prabhakar*. Approximation-based Verification of Hybrid Systems. *California Institute of Technology*, Pasedena, CA, USA, September 2012.

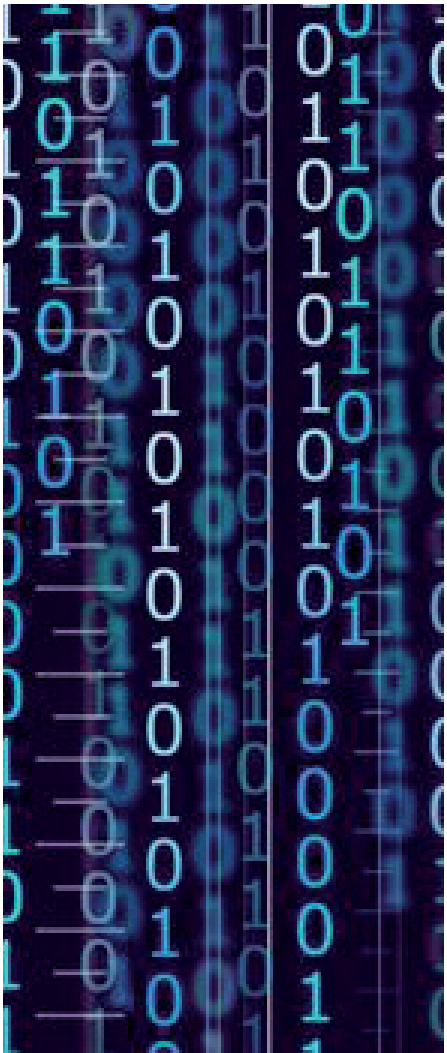

6.2.3. Invited Speaker Series

During 2012, 17 external researchers were invited to give altogether 24 talks at the IMDEA Software Institute. The following is the list of researchers and their talks:

1. *Neil Jones*, Professor, DIKU University of Copenhagen, Denmark: Introduction to partial evaluation.

2. *Neil Jones*, Professor, DIKU University of Copenhagen, Denmark: Obfuscation by Partial Evaluation of Distorted Interpreters.

3. *Neil Jones*, Professor, DIKU University of Copenhagen, Denmark: PROGRAMS = DATA = FIRST-CLASS CITIZENS IN A COMPUTATIONAL WORLD.

4. *Neil Jones*, Professor, DIKU University of Copenhagen, Denmark: Superlinear speedup (Work in Progress).

5. *Neil Jones*, Professor, DIKU University of Copenhagen, Denmark: Kleene's Second Recursion Theorem (Work in Progress).

6. *Neil Jones*, Professor, DIKU University of Copenhagen, Denmark: Some remarks on the spectrum problem.

7. *Markus Rabe*, PhD Student, Saarland University, Germany: Temporal Information Flow.

8. *Schmuel Mooly Sagiv*, Professor, Tel Aviv University, Israel. COLT: Testing and Verifying Atomicity of Composed Operations.

9. *Yuriy Brun*, Post-doctoral Researcher, University of Washington, USA: Speculative Analysis - What's Wrong with the Program I Haven't Written Yet?

10. *Earl Barr*, Post-doctoral Researcher, University of California at Davis, USA: Techniques and Tools for Engineering Robust Numerical Software.

11. *Earl Barr*, Post-doctoral Researcher, University of California at Davis, USA: On the Naturalness of Software (to appear at ICSE 2012).

12. *Mike Dodds*, Post-doctoral Researcher, University of Cambridge, UK: Recovering Disjointness from Concurrent Sharing.

13. *Mike Hicks*, Associate Research Professor, University of Maryland, USA: Polymonads - reasoning and inference.

14. *Pietro Ferrara*, Post-doctoral Researcher, ETH Zurich, Switzerland: TVLA and Value Analyses Together.

15. *Geoffrey Smith*, Associate Research Professor, Florida International University, USA: Measuring Information Leakage using Generalized Gain Functions.

16. *Geoffrey Smith*, Associate Research Professor, Florida International University, USA: Some Recent Results on Min-Entropy Leakage.

17. *Diego Garbervetsky*, Professor, University of Buenos Aires, Argentina: Quantitative Analysis of Java/.Net like programs to understand heap memory requirements.

18. *Jérôme Feret*, Researcher, École Normale Supérieure, France: Formal model reduction.

19. *Kathryn Francis*, PhD Student, The University of Melbourne, Australia: Optimization modelling for software developers.

20. *Jan Reineke*, Assistant Professor, Saarland University, Germany: Timing Analysis of Embedded Software for Platforms.

21. *Roberto Di Cosmo*, Professor, Université Paris Diderot, Director IRILL: Analysing co-installability of software components.

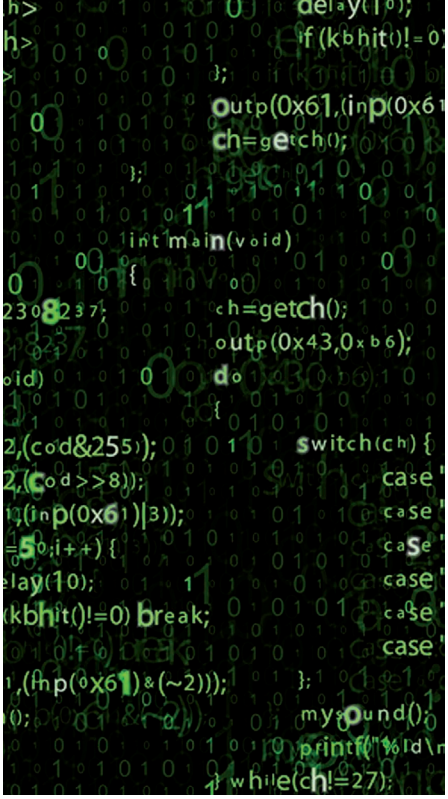

22. *Judith Bishop*, Researcher, Microsoft Research: Microsoft Research connecting to Labs and Academia.

23. *Derek Dreyer*, Assistant Research Professor (tenure-track), Max Planck Institute for Software Systems, Germany: Logical Relations for Fine-Grained Concurrency

24. *Vítor Santos Costa*, Associate Research Professor, University of Porto, Portugal: Learning with Prolog.

6.2.4. Software Seminar Series

The Institute also holds an internal seminar series to foster communication and collaboration. A total of **20** seminars were given in 2012.

6.3. Scientific Service & Other Activities

6.3.1. Participation in Program Committees

Anindya Banerjee:

1. *6th Indian Software Engineering Conference* (ISEC-2013).

2. *Workshop on Logics for Systems Analysis* (LfSA, satellite of CAV 2012).

3. *Workshop on Interference and Dependence* (satellite of POPL 2013).

Gilles Barthe:

4. *First Conference on Principles of Security and Trust*, (POST 2012, part of ETAPS 2012).

5. *37th International Symposium on Mathematical Foundations of Computer Science* (MFCS 2012).

6. Tool Demonstrations, *27th IEEE/ACM International Conference on Automated Software Engineering* (ASE-Tools 2012).

7. *50th International Conference on Objects, Models, Components, Patterns* (TOOLS Europe 2012).

