Fostering Systems Research in Europe

A White Paper by EuroSys,
the European Professional Society in Systems

Peter Druschel
Thomas Gross (chair)
Rebecca Isaacs
Marc Shapiro

21 April 2006

1 Executive Summary

The Computer Systems discipline (which encompasses the sub-areas of operating systems, distributed, embedded, real time and pervasive systems) constitutes a central pillar of computer science. Systems research is the scientific study, analysis, modeling and engineering of effective software platforms. Its challenge is to provide dependable, powerful, performant, secure and scalable solutions within an increasingly complex IT environment. As toolsmiths fueled the Industrial Revolution, today Systems researchers lay the foundation for IT services and applications in the Knowledge Era. Healthy research in Systems is therefore essential for the success and continuing innovation of the IT-based industry (be it proprietary or open source) in Europe.

Europe contributed many early innovations in Systems and continues to produce significant successes; yet it tends to be overshadowed by research in the US. We find several systemic reasons for this, which need to be addressed. Among others: (1) Overall, Europe under-invests in fundamental research in Systems; (2) the structure and culture of academic institutions do not consistently foster excellence at all levels of Systems education and research; (3) In general, European research groups are isolated, and need to network more effectively amongst each other, with their peer groups in the US and other parts of the world, and with the IT industry. If nothing is done, Systems research in Europe will decline, drying up the roots of innovation. This will negatively impact, not only the European IT industry, but beyond it, all sectors that are IT-based, e.g., financial services, government, health care, education, and manufacturers of high-value products such as aircrafts and cars.

The European Systems community has started to address these issues through improved networking and by raising awareness among leaders in the business community, at universities, at funding agencies and among policy makers. This is a good start, but to excel, more is needed: changes by all players are necessary to improve the Systems research landscape. We make specific recommendations, which are detailed and justified in the main body of this paper. Here is a short summary:

- Universities: The top priority is to foster excellence at all levels of education and research. For students, we make the following recommendations: establish "Research Masters" programmes feeding into a PhD; ease time limits on PhDs; generalise doctoral internships; encourage student exchanges. For faculty, we recommend outside hiring, evaluation involving outside peers, and evaluation metrics adapted to Systems. To compete for the best young talent, institutions should offer competitive working conditions, including stability, responsibility and significant career prospects. In particular, junior faculty should have modest teaching load and receive mentorship, while enjoying the freedom to pursue their own research agenda.
- Industry: Our proposals aim to encourage to innovation and technology transfer. The Systems research community and industry need to improve their interaction. Each side needs to better appreciate the other's needs and roles; e.g., intellectual property issues and the value of fundamental, risky, long-term research and publication. The European IT industry should offer internships for PhD students and hire more PhDs.
- Funding agencies: Funding agencies should support long-term, focused, risky and fundamental research projects. Systems research and infrastructure investment need to be sustained over sufficiently long periods. Funding decisions need to be based primarily on technical criteria, such as quality and impact; political criteria (such as balance between EU countries) must come second for research projects.

There is currently a window of opportunity for attracting talented researchers to Europe and to establish Europe as the leading location for high-quality, high-impact Systems research. But to take advantage of this opportunity, the issues we raise need to be addressed now.

2 Introduction

Computer applications are so essential to modern life that it is easy to take them for granted and lose sight of their intricate design. "Systems" (or, more pedantically, Computer Systems) constitute a basic pillar of IT. The Systems discipline aims to tame the growing complexity of software, and is concerned with ensuring fundamental properties of a software system, including correctness, reliability, availability, usability, persistence, security, privacy, integrity, interoperability, portability, mantainability, scalability and performance. For this it uses well-founded techniques, methods and principles such as abstraction, virtualization, componentization, redundancy, logging, cryptography, authentication, and algorithmic theory. The Systems area covers, broadly, operating systems, middleware, distributed, pervasive and embedded systems, databases, filesystems and archival storage systems.

Just about every economic sector benefits from the considerable advances made in Systems research. For instance, the fact that the same application code runs unchanged, whether on a hand-held computer or on a 2048-node cluster (though at different speeds), is an example to the success of the Systems approach.

Another example is the ability for a user to roam across the world and access data and applications anywhere.