8. *18th Conference on Logic Programming and Automated Reasoning* (LPAR 2012).

9. *8th International Workshop on Security and Trust Management* (STM 2012).

Juan Caballero:

10. *33rd IEEE Symposium on Security and Privacy* (S&P 2012).

11. *15th International Symposium on Research in Attacks, Intrusions and Defenses* (RAID 2012).

12. *XII Spanish Meeting on Cryptography and Information Security* (RECSI 2012).

13. *12th Annual Digital Forensics Research Conference* (DFRWS 2012).

14. *9th Conference on Detection of Intrusions and Malware and Vulnerability Assesment* (DIMVA 2012).

15. *10th Annual Conference on Privacy, Security and Trust* (PST 2012).

Manuel Carro:

16. *28th International Conference on Logic Programming* (ICLP 2012).

17. *Seventh International Workshop on Declarative Aspects of Multicore Programming* (DAMP 2012).

18. *XII Jornadas sobre Programación y Lenguajes* (PROLE 2012).

19. *10th International Conference on Service Oriented Computing* (ICSOC 2012).

20. *The 10th IEEE International Symposium on Parallel and Distributed Processing with Applications* (ISPA 2012).

21. *1st International Workshop On Monitoring And Prediction Of Cloud Services* (MoPoC 2012, satellite of ICSOC 2012).

Manuel Clavel:

22. *9th International Workshop on Rewriting Logic and its Application* (WRLA 2012).

23. *2012 Workshop on OCL and Textual Modelling* (OCL 2012).

24. *Model-Driven Security Workshop* (MDSec 2012, in conjunction with MoDELS 2012).

John Gallagher:

25. *22nd International Symposium on Logic-Based Program Synthesis and Transformation* (LOPSTR 2012).

26. *Partial Evaluation and Program Manipulation* (PEPM 2012).

27. *50th International Conference on Objects, Models, Components, Patterns* (TOOLS Europe 2012).

28. *The Fourth International Workshop on Numerical and Symbolic Abstract Domains* (NSAD 2012).

Pierre Ganty

29. *14th International Workshop on Verification of Infinite-State Systems* (INFINITY 2012).

Alexey Gotsman:

30. *28th Conference on the Mathematical Foundations of Programming Semantics* (MFPS 2012).

31. *40th International Colloquium on Automata, Languages, and Programming* (ICALP 2013).

Manuel Hermenegildo:

32. *The Alan Turing Centenary Conference* (Turing 100 2012).

33. *18th Conference on Logic Programming and Automated Reasoning* (LPAR 2012).

34. *International Workshop on Scripts to Programs* (STOP 2012, co-located with ECOOP 2012).

Boris Köpf:

35. *8th International Conference on Information Security Practice and Experience* (ISPEC 2012).

36. *7th International Symposium on Trustworthy Global Computing* (TGC 2012).

37. *IEEE Workshop on Privacy and Anonymity for the Digital Economy* (PADE 2012).

38. *Workshop on Quantitative Aspects in Security Assurance* (QASA 2012).

Mark Marron:

39. *Workshop on Developing Tools as Plug-ins* (TOPI 2012, co-located with ICSE 2012).

40. *2013 ACM SIGPLAN International Symposium on Memory Management* (ISMM 2013)

41. *8th Workshop On Bytecode Semantics, Verification, Analysis And Transformation* (BYTECODE 2013).

Laurent Mauborgne:

42. *40th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (POPL 2013).

43. *5th Workshop on Numerical and Symbolic Abstract Domains* (NSAD 2013).

José Francisco Morales:

44. *28th International Conference on Logic Programming* (ICLP 2012).

Juan José Moreno:

45. *XII Jornadas sobre Programación y Lenguajes* (PROLE 2012).

46. *XVII Jornadas de Ingeniería del Software y Bases de Datos* (JISBD 2012).

47. *VIII Jornadas de Ciencia e Ingeniería de Servicios* (JCIS 2012).

Aleksandar Nanevski:

48. *6th ACM SIGPLAN Workshop Programming Languages meets Program Verification* (PLVP 2012).

49. *6th International Joint Conference on Automated Reasoning* (IJCAR 2012).

50. Track B, *39th International Colloquium on Automata, Languages and Programming* (ICALP 2012).

51. *1st ACM SIGPLAN Workshop on Higher-Order Programming with Effects* (HOPE 2012).

52. *22nd European Symposium on Programming* (ESOP 2012).

Pavithra Prabhakar:

53. *Hybrid Systems: Computation and Control* (HSCC 2013).

César Sánchez:

54. *19th International Symposium on Temporal Representation and Reasoning* (TIME 2012).

55. *XII Jornadas sobre Programación y Lenguajes* (PROLE 2012).

56. *5th IPM International Conference on Fundamentals of Software Engineering* (FSEN 2013).

6.3.2. Conference and Program Committee Chairmanships

Gilles Barthe:

1. Program Co-Chair, *2012 International Symposium on Engineering Secure Software and Systems* (ESSoS 2012).

Manuel Hermenegildo:

2. Program Chair, *2012 European Computer Science Summit* (ECSS 2012).

Aleksandar Nanevski:

3. Program Co-Chair, *2012 Workshop on Syntax and Semantics of Low-Level Languages* (LOLA 2012).

6.3.3. Editorial Boards and Conference Steering Committees

Anindya Banerjee:

1. Associate Editor, *Journal of Higher Order and Symbolic Computation*.

2. Editor, *Foundations and Trends in Programming Languages*.

John Gallagher:

3. Area Editor (Technical Notes and Rapid Publications), *Theory and Practice of Logic Programming*.

4. Steering Committee Member, *Conference on Partial Evaluation and Program Manipulation* (PEPM).

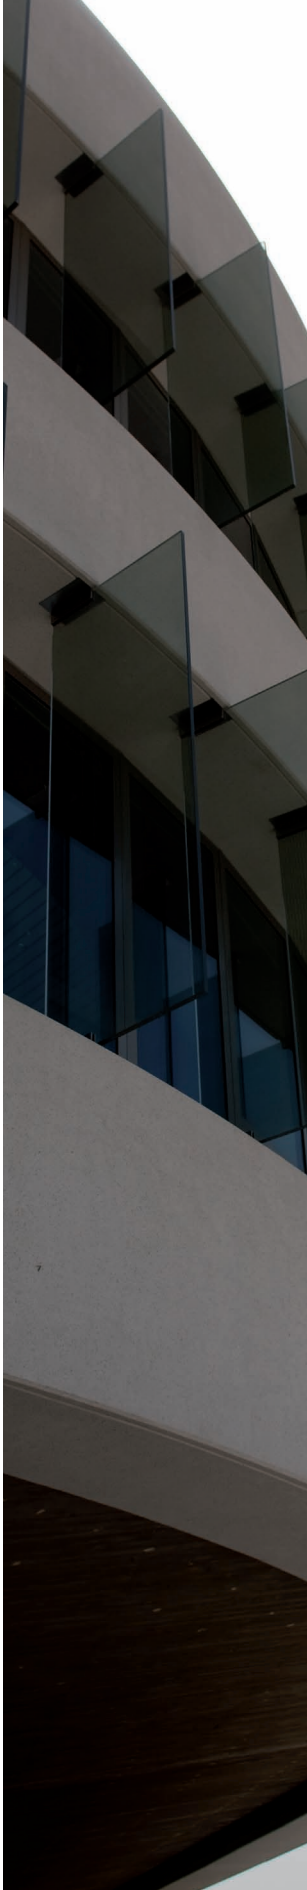

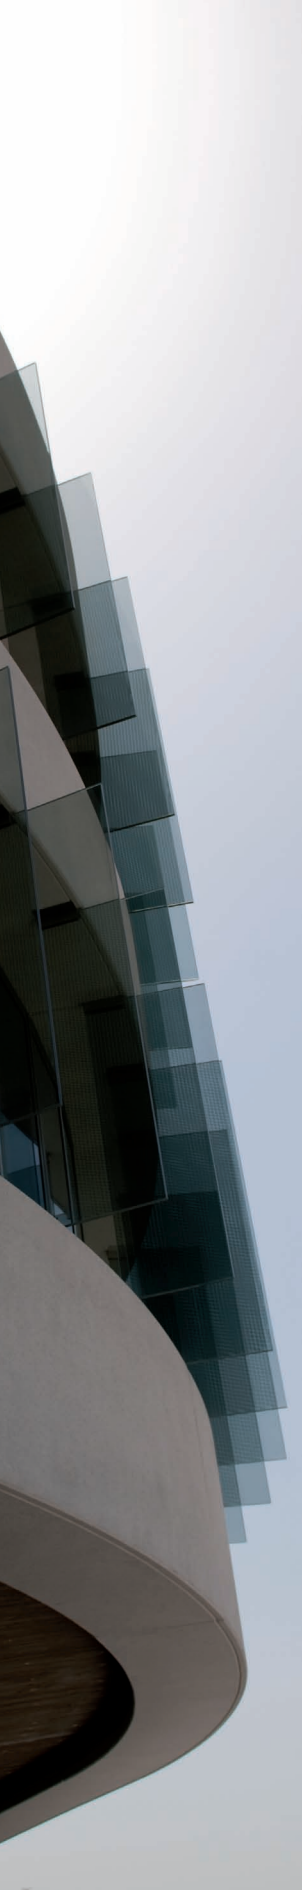

Manuel Hermenegildo:

5. Editorial Adviser, *Theory and Practice of Logic Programming*.

6. Area Editor, *Journal of Applied Logic*.

7. Associate Editor, *Journal of New Generation Computing*.

8. Board Member, *Journal of Algorithms in Cognition, Informatics, and Logic*.

9. Steering Committee Member, *ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages* (POPL).

10. Steering Committee Member, *International Static Analysis Symposium* (SAS).

11. Steering Committee Member, *International Symposium on Functional and Logic Programming* (FLOPS).

12. Steering Committee Member, *International Conference on Verification, Model Checking, and Abstract Interpretation* (VMCAI).

6.3.4. Association and Organization Committees

Gilles Barthe:

13. Co-Chair of the Working Group on Concurrency and Distribution for COST Action 0701: *Formal Verification of Object Oriented Software*.

Manuel Carro:

14. Deputy representative at Informatics Europe.

John Gallagher:

15. Executive Committee member for the Association for Logic Programming.

Alexey Gotsman:

16. Steering Committee member of the EPSRC Programme Grant on Resource Reasoning.

Manuel Hermenegildo:

17. Member of the Academia Europæa.

18. Member of the IFCoLog (International Federation in Computational Logic) advisory board.

19. Elected President of SpaRCIM (Spanish Research Consortium for Informatics and Mathematics).

20. Member of Informatics Europe department evaluation advisory board.

21. Member of IRILL (Free Software Institute) scientific board (France).

22. Member of Dagstuhl scientific board (Germany).

23. Member of CSIC (National Research Council) scientific board.

24. Member of the Comunidad de Madrid high school honors program faculty selection committee board.

Boris Köpf:

25. Co-organization of Dagstuhl Seminar 12481: *Quantitative Security Analysis*.

Juan José Moreno:

26. Head of the Organizing Committee, *Turing Year in Spain*.

César Sanchez:

27. Action Chair for COST Action IC0901: *Rich-Model Toolkit - An Infrastructure for Reliable Computer Systems*.

6.3.5. Awards

Conference paper awards

1. *Alexey Gotsman*, Hongseok Yang. *Linearizability with ownership transfer*. Proceedings of the 23rd International Conference on Concurrency Theory (CONCUR'12), Newcastle upon Tyne, UK, LNCS, Vol. 7454, pages 256–271, Springer, 2012. Best paper award.

2. *Pavithra Prabhakar*, Geir E. Dullerud, Mahesh Viswanathan. *Pre-orders for reasoning about stability*. Hybrid Systems: Computation and Control (part of CPS Week 2012), HSCC'12, Beijing, China, April 17-19, 2012, pages 197–206, ACM, 2012. Honorable mention award.

Research Awards

1. *Alexey Gotsman* received one of the 10 Software Engineering Innovation Foundation (SEIF) Awards given by Microsoft in 2012 for his research project *Specifying and Validating Components on Memory Models of Mobile Platforms*.

2. *Mark Marron* received another one of the Software Engineering Innovation Foundation (SEIF) Awards 2012 given by Microsoft in 2012 for his research project *MemAlyzer: Finding and Fixing Memory Usage Problems*.

Other Awards

1. *Juan José Moreno* received the Silver Medal in Psychology in June 2012 awarded by Spanish Official Psychological Association.

Researcher's Night

In September 2012, the IMDEA Software Institute participated in the European-wide initiative "Researchers' Night". The panel "Science is my Life" included IMDEA Software Institute researcher Boris Köpf (and researchers from other IMDEAs), and was moderated by Manuel Carro, the Deputy Director of the IMDEA Software Institute.

Scientific Highlights

annual report 2012

towards "gree

verifying & controlling computing resource consumption

Towards "Greener" Software: Verifying & Controlling Computing Resource Consumption

Energy consumption and the environmental impact of computing technologies have become a major worldwide concern. This is motivated in part by recent sustainability studies and also by the need to optimize the energy consumption of increasingly demanded complex computing systems which have to operate with limited batteries (e.g., implantable/portable medical devices or mobile computing). The growth of cloud computing-related energy consumption and Internet traffic is not sustainable with the energy efficiency levels of current computing technology. Energy consumption has also become a major concern in high-performance computing, distributed applications, and data centers.

In spite of the huge recent advances in power-efficient hardware, most of the potential energy savings are wasted by software that does not exploit these hardware energy-saving features, and by poor dynamic management of tasks and resources. Facing this challenge, the IMDEA Software Institute aims at promoting energy efficiency to a first-class software design goal, and, in a more general context, at developing tools that facilitate the production of "greener" devices, i.e., devices that make a certifiably more efficient use of their available resources *resources* (e.g., energy, execution time, memory, as well as other user-defined magnitudes like network accesses or transactions). IMDEA Software researchers are working towards achieving radical advances in energy-aware software design and management. These advances include the explicit exploitation of power-efficient features offered by hardware supporting conventional computation models, as well as by emerging approaches such as massively parallel systems and biologically inspired computation models.