Systems research has enormous technical and economic importance. Witness the major companies that owe their success to their Systems expertise, for instance Akamai, Amazon, Apple, eBay, Esmertec, Google, Groove, IBM, Iona, iRobot, Jaluna, Microsoft, RedHat, SMC, Sun, Suse, Symbian, VMWare, or XenSource. Many of these leverage years of academic research, e.g., VMWare and Google from Stanford, or Akamai and iRobot from MIT. Many of these companies have their home in the US, although there are also European successes (either independent companies or units of a US company) that contribute to the creation of high-tech jobs in Europe.

Commercial achievement is not the only measure of success. Other metrics include publications at premier conferences, membership in prestigious programme committees, or prizes for doctoral students.

In specific areas, European Systems research efforts have significant impact. For instance, European teams play a major role in cluster and grid computing. System advances have made vital contributions to the telecommunications and embedded systems area, a field where Europe is extremely competitive, with examples such as GemPlus, Siemens, Infineon, ST-Microelectronics, Nokia, Airbus, etc. Another European strength is the open software community, largely kick-started by Linux, in which Europeans are very active.

Nonetheless, the overall impact of European Systems research is low, compared to the US. There is widespread agreement that Europe is falling behind the rest of the world in this area. Systems research in Europe is not fulfilling its potential to influence the computing products that we use daily. And indeed, Europe generally has little influence at the core of the software industry, despite many niche successes. This white paper attempts to assess the situation, provide an analysis and offer suggestions for improvement. Beyond a mere call for funding, we cast a critical eye on our own community and its academic, industrial, and institutional environment. Our aim is to improve both the production of knowledge in Europe and its transfer to European industry.

This white paper develops as follows. In Section 3, we define the term Systems research more precisely; Section 4 illustrates the area using a selection of "grand challenges", fundamental problems that remain to be solved. In Section 5 we provide an assessment of current European strengths and weaknesses. Then we make proposals to improve the situation, focusing on the different actors. Section 6 suggests improvements that the European Systems community can implement itself. Section 7 contains institutional recommendations to improve the training of PhDs and junior researchers. Proposals to encourage and reward junior faculty are found in Section 8. Section 9 discusses building relationships with the business community. Finally, we conclude with a summary in Section 10.

This white paper mirrors some of the conclusions of the recent Zürich Summit of Computer Science Faculty and Department Heads [13], but with a focus on the Systems area.

3 Systems research — what is it?

Systems research is concerned with the analysis, design, implementation, deployment and evaluation of complex software systems. Systems are a core area of Computer Science.[1] Compared to other areas, Systems research has strong distinguishing characteristics. First and foremost, Systems is an experimental discipline. The primary measure of success of a Systems project is whether it provides a practical and efficient solution to a real problem. A good design, a solid theoretical foundation, preliminary validation by formal proof and simulation are all important, but not sufficient. There is no alternative to implementation and experimentation to judge how well a design succeeds in practice.

It follows that a Systems project has large resource demands, in terms of human investment, resources and time. It involves designing and building a software artefact, integrating it into its environment, and evaluating how it performs its intended purpose. The art of Systems research demands the ability to understand complex systems, to design and implement systems that work, and to evaluate how an actual deployment will perform. This knowledge can only be acquired by practice, which takes time and effort.

For a potent example of the Systems approach at work, consider distributed system security. Since the 1980's our community has been at the forefront of research in networks and distributed systems. Security had been studied extensively in classical operating systems, yet the new world of open, co-operating, non-trusting distributed entities raised new, deep security issues. Three Systems researchers, Burrows, Abadi and Needham proceeded to study the fundamental issues; their study was published first at SOSP, the flagship Systems conference, then later in TOCS, our major journal [3, 4]. This work, commonly known as the BAN Logic of Authentication, has become a major intellectual underpinning for the theory and practice of security. It has been widely studied theoretically, and practical implementations are in use daily by millions of users.

4 Current research issues

To crystallise the most pressing issues in Systems research, we can envisage applications of great social value that appear "almost" within reach of current technology. In addition to being worthy and interesting engineering projects, they will serve (later in this section) to illustrate the range of profound and exciting challenges that must be addressed by Systems research.[2]

- Pervasive democracy: to open all EU information to all EU citizens; to promote efficiency through controlled information sharing across government organisations; to enable citizens' timely access to, and input on, decision making; yet to protect the system from manipulation and to protect citizens' privacy.