Most of this research on "greener software" is being performed within the scope of the recently granted project ENTRA, a European Union Seventh Framework Programme (FP7/2007-2013) FET project, in collaboration with XMOS (UK), The University of Bristol (UK) and Roskilde University (Denmark). The ENTRA project proposes a novel, holistic, energy-aware system development approach that covers hardware, software, and the run-time environment, making information on energy usage available through-

ner" software

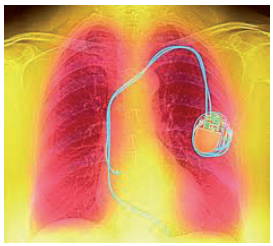

out the system layers and promoting optimizations both during code development and at runtime. This approach requires an important effort on developing novel analyses, combined hardware-software energy modeling, and verification and optimization techniques.

IMDEA Software researchers have already extended conventional debugging and verification techniques to deal with resource usage properties, and developed automatic optimization techniques that reduce significantly the resource usage of programs, in particular the total execution time and power use.

All the developed techniques are implemented and integrated into IMDEA's state-of-the-art tools, which are demonstrated to industrial collaborators and tested on concrete examples extracted from their application codes. For example, the pioneering CiaoPP system provides a general framework for estimating with high precision the resources consumed by a given piece of software and for debugging/certifying such consumption with respect to specifications. The tool is highly adaptable to different languages, hardware, and resources because it is built around a customizable static analyzer with a versatile and well-defined assertion language.

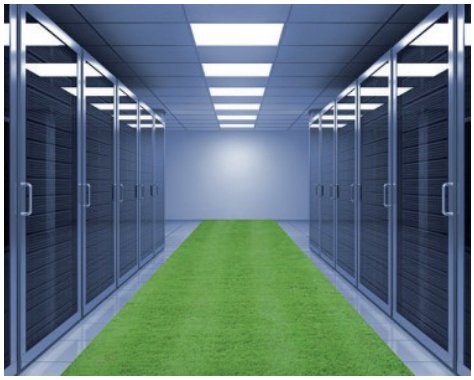

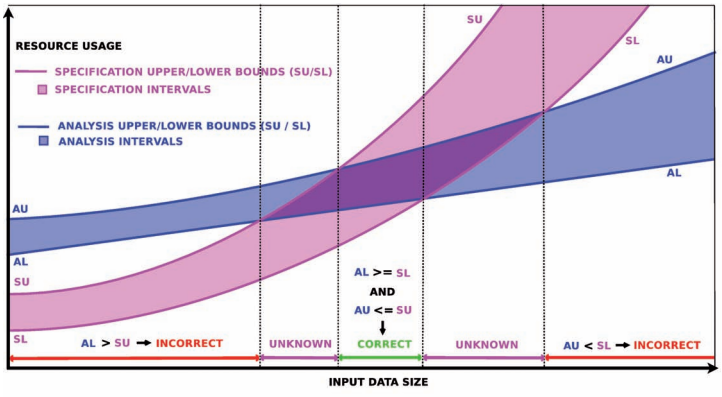

understanding the role o

Understanding the Role of Malware in Cybercrime

Cybercrime, criminal activity conducted via computers connected to the Internet, is a growing threat for developed regions like Europe where nearly three quarters of the households and a large number of the infrastructures are connected to the Internet, and an increasingly number of services and transactions happen on line.

At the core of most cybercrime operations is the attacker's ability to install malicious programs (i.e., malware) on Internet-connected computers without the owner's informed consent. Malware includes bots, viruses, trojans, rootkits, fake software, and spyware. Malware enables attackers to establish a permanent presence in the compromised computer and to leverage it for their cybercrime operations. The target of those operations may be the compromised computers themselves (e.g., stealing from the computer an organization's intellectual property or a user's banking credentials), but also third parties. In the latter case, the compromised computers are simply assets, which the attacker employs to launch malicious activities such as sending spam, launching denial-of-service (DoS) attacks, faking user clicks on online advertisements (i.e., click-fraud), or simply as a stepping stone to hide its location.

The goal of the MALICIA project at the IMDEA Software Institute is to study the crucial role of malware in cybercrime and the rise in recent years of a far-reaching "underground economy" associated with malware and the subversion of Internet-connected computers. Long gone are the days where attackers compromised computers and built malware to show off their skills to peers. Nowadays, the malware ecosystem revolves around cybercrime and the monetization of compromised computers.

As the malware ecosystem has grown larger and more profitable, specialization has come to the marketplace. Attackers have understood that tackling the entire value-chain from malware creation to monetization poses a daunting task requiring highly developed skills and resources. As a result, specialized services have been created at all stages in the malware-monetization chain, such as toolkits to automate the construction of malware, program encryption tools to evade antivirus (AV) software, "bullet-proof" hosting, and forums for buying and selling ill-gotten gains. Specialized services lower the barrier to enter the malware ecosystem. However, defenders can also take advantage of specialization, as disrupting the specialized services disrupts the different malware operations using them.

As a first step in the MALICIA project, we have collaborated with researchers at the University of California, Berkeley and the International Computer Science Institute to

f malware in cybercrime

investigate the commoditization of malware distribution in the form of *pay-per-install* (PPI) services. PPI services offer criminals a simple way to outsource the distribution of their malware. The clients provide their malware to the PPI service and select the number of desired installations (called *installs*) in each geographical area. The PPI service takes care of installing the malware on compromised computers in exchange for a small fee that ranges from $180 for a thousand computers in some European countries and the US, down to $7 for a thousand computers in Asia.

To satisfy the clients' demand for installs, the PPI provider typically outsources malware distribution to third parties called *affiliates*. PPI providers pay affiliates for each compromised computer, acting as a middle man that sells installs to the clients while buying installs from affiliates. Each affiliate may specialize in some specific malware distribution method (e.g., bundling malware with a benign program and distributing the bundle via file-sharing networks; exploiting web browsers through drive-by-downloads; or social engineering). The PPI service gives each affiliate a downloader program customized with a unique affiliate identifier. When the affiliate installs the downloader in a compromised computer, the downloader connects back to the PPI service to download the client programs. After installing the client programs on the compromised host, the downloader reports the affiliate identifier and the affiliate is credited with an install.

To understand the PPI market we *infiltrated* four PPI services. For this, we developed infrastructure enabling us to (1) interact with PPI services by mimicking the protocol interactions they expect to receive from affiliates, and (2) collect and classify the malware being distributed by the PPI services. Using this infrastructure we harvested over a million malware programs and classified them by malware family as well as monetization methods. Our analysis revealed that of the world's top 20 malware families, 12 employed PPI services for their distribution. It also revealed that some malware families exclusively target the US and a variety of European countries. The monetization methods in use are wide including: spam, installing fake antivirus software, information-stealing, denial-of-service, click-fraud, and adware.

Much remains to be learnt about the malware ecosystem and the specialized economy supporting cybercrime. Our current work strives on deepening our understanding of other parts of the ecosystem. One overarching goal is evolving malware analysis from understanding what a malware program does, to also cover *why* it does it, i.e., what role the malware program plays in the cybercrime operation where it is used.

The research in the MALICIA project has so far received an "Outstanding Paper" award at the 2011 USENIX Security Symposium and has been covered by MIT's Technology Review.

secure data

from models to code (automatically)

Secure Data Management: From Models to Code (Automatically)

Data-management applications are focused on the so-called CRUD actions that create, read, update, and delete data. Such applications are often implemented as multi-tier systems where the application manipulates data stored in a database and interacts with users through a graphical interface (GUI). When the data managed is sensitive, then security is a concern.

Implementing access control over data is nontrivial. Fine-grained access control policies may depend not only on the user's credentials but also on the satisfaction of constraints on the state of the persistence layer, i.e., on the values of stored data items. In such cases, authorization checks are typically implemented programmatically, by directly encoding them at appropriate places in the application. This is cumbersome, error prone, and scales poorly. Moreover, it is difficult to audit and maintain as the access checks are spread throughout the code and security policy updates require code changes.

We are creating a methodology for the model-driven development of secure data-management applications. It consists of formal languages for modeling multi-tier systems, and tools for generating these systems. Within our methodology, a secure data-management application is modeled using three interrelated models: a data model, a security model, and a GUI model. The heart of this methodology, illustrated in Figure 7.1, is a well-defined model-transformation function that automatically lifts the policy that is specified in the security model to the GUI model. The resulting GUI model is security aware.

Security-aware GUI models are platform independent and can be mapped to implementations employing different technologies, including desktop applications, web applications, and mobile applications. As part of our research, we have built the ActionGUI Toolkit (http://www.actiongui.org), which automatically generates Java web applications from security-aware GUI models. The ActionGUI Toolkit features model editors for constructing and manipulating data, security, and GUI models. Moreover, it implements our model transformation to generate security-aware GUI models. Finally, it includes a

management

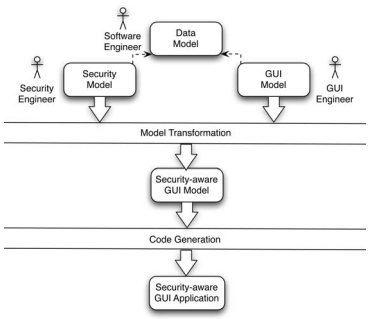

Model-driven development of security-aware GUIs.

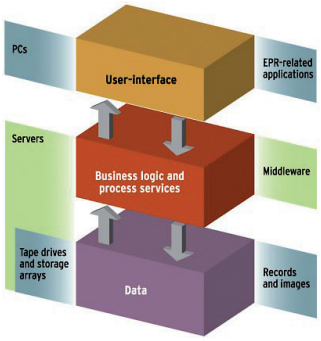

3-tier architecture

code generator that, given a security-aware GUI model, automatically produces a Java web application, ready to be deployed in a web application container (e.g., Glassfish or Tomcat).

Overall, the contribution of this research is two-fold. First, our methodology offers model-driven engineering's purported benefits for data-management systems. By working with models, designers can focus on the application's data, behavior, security, and presentation, independently from the different, often complex, technologies that are used to implement them. Second, the use of model transformations leads to modularity and separation of concerns: the GUI model and the security model can be changed independently and by different developers, if desired. This avoids the problems associated with brittle, error prone, hardcoded security policies that are difficult to maintain and audit. Of course, the proof of the pudding is in the eating and we report on the ActionGUI's web portal (http://www.actiongui.org) about example applications that we have already developed using the ActionGUI Toolkit. This work is performed in close collaboration with the Information System Group at ETH Zürich, led by Prof. David Basin. ActionGUI-1.0 is registered software from the IMDEA Software Institute and the Institute is working on its commercialization, in collaboration with ETH Zurich.

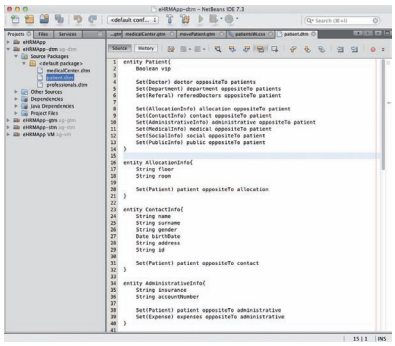

ActionGUI: An example of a data model (screenshot)

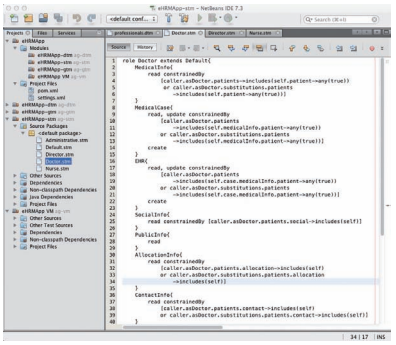

ActionGUI: An example of a security model (screenshot)

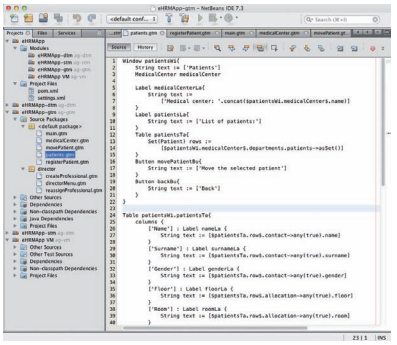

ActionGUI: An example of a GUI model (screenshot)

when security

computer-aided cryptographic proofs

When Security Matters: Computer-Aided Cryptographic Proofs

Modern cryptography is the science of developing methods for protecting information and communication against misbehaving parties. Initiated with the pioneering work of Shannon on secrecy in encryption systems in 1949, modern cryptography has become an active field of research with the discovery of public-key cryptography by Diffie and Hellman in 1976, the invention of the RSA algorithm by Rivest, Shamir and Adleman in 1978, and the conception of provable security by Goldwasser and Micali in 1984. Initially focused on encrypted communications over insecure channels, cryptography has expanded considerably to achieve a broad range of security goals, from basic ones such as confidentiality and integrity, to elaborate goals such as proofs of knowledge. In parallel, applications of cryptography have outgrown the domain of military and diplomatic communications, to play a central role in the Internet, the Cloud, and more generally in any massively distributed infrastructure that can store and process huge quantities of data and computations. In effect, billions of individuals, companies, and institutions use cryptography routinely for interacting with each other. Online banking systems, electronic health records, cash machines, cell phones, or digital identity management systems are only a few examples of the numerous applications that rely on cryptography.

German Enigma machine.

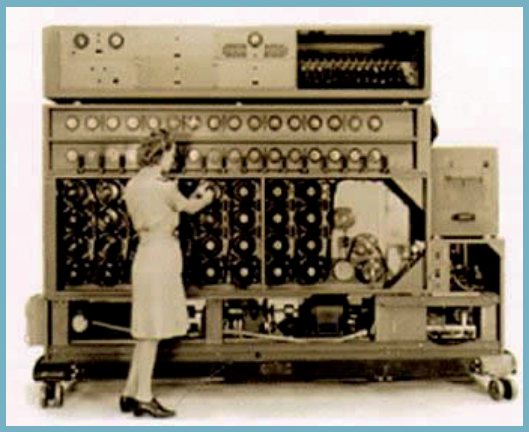

Polish/British special purpose Enigma code-breaking machine, called "bombe".

matters

Contactless smart-cards have become a standard for biometric identification documents worldwide.