[1]For the purposes of this paper we use the terms Computer Science and Informatics interchangeably to refer to the scientific study of information processing.

[2]These examples focus on the area of distributed computing; there are comparable challenges in other Systems areas, e.g., embedded and real time systems. This list is inspired in part by discussions at a recent NSF workshop [17]. Another useful reference is the UK Grand Challenges exercise [20].

- IT for saving lives: to establish a pervasive medical information system, enabling doctors to monitor patients in real time, and to collect public health information at a large scale, making aggregate, anonymised health information readily available to the medical research community and health authorities.
- Disaster relief: to be able to quickly set up robust, secure data communication in disaster areas, and to securely share information in real time between relief teams, government and citizens.
- Monitoring our planet: setting up a world-wide, self-managing sensor network to enable scientists to understand the planet in real time; the network shall aggregate, merge and reduce information, and help scientists identify important trends and developments.
- Eternal memory: Managing a lifetime of personal data, reliable, securely, anywhere, anytime. An individual should be able to share the information of his or her choice, while being assured of privacy.
- Federation of academic libraries: To make academic literature available to all users world-wide, regardless of the location of the document or of the user; co-operative storage and search of large amounts of scientific documentation.

These examples highlight the many diverse aspects of Systems research, and underscore a number of *fundamental challenges* that must be addressed:

- **Security** is a major issue in open networks. Security must be considered in system design from the start, and cannot be added after the fact. Designs should minimise trust assumptions and provide isolation and virtualisation properties. There is a need for both static verification and run-time enforcement. Conversely, traceability and accountability are necessary, both of users who might abuse the system, and of the designers of the system who might cut corners. A system should have reflexive, self-inspection interfaces; attestation and auditing are essential. Privacy policies must be enforceable and open to citizen scrutiny.
- **Reliability:** Computers and networks are pervasive and are being used in applications that are mission critical and of high economic value. The Internet itself, initially designed as "best-effort," is quickly becoming a utility (just like electricity and water); even a brief outage has significant economic cost. Complex systems are failure-prone, yet the natural redundancy of the Internet can be leveraged to provide high reliability: for instance, peer-to-peer systems demonstrate self-administration and self-healing properties.
- **Scale and diversity** in pervasive computing: today's networks encompass a large variety of machines and administrative domains, where essential properties such as connectivity, computing power, administration and policies, etc., change from place to place and from time to time. Such diversity increases complexity. We need to design networked systems that scale up to encompass the globe, and scale down to run effectively inside the home. We need diverse devices to work together seamlessly when requested by their legitimate users.

- **Data storage and filtering in the network** provide new capabilities. They provide users with seamless access to diverse sources of information, with the power to process information and find what is relevant, and with virtually unlimited, reliable storage. Grid/peer-to-peer storage enables new applications, and reduces network and administration costs for everyone. Information should be shared securely and be highly available anytime, anywhere to all legitimate parties. Filtering capabilities find relevant data; the system provides functions to aggregate and anonymise data. The provenance of data must be tracked and the system should provide data auditing services.
- **Management complexity.** Current systems are rife with complexity issues such as software patching, feature interaction, network management, backup, etc. It must be possible to diagnose faults and understand why performance is not as it should be. End users want computer systems that are convenient, easy to use and predictable. Enterprise system administrators want to be able to set and enforce management, sharing and security policies. Most importantly, we need tools that enable us to *design* simple, well-isolated components, that are robust, reliable, and understandable.
- **Experimental facilities: Testbeds and workloads.** To enable the scientific study of Internet-scale phenomena, and to verify new network software before actual deployment, we need facilities to test and simulate large-scale systems under controlled conditions. Such testbeds are expensive and labour-intensive. A major issue is the lack of realistic workloads, for which the collaboration of vendors and service providers is needed (e.g., Windows crash traces or anonymized session data from telcos and ISPs).

The challenges are immense and have the potential to affect daily life. They cover a range of computer science fields, but it is up to the Systems community to provide the infrastructure. To tackle them, large-scale research is needed. Consider for instance the GENI Initiative being launched by the US National Science Foundation (NSF) [16], which sets out to re-examine the design principles of the Internet and to address many of the above issues. There is no comparable effort in Europe.

5 Status of European research in Systems

There was a time when European research efforts in Systems led the field. Core concepts such as operating systems, virtual memory, capabilities, local area networks, network file systems, arrays of PCs, etc., were invented in Europe. Researchers in Europe made fundamental contributions to process synchronisation, multi-processor caching, and compilers for parallel computers. But after the 1960s the centre of Systems research moved to the US. Now, despite great potential, European Systems research is under-performing, has low visibility, and has not sparked the creation of startup companies in the same manner as its US counterpart.

5.1 European strengths

European Systems researchers tend to be isolated from the main research community in North America. Paradoxically, this can be an advantage, as European researchers are not boxed in to the prevailing thinking of the US. Our research is incredibly rich and diverse. We see more out-of-the-box, original, blue-sky thinking in Europe than in US. Some recent examples of very original results (in no particular order):

- MARS: a real time operating system based on the time-triggered paradigm; had significant industrial impact on the European embedded-systems industry [9].
- The L4 micro-kernel, a different approach to high-performance operating system kernels [22].
- Xen: an original approach to virtual machine monitors that eschews full hardware virtualisation for a simpler, more efficient, elegant hardware-software emulation [21].
- Camille: an OS for embedded devices (smart cards) that leverages language and programme proof techniques to make programme code very small and highly secure [10].
- LISFS: a "logic file system," whose design is entirely based on a logic theory [18].
- Pastry: a substrate for building distributed systems, using self-administering and self-healing peer-to-peer techniques that scale to very large networks [14].
- Esterel: a language for describing the interactions between software components and the real world through synchronous events [8].

Additionally, as is well known, the World Wide Web originated at CERN in Switzerland, and the Linux [12] and Minix [15] operating systems in Finland and in the Netherlands, respectively. European Systems research is a reservoir of untapped ideas that have real industrial potential.

The funding situation is relatively favorable in Europe. There are many sources of EU and national funding. Although the system is complex, funding is currently more stable than in the US. European researchers tend to devote a smaller fraction of their time to fund-raising. Typically a European university imposes overhead rates that are much less than US universities, and faculty members at a European university are not required to raise a part of their salary through research contracts and grants.

Europe has a vibrant open source community with many contributors. This community sharpens its skills working with open source software, lessens our dependence on commercial software, and occasionally gives birth to successful commerical enterprises, e.g., SuSe, a major Linux distributor.

5.2 European weaknesses

5.2.1 Recognition

An important problem is the poor recognition in European academia and funding organisations, for Computer Science in general, and for experimental and Systems research in particular. Systems work is often considered "just hacking" and is not recognized as a field of scientific interest. The unique characteristics of Systems work are not considered in the recruitment and promotion process.

In contrast, top universities in the US such as MIT, Harvard, CMU, Stanford, and Berkeley all value their Systems researchers highly.

5.2.2 Duration of PhD research

A major hurdle is the short duration of PhDs in many European countries: typically, three years. A good Systems thesis will take five or six years in the US, sometimes more. This is what it takes to achieve the required level of maturity; then come up with a good, original idea, prove its soundness, package it as a usable abstraction, evaluate its performance, implement and measure it in a realistic context, and publish the results.[3] The three-year limit makes it impossible to complete the cycle; often the work is incomplete or unpublished when the student's contract expires.

A PhD student in the US typically engages in research work under his or her advisor's supervision during the Master's studies. This gives him a significantly longer timespan to achieve the required level of expertise and complete the research. In Europe, in general, the student chooses a dissertation topic only after the Master's. We are concerned that the Bologna Agreement encouraging student mobility, might have an (unintended) negative effect on research, by deepening the split between a Masters programme and the corresponding PhD. Instead, we need "Research Masters" programmes that start research work early and feed into a PhD programme at the same institution.

5.2.3 Status and role of PhD students

In the US, PhDs are considered students and are supported by their university via teaching assistantships, grants, or combinations thereof. The student is attached to an academic department, with a variety of mechanisms (such as "comprehensive" or "qualifying" exams) to establish a solid department-wide level of knowledge. Students are encouraged to explore outside of their narrow area of expertise and to interact with other faculty.