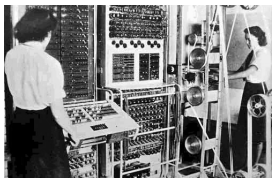

Modern cryptography achieves its goals by extracting from specialized branches of pure mathematics, such as number theory, very efficient algorithms that can be used in practical systems for the purpose of granting them an adequate level of security. The downside to such an amazing feat is that cryptographic algorithms are very difficult to analyze. Thus, cryptographers devote a significant amount of time to building rigorous mathematical proofs of the security of a cryptographic scheme. The goal of these proofs is to show that the cost of attacking a cryptographic system is prohibitive, and to determine the parameters that must be used—for example, the size of keys—in any practical realization. Building such cryptographic proofs is far from being an academic problem: in fact, there are critical gaps in the security proofs of the most widely used cryptographic schemes and protocols. For instance, researchers from IMDEA Software Institute, INRIA, and Université of Grenoble have recently unveiled an inaccuracy in the proof of the RSA-OAEP, a widely deployed encryption scheme that is recommended by several standards, including IEEE P1363, PKCS, ISO 18033-2, ANSI X9, CRYPTREC and SET.

In order to ensure that cryptographic systems achieve their purported security requirements, researchers from IMDEA Software Institute, INRIA and Microsoft are developing EasyCrypt, an automated tool that supports a radically new approach to cryptographic proofs. While adhering to the principles and the methods of provable security, EasyCrypt takes the view that cryptographic proofs should be treated in a manner similar to high-integrity software, so that confidence in the design of a cryptographic system is no lower than confidence in the software systems that use it. In order to realize its ambition, and to provide working cryptographers with practical tools for building trustworthy and verifiable proofs, EasyCrypt builds upon state-of-the-art verification tools, including SMT solvers, automated theorem provers, and proof assistants.

Although its development is in its initial stage, EasyCrypt has attracted considerable interest from the academic and industrial communities. The article “Computer-Aided Security Proofs for the Working Cryptographer,” by Gilles Barthe, Benjamin Grégoire, Sylvain Heraud, and Santiago Zanella Béguelin, received the Best Paper Award at CRYPTO’11, the premier conference in cryptography; besides, Santiago Zanella Béguelin received the 2011 EAPLS Best PhD Dissertation Award. Researchers at the IMDEA Software Institute and INRIA have also initiated collaborations around EasyCrypt with many leading researchers in several European and US academic and research institutions.

taming the

Taming the Marriage of the Cyber and the Physical Worlds

of the cyber and the physical worlds

Today, computers are no longer standalone devices sitting on our desktops, but are increasingly found embedded in everyday devices, systems and structures. These include smart phones, smart buildings, smart medical devices, next generation air-traffic control systems and automobiles, and smart grids. The drastic reduction in the cost of sensing, actuating, computing and communicating technology has enabled the proliferation of this new genre of engineered systems in which a network of embedded processors interact tightly with the physical world to achieve complex functionalities. These systems are referred to as Cyber-Physical Systems owing to the coupling between the networked computation and physical systems.

Cyber-physical systems are the systems of the future and will have an impact on the engineering systems technology comparable to the impact the internet had on the information systems. They have applications in a wide range of systems spanning communication, infrastructure, energy, health-care, manufacturing, military, robotics, and transportation. Governments around the world have taken several initiatives to exploit this potential. The report of the US Presidents Council of Advisors on Science and Technology (PCAST) has placed Cyber-Physical Systems on the top of the priority list for federal research investment. The European Union has recognized the strategic importance of Embedded Computing Systems and has launched the ARTEMIS Joint Technology Initiative (JTI) as part of the FP7 program.

Cyber-Physical Systems have immense potential for a long-term impact on the society. At the same time, the unprecedented complexity arising due to the interleaving of the cyber and the physical components is overwhelming. On one hand, digital systems operate in a discrete manner, where computation and communication proceed in synchronization with the processor cycles. On the other hand, physical systems execute continuously in dense real-time. Hence, cyber-physical systems are hybrid systems exhibiting both discrete and continuous behaviors, and are networked and/or distributed with possibly humans in the loop. The grand challenge of the near future is the development

marriage

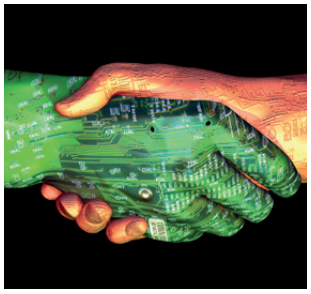

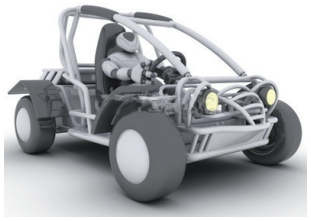

of design methodologies and tools to cater to the development of reliable cyber-physical systems.

The realization of high-confidence cyber-physical systems requires addressing a plethora of issues ranging from specification, modeling and analysis of heterogeneous models, networking, interoperability, time-synchronization, scalability and complexity management through modularity and composability, and validation and verification. This, in turn, requires the interaction of several communities of researchers and developers working in areas which overlap with cyber-physical systems. This includes embedded and real-time systems, control systems, formal methods, hybrid systems, distributed systems, and application domains such as aeronautics, automotive, and robotics.

In particular, the interaction between the cyber and the physical worlds brings new challenges. Firstly, there exist strong foundations for the analysis of purely discrete systems in computer science, and similarly, for purely continuous physical systems in control theory. However, it is often not straight-forward to extend the techniques to the systems with coupled cyber and physical behaviors. Secondly, there arises new concepts and notions from one world which are almost alien to the other. For instance, stability is a well-studied problem in control theory, however, it is a totally new concept for discrete systems. In summary, the opportunities for the scientific and technological research and development in bridging the gap between the two worlds are huge.

Researchers at IMDEA are actively involved in this exciting new area by focusing on the development of the state-of-the-art technology for verification of cyber-physical systems, particularly, in the early design phase which is seen to have a huge impact on the development cost of the products. They are addressing problems related to scalability of the current verification techniques by designing novel state-space reduction algorithms and tools. To this end they are collaborating with researchers from several disciplines around the world including researchers at the University of Illinois at Urbana-Champaign and California Institute of Technology in the US, University of Sheffield in UK, and VERIMAG in France.

the many shad

The Many Shades of Security: When Security is Not Just Yes or No

when security is not just yes or no

An ever-increasing number of security-critical processes rely on software, and a growing amount of personal data is stored in the cloud. This pervasiveness of software systems opens up tremendous opportunities for society, economy, and individuals—but it also exposes users to attacks against their security and privacy. For the information society to thrive, we need to rigorously secure our software systems. Unfortunately, security is typically in conflict with requirements on functionality, performance, usability, and cost:

- The release of sensitive information such as medical records or search histories severely threatens the privacy of individuals and organizations. However, society and the economy benefit from utilizing this sensitive data, for example, for computing health statistics or for targeted advertisement.

- Techniques such as caching and pipelining dramatically improve the performance of today's microprocessors. However, they also introduce large variations into the execution time and power consumption of computations. Such variations can be observed and exploited by so-called side-channel attacks to recover sensitive information, such as cryptographic keys.

- Password policies, such as the requirement that user passwords contain numbers and upper-case letters, aim to decrease the likelihood of password compromise. However, overly restrictive password policies can decrease usability and negatively impact user productivity.

In the presence of such conflicting requirements, *perfect* security is impossible or simply too expensive to achieve. Rather, the challenge is to identify a trade-off in which the system meets its requirements and at the same time provides a *sufficient* degree of security. Today's methods for evaluating the degree of security are based on parameters such as infection rates or the number of reported security vulnerabilities, which can be obtained by mining system logs and vulnerability reports. However, the counting of

es of security

security incidents is not an acceptable basis for evaluating the security of systems in application domains in which a single incident can already imply critical loss.

Researchers at the IMDEA Software Institute are working on the next generation of security metrics, which are suitable for evaluating safety-critical systems. One key requirement is that these metrics be computable statically from the program code and the platform specification, that is, before the system is under attack. Another requirement is that these metrics come with operational interpretations (such as lower bounds for the effort required to break into a system) that can be used as a basis for managing security risks in application domains where post-mortem analysis is not an option.

One of the highlights of the work on security metrics at IMDEA Software is the development of CacheAudit, a tool for the automatic, static quantification of microarchitectural side-channels. Microarchitectural side-channels are an growing concern in the cloud, where they can be exploited to mount attacks across virtual machine boundaries. CacheAudit allows to derive quantitative guarantees against this kind of attack; for example, it has been used to derive the first proof of resilience against cache attacks of central parts of a widely deployed cryptographic library. CacheAudit is being developed in active collaboration with researchers from Siemens AG and Saarland University. Further collaboration partners on security metrics for privacy-preserving computations are Microsoft Research Cambridge and the Max Planck Institute for Software Systems.

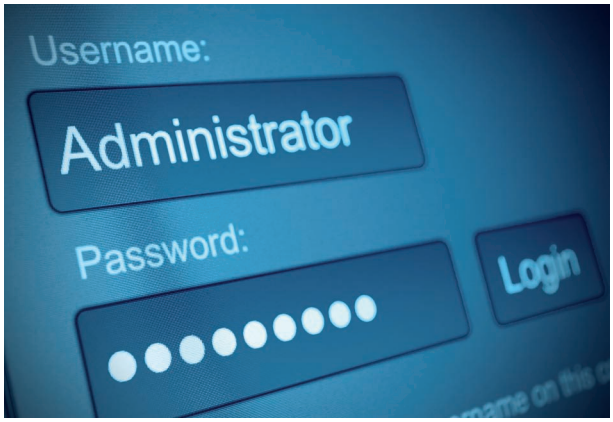

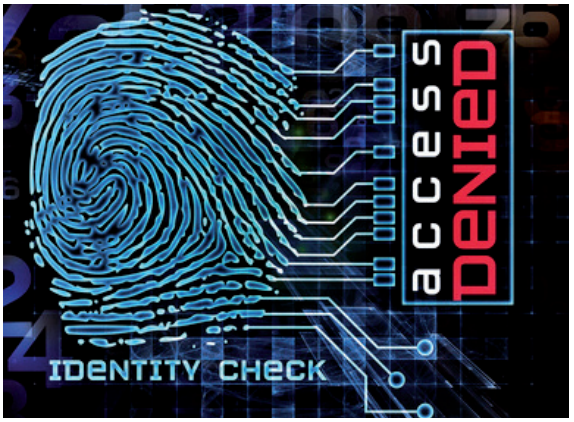

architecture-dri

Architecture-Driven Verification: Tackling The Complexity Of Modern Software

tackling the complexity of modern software

The modern information society critically relies on software systems, with some of its most vital processes, such as communication, power grids, railroads and finance, controlled by software. At the same time, software is notoriously unreliable: errors in it can cost billions of euros or even lives. This is a consequence of a systemic problem. Recent studies have put the cost of software bugs to the economy of the US alone at $60 billion a year or about 0.6% of GDP. A whopping 80% of the software development costs of a typical project are spent on identifying and fixing defects. Given the current trends in software development, in the future the cost and dependability problems will only be exacerbated. Extensive testing and carefully structured development processes can improve the quality of software, but do not guarantee that it is bug-free. The growing dependence of modern society on software systems makes this situation unsustainable.

Software verification is an area of computer science that has the potential to resolve this problem: its goal is to ensure the correctness of software by proving that it satisfies a given property in *all* possible situations. Formerly a purely theoretical area, since the year 2000 software verification has experienced a resurgence of interest from practitioners and is now emerging as a cutting-edge approach to improving software quality. However, even though there is much excitement in the verification area, there is still a huge distance to go before we will be able to verify pieces of software as complex as a modern operating system kernel. This is because the cost-benefit ratio of current verification technology is not good enough to scale it to major software systems. Software verification is currently good at dealing with programs that are either big, but simple, or complicated, but small. Unfortunately, modern software is both big and complicated. IMDEA Software Institute researchers are developing radically new verification methods aiming to overcome this problem.

The hypothesis underlying this research is that the main reason for the inadequacy of the existing verification approaches when dealing with complex software is their generality. The techniques they suggest are based on generic principles that come from properties of programming languages, which allows applying such techniques to arbitrary programs. However, since they cannot take advantage of the particular ways in which programs are usually written in those languages, they require too much labor and do not scale to

ven verification

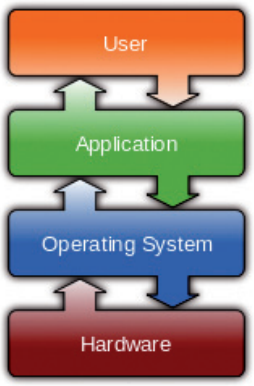

big and complicated systems. At the IMDEA Software Institute, we are developing methods and tools for cost-effective verification of real-world systems software by exploiting the way programmers write it: in practice, they stick to informally described patterns, idioms, abstractions and other forms of structure contained in their software, which are together called its *architecture*. IMDEA researchers harness this trend to develop verification methods and tools that are tailored to the architectures used in modern systems software. This is a radically novel approach to the problem of verifying complex software: instead of building generic verification tools, we are attempting to give a formal footing to software engineering concepts used by systems programmers to reason about their software informally, and use the results to drive the design of verification techniques.

Such architecture-driven techniques have the potential to result in verification tools that require a minimal and intuitive input from their user—no more than a formal version of the high-level informal specifications programmers already have in mind when developing software. In time, this can yield a dramatic leap in the cost-benefit ratio of the verification technology, allowing it to scale to systems of real-world size and complexity that have so far been beyond the reach of quality assurance methods guaranteeing correctness.

Most of this research will be performed within the scope of the recently granted EU project ADVENT, an FP7 FET Young Explorers project including Katholieke Universiteit Leuven (Belgium), Max Planck Institute for Software Systems (Germany) and Tel-Aviv University (Israel). The research is also supported by a Microsoft Software Engineering Innovation Foundation Award and a Microsoft European PhD Scholarship.