In contrast, in many European institutions, a PhD is considered an assistant to his advisor. This system has a number of disadvantages. Instead of encouraging diversity, breadth and exposure to other ideas, it links the individual strongly to

[3] Six years is a long time, yet Simon and Hayes have demonstrated that it takes an individual ten years of intensive training to reach a level of expertise that enables him to make significant contributions. This ten-year span holds across many areas, including science, literature and music, and across time (the last two centuries). See Simon and Hayes [7, Ch. 11] for a discussion of this phenomenon for music.

a specific supervisor. Treatment of PhDs as employees subjects them to labour rules that were developed for a different target group. Worse, it may instill the view of graduate studies as a "work for pay" employment, rather than a gift from society to further one's education.

5.2.4 Funding constraints

Variable, complex funding with an emphasis on specific deliverables makes it difficult to run long-term projects with a Systems focus. To obtain funding, groups focus on short-term objectives, at the expense of highly visible publications: the infrequency of the major conferences,[4] and the long turn-around time of top-tier journals often makes it impractical to target these venues within the timescales of a project of two or three years. Instead, groups resort to publication in small, local, or peripheral venues.

5.2.5 Industrial relations

Systems work has a natural vocation to impact the way computers are used in daily life. Hence lively interaction with the commerical sector is essential – both with the ITC industry, and with users such as financial services, manufacturers, and government or (say) the health care sector. However, Systems research must not be confused with engineering or development. The roots that scientific excellence and innovation feed upon are basic, long-term, high-risk research.

In the US, the Systems-related industry is very strong and values fundamental, publishable research, demonstrated by the sizable investments of Microsoft, Intel, IBM, ATT, Google, VMware, etc., in Systems research. These companies have a vigourous, healthy relationship with academic Systems groups, based on long-term commitments. A recent example: Google, Sun and Microsoft together committed $7.5 million to the new RAD laboratory at Berkeley [11] for open source research in distributed systems.

In Europe, the IT industry has few big players, being made up mostly of SMEs (Small and Medium-sized Enterprises), which are too small to engage in research. European IT strongholds such as gaming, embedded systems, telephony or entreprise decision support rely heavily rely heavily on the results of Systems research. Industrial sectors such as transport, packaging, machinery or banking are more and more influenced by IT. For all these sectors, access to the output of Systems research (either in the form of IT components or in the form of highly-skilled personnel) is crucial.

The European IT industry rarely does research of its own, and under-appreciates the value of fundamental research and scientific publications in Systems. For instance, although the last SOSP conference (the top venue in Systems) took place in Europe, the European software industry was conspicuously absent, whereas several US and Indian companies attended in numbers and provided substantial sponsorship. The European software industry rarely employs research PhDs.

[4] Many major venues such as SOSP [2], OSDI [24], ASPLOS [1], or HotOS [23], meet only once every 18 to 24 months.

Joint projects with academia or research contracts are not uncommon, but Systems work is usually limited to short-term development. If this trend continues, the roots of innovation will dry out.

5.2.6 Isolation

The European Systems community is small, dispersed, and out of the decision loop. The leading circles in Systems research are located in the US and this is where the consensus forms on what are the current "hot" topics and who are the star students and researchers. There are good Systems groups across Europe, but they are smaller than US groups and do not interact with other (European or American) groups as closely as the US groups. Therefore these groups don't have critical mass and do not leverage research funding as well as their overseas competitors.

6 Recommendations to the Systems community

Improvements to the situation described so far must be discussed, proposed and implemented primarily by the profession itself.

The very first step is to organize both individuals and institutions across Europe, in Systems and more largely in Computer Science, to take initiatives, to share information and statistics, and to influence decision-makers. Indeed this process has already started, for instance in the form of the present paper, the creation of EuroSys, and the recent Zürich summit.

The EuroSys professional society was established as a voluntary organisation of Systems researchers and practitioners in Europe [5]. It aims to improve the quality and impact of European Systems research, by organising the community, exchanging information, and helping set up the other activities listed hereafter.

The summit of Informatics deans and department heads was organised at ETH in Zürich in September 2005 [13]. Some of the issues discussed were improving student mobility at the Bachelors level, sharing and coordination of curricula between European universities, and evaluation of researchers in Informatics.

A related initiative is the UK "Grand Challenges Exercise" to collect challenges for Computer Science that contribute to the long-term advancement of the subject, are directed towards a revolutionary advance, and could be a basis for funding policies [20]. Their topics emerge from consensus of the general scientific community, to serve as a focus for curiosity-driven research or engineering ambition.

Our professional organisations need to take the lead in a number of initiatives, with the goal to improve the professionalism, impact and excellence of Systems faculty. Some examples include: (i) Teaching collaborations, in order to improve effectiveness in teaching. (ii) Mentoring of junior researchers by seniors, from their own university or from a different one. (iii) Senior workshops to organise the circulation of information, personnel, and best practices. (As a case in point, the present paper grew out of the EuroSys Senior Workshop in Lisbon [6].)

To be successful, these initiatives need to be encouraged by universities, employers and funding agencies by providing time and career advancement to volunteers.

The academic community also needs to appreciate the value "real-world" partners can bring to research. Take-up of the results of academic research by real users, be it other academics, the IT industry, non-profit organisations, or informal communities, Should be considered when assessing the quality of research.

7 Institutional recommendations: training researchers

We now make a number of institutional recommendations. In the present section, we focus on fostering and demanding excellence in PhDs. We want to encourage quality and breadth, and to improve their job prospects. Implementation depends on take-up by faculty, universities, academic administrations, and funding agencies.

7.1 Duration and organisation of PhD studies

As explained in Section 5.2.2, a three-year limit to finish the PhD is too short to make an impact. Systems research is experimental and labour-consuming, making it extremely difficult to achieve a sufficient degree of completeness, depth and maturity in three years. This is a major obstacle to publication in the top Systems venues, yet experience shows it is essential to build a substantial publication portfolio early in a career. Thus doctoral students in Europe do not manage to attain international recognition.

We offer two solutions. First, to make the time limit more flexible, by lifting administrative barriers and by allowing PhD grants to extend for a longer time when warranted. Second, universities should offer combined Research Masters-PhD programmes, over a typical duration of five years, extended when justified. Qualified students are matched with an advisor within the first year and begin research work as soon as possible. Students must be able to change topics or advisors.

7.2 Encouraging breadth and innovation

Mechanisms are needed to expose students to a breadth of subjects, in order to avoid over-specialisation and to encourage transversal research. One idea is the "doctoral school," bringing a diverse group of talented individuals up to speed and to get them into a position to undertake research. Small doctoral-level classes should organise projects, based on the students' proposals (with appropriate faculty feedback). This is a way to start new ideas while building critical mass, as students of the same class work together, and to instill an esprit de corps. The goal of such a class is to investigate research issues, not to "cram" pre-digested knowledge.

Regular, independent reviewing is essential to encourage quality. A thesis should be reviewed regularly by a committee of three or four researchers. This could be the advisor, another Systems faculty, a faculty member from another area, and an outside member (of PhD stature), either from another university or from industry.

7.3 Student exchanges

Mobility programmes such as Erasmus are extremely positive at the Bachelor and post-doctoral level. We must encourage Bachelors students to apply to the best institutions for their chosen research area. Indeed, even the most talented student can excel only in an environment that practices excellence in his field. This attitude is common in other areas (e.g., music or sports); in the US it is considered an honour to send an exceptional student to a top university. This is in contrast with the current wide-spread practice of keeping good students, even at an institution that cannot offer them a comparable environment.

However we reiterate our opinion (Section 5.2.2) that Systems research should start at the Research Masters level and students should remain in the same institution for their PhD, whenever possible. Instead, we are in favour of short-term mobility *during* doctoral studies. Doctoral students benefit hugely from extended visits to a different institution. They help fight the often parochial view that European students have of their research community. Exchanges allow them to experience other environments and develop contacts important for their future career.

7.4 Doctoral internships

Doctoral internships allow students to work for typically three months at an external research laboratory (often an industrial lab) on a different project. Doing this once or twice as a doctoral student has proved immensely valuable. An internship is mind-opening, as it exposes the student to a new atmosphere and new problems. Often a real world problem encountered during the internship provides inspiration for further doctoral work and publication.

The experience of an internship also improves the student's chances in securing a job later in his or her career. It makes him more experienced, shows that he is not too narrowly specialised, provides a fresh source of recommendation letters, and puts him in direct contact with potential employers.

Doctoral internships are still uncommon in Europe. Some advisors consider them an unnecessary distraction (whereas, on the contrary, an internship can provide new research and publication opportunities). They conflict with the fixed three-year duration for PhD studies. Grants generally don't know how to allow the student to be paid by another body for a short duration. Internships need to be encouraged; they should be integrated into PhD programmes and grants should accommodate them.