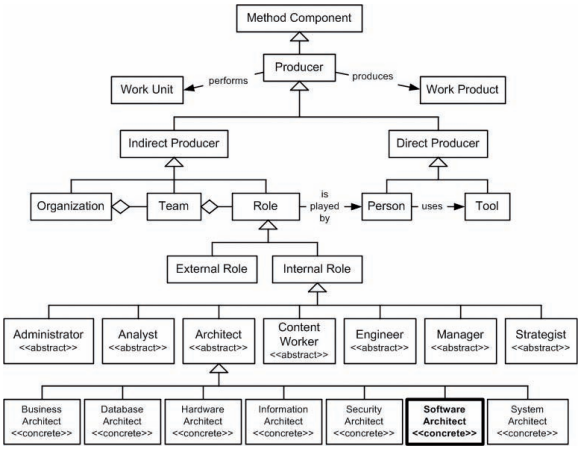

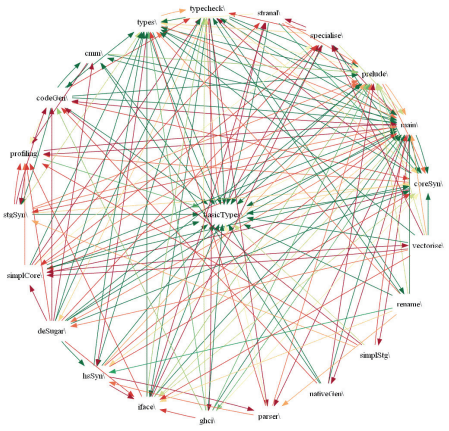

high integri

when software must not fail

High Integrity Software: When Software Must Not Fail

Some software cannot fail, in some real world applications. This software is called high integrity software and must be trusted to work dependably in some critical function. Failure in these programs may have catastrophic results in terms of lives or have high economic cost. For example, failure in a program used by air traffic controllers could lead to fatal accidents; a failure in a medical system or a medical device could lead to irreversible damage; failures in parts of automobile systems such as brake controllers, apart from being potentially dangerous, could lead to massive and costly recalls. In fact, all of these scenarios have occurred already.

Examples of high integrity software include safety systems of nuclear power plants, medical devices, air traffic control, automated manufacturing and satellite control. In software that runs power grids, financial systems, water treatment plants and other critical elements of a country's national infrastructure another dimension of high integrity must be guaranteed: such software must further be secure against cyberattacks and thus must guarantee that they meet security requirements.

Current software engineering practices balance between the cost (and time) to complete a project on the one hand, and software quality on the other. The rapidly growing demand for software that provide ever-more complex functionality has increased industry demand for software developers. In turn, this need has motivated a trend towards

Satellites have to be autonomous up to a certain degree. The programs running in their computers continuously monitor for deviations from their scheduled trajectories and take the appropriate decisions to correct them.

ty software



reducing the training necessary for software engineers and developers to enter the job market. However, at the same time, the quality of the software produced has become more and more difficult to guarantee. The issue with software quality is witnessed by the fact that the dominant factor of the overall cost of current industrial software projects is testing, and not building the product itself. Even in non-critical projects testing dominates more than 90% of the total cost.

High integrity software requires the replacement of disclaimers (as is current practice) by guarantees such as functional correctness and security. The US government is currently discussing legislation to ensure such guarantees from software vendors.

The quality and reliability requirements of high integrity software justifies the investment in scientific undertakings to create a body of knowledge about how to build more reliable software. These new techniques intend to provide better guarantees of quality, at the price of using more sophisticated methods and tools by properly trained software engineers. These methods encompass foundational theories, tools and convincing experiments and their importance cannot be overstated. While in the “here and now” they lead to more productive software processes for industry, their long term goal is to develop processes that can be applied to the widest possible range of application software, far beyond current demands of industry or popularity in the market place.

Researchers from the IMDEA Software Institute have developed – and continue to develop – cutting-edge technologies for high-integrity software following two approaches.

The first approach is foundational: it is aimed towards creating the basic science, based on rigorous mathematics, that can be used to craft the high-integrity software of the future. These techniques are designed to provide the best guarantee of adherence to intended behavior and are at the heart of the robust, scalable tools, that serve as testbeds for our foundational approach. Completed and ongoing projects include the use of high-order theorem proving to verify programs and libraries, static analysis for functional and non-functional properties of real-time, embedded and reactive systems.

The second approach concerns the development of novel lightweight and applicable techniques that can be directly incorporated to improve existing software practices: advanced visualization of heap-manipulating programs, debugging of production system programs, and online monitoring of embedded reactive programs based on runtime verification.

Institute researchers have also launched a significant effort towards verifying concurrent software, particularly concurrent data type libraries and operating systems software. This is a particularly challenging area and requires a foundational re-think from current practices which, in industry, are based predominantly on testing. For instance, it is well known that concurrency bugs owing to race conditions and deadlocks are very hard to detect. Even the seemingly easy task of reproducing a bug becomes a challenge due to influence of factors such as the current CPU workload.

The IMDEA Software Institute is collaborating with the leading aerospace company Deimos, located in the area of Madrid, on the technology transfer of these techniques. This continuing effort started with the rigorous and systematic development of software for satellite image processing. The aim of this project is to develop the tools to interactively synthesize provably correct software, based on a formal approach to software families, applied to image processing. Using these tools software engineers can develop very efficient parallel software that can be verified with independent tools. Moreover, different projects can experience dramatic cost gains by the increase in the level of reuse by the use of software families.

As mentioned before, an important dimension of high integrity is that software meets security requirements. Current computing environments and infrastructures are increasingly heterogeneous and dynamically changing. Executable programs are everywhere: web pages, email, plug-and-play extensions, JavaScript, on-line games, Word and PowerPoint documents and attachments, electronic banking, etc. Software is constantly being updated and downloaded over the Internet, sometimes without our knowledge or consent.

Yet, today's security architectures provide poor protection from faulty software, and even less from malicious software. These security architectures were developed at a time when software was managed and updated infrequently by an experienced administrator, when we trusted the (few) programs we ran, when physical access to the systems

was required to cause any damage to the data, and crashes and outages did not cost billions.

As none of these conditions is valid anymore, our information systems have become increasingly susceptible to attacks with potentially devastating consequences.

To accommodate for the new trends in software use and deployment, researchers at the IMDEA Software Institute are working on theories and tools for building trustworthy software that are are well suited for networked computing systems built from diverse and extensible components. By leveraging techniques from programming languages and program logics researchers are addressing the following fundamental issues: (a) what is the precise meaning of security policies, (b) how to correctly specify security policies, (c) how to prove that programs respect the policies and (d) how to provide verifiable evidence, checked by a machine, of proofs of conformance of programs to security policies.

Institute researchers have also shown that such foundational security infrastructure can be put to use in practice. For example, in the Mobius project, jointly with France Telecom and INRIA, they have shown the feasibility of on-device checking of mathematical proofs, using dedicated checkers developed and extracted from rigorous mathematical formalizations in the proof assistant Coq.

editor
imdea software institute

graphic design
base 12 diseño y comunicación

photos on pages 11, 12, 19, 57 & 82
Daniel Schäfer

legal deposit number
M-15.509-2013