8 Institutional recommendations: Faculty and researchers

In this section, we give some recommendations to university administrations and funding agencies, regarding faculty and senior to encourage the development of Systems research in Europe. First we touch upon funding. Then we focus on the careers and evaluation of Systems faculty. Implementation of this section's recommendations falls mostly on

8.1 Funding vehicles

European research funding places an emphasis on broad, multi-institutional, collaborative projects with many senior investigators, large project meetings and predefined deliverables. While such projects are of considerable value, this encourages a relatively low-risk, conservative approach to ensure that the expected results can be delivered on time. Long-term or revolutionary, high-risk approaches are difficult to explore in this context.

To ensure sustained innovation in Systems, however, it is critical that the European academic research community also conduct research on long-range and/or high-risk ideas with high potential impact. To support this, the funding agencies should support single investigator or small-team grants for high-risk research. Furthermore, Systems need long-term investment; consider the example of XenSource, based on the results of many years of funded research [21]. Funding agencies should value long-running projects, along the lines of EPSRC Platform Grants for research infrastructure in the UK, and their 6-year grants for inter-disciplinary projects, which have been very successful. A 6–10 year duration allows the researchers to focus on research issues and provides them with the time to develop substantial prototypes, yet the time span is short enough to keep abstractions and prototypes grounded on reality. Of course, this time period can be divided into phases with appropriate reviews.

The driving criterion for allocating funding needs to be technical: the keywords are excellence, depth and impact. Other considerations, such as balanced funding across countries or regions, should remain secondary.

8.2 Evaluation criteria

For Systems research to be successful and to influence European companies and society, the evaluation criteria must be both fair and appropriate. As in other disciplines, the objective of research is to contribute to knowledge. However, the metrics for measuring success or impact are different in the Systems field [19]. As explained earlier (Section 3) Systems are an experimental science. Even more so than other branches of Informatics, Systems research involves substantial human investment; it involves the construction of an experimental apparatus or a research prototype, requiring a significant development effort. As a result, the rate of publication is lower than in other areas: two full-length publications in a leading venue per year is considered a good rate of publication.

Moreover, the preferred venues of publication are highly competitive conferences such as SOSP [2] or OSDI [24]. A publication in a leading conference in the Systems area is more prestigious than several in a secondary journal. The visibility of the publication venue, as well as the quality and the impact of the individual publication should be taken into account.

Another very appropriate metric of impact is the use of the researcher's artifacts by other research teams, industry or open software projects. This should be considered in the evaluation.

8.3 Junior faculty

Diversity is a key to innovation. Yet many universities in Europe still hire their own students as faculty. While it is comfortable to work with people one already knows, what research needs is challenge. To excel, European universities must collectively determine to discourage this practice.

Some countries or institutions show a distrust of their junior faculty, where they have limited rights with regard to advising PhD students, or teaching obligations that make it impossible to devote time to substantial efforts in Systems research. This set-up deprives the European research scene of productive junior researchers, and inhibits mobility at the junior faculty level.

Academic researchers in Systems must be able to devote sufficient time to research and student supervision. With a teaching load of more than 15 hours per week, including contact and preparation time, it is very difficult to perform and supervise Systems research at the highest level.

In Systems, asking the right questions, formulating an appropriate hypothesis, designing and implementing an experimental apparatus, analysing the results and drawing conclusions are all significant challenges that require experience. Therefore, Systems students require close supervision and guidance, more so than in other areas of informatics.

Junior faculty should be involved in the department, lab or institute leadership to strengthen their identification with the institution, to motivate them, and to allow them to grow into leadership roles.

8.4 Promotion

Institutions that want impact in Systems research need to compete for the best young talent on the global market.[5] To keep the best young researchers in Europe, or to attract more from abroad, we need to provide them with motivating working conditions and a career path that enables them to develop and lead.

We briefly describe hereafter one mechanism that has been highly successful in this regard: the tenure-track system. Several European institutions have adopted it, for instance ETH Zürich, EPFL Lausanne and the Max Planck Institute for Software Systems. However, more than a specific mechanism, what

[5] There is currently a window of opportunity for attracting talented Systems researchers from elsewhere, as funding and visas are harder to get in the US.

counts is the objective of encouraging and rewarding excellence in Systems research.

The tenure system provides significant milestones for promotion. In between, researchers are guaranteed stability and time to develop their research, with modest teaching, administrative and service duties. Junior faculty are fully responsible for their research programmes and are free to pursue their own vision. At the same time, they should be mentored by senior faculty members, to advise them on issues such as research strategy, publication, fundraising, personnel issues and on setting priorities, whenever they request such advice.

Since Systems research is labour intensive and has long lead times, the process should allow tenure-track researchers at least six years to establish an appropriate record. An intermediate review occurs after approximately three years, based mostly on local reviews.

External evaluations from leaders in the candidate's area of reseach is the major factor in promotion decisions. A tenure committee should solicit evaluations from peers at other institutions (academic and research laboratories) as well as from local experts. The review takes into account the candidate's outside recognition and research impact, measured by metrics appropriate for the Systems field (see Section 8.2). Other factors include service, teaching, and ability to attract funding. To attest recognition, the review should request letters of reference from at least six internationally leading experts in the candidate's area.

9 Recommendations to industry

There are strong incentives for industry to engage with Systems research. Not least of them is access to the intellectual force and creativity of the university environment: with its free flow of ideas and sharing of knowledge, academia encourages innovative thinking and cross-disciplinary research. The benefits of such interaction may appear intangible in the short term, but many significant commericial innovations have originated in academic research groups.

9.1 Expose academia to hard problems

Companies are a source of hard problems in the Systems area, with geographically large software deployments, thousands of users and often stringent demands on performance, reliability, security etc. For companies, a community in which many smart people are prepared to spend a great deal of time and energy on solving a problem is an ideal place to share intellectually challenging requirements. Conversely, the quality of academic research benefits from exposure to hard real-world problems that cannot be abstracted away. The problems can be discussed, and possible solutions can be scrutinised for flaws, in conferences, round-table discussions and trade shows. Solutions to previously ignored problems may have far-reaching impact.

Some information may be too sensitive to communicate to outsiders. In this case, commercial enterprises can still contribute to research, and benefit from valuable feedback, by disclosing (possibly in collaboration with a researcher bound by a non-disclosure agreement) a “sanitised” version. Conversely, the business community can invite research groups to exhibit their prototypes to learn about outside developments.

A common issue for researchers is lack of equipment or data from a real world deployment. Further impetus to address a particular problem can come from industry in the form of access to infrastructure, real data sets and so on.

Academic research and industry have differing interests regarding intellectual property (IP). Academics need to publish and to share knowledge freely and in a timely manner, whereas companies assert the right to protect their IP. Furthermore, an academic working on a particular problem might create new IP which he might want to exploit. However, solutions to these conficts can generally be negotiated. A satisfactory balance that protects critical IP while retaining some freedom to publish can often be found. Furthermore, there are prominent examples of commercial successes built on a technology where protected property rights are shared, for instance, the Ethernet.

9.2 Hire PhD graduates

The preconception that a PhD graduate in Systems is fit only for teaching is clearly wrong, as witnessed by the make-up of enormously successful software companies such as Google, Amazon or Microsoft. Indeed, a PhD is trained to a high level of expertise and brings skills, such as capacity for independent and abstract thought, written and verbal communication, that companies should take advantage of. PhDs were selected for their inquisitiveness and performance, have been trained to be rigourous, and have been exposed to a wide variety of new ideas.

Hiring PhDs is also beneficial in the long term, by increasing personal ties between industry and academia.

Companies should not wait until a student completes her degree before they start interacting with her. Companies can participate in mentoring programmes and provide role models for PhD students. Such role models can communicate directly what a company expects from a recent graduate entering the workforce.

9.3 Offer doctoral internships

The practice of doctoral internships is widespread in the IT industry in the USA, both in industrial research labs and in production teams. Students are usually in their second or third year of PhD studies, hence already know how to make progress independently, have expertise in their area, and some development experience. Another, high-impact example is the “summer associates” programmme of consulting companies.

Internships are beneficial both for the industry and for academics. Industry benefits from a regular influx of fresh ideas and enthusiasm. The host can try out new or tentative projects with little risk. The host can assess promising future employees, while conversely the student gains knowledge of possible future employment. Finally, internships strengthen the ties between industry and academia through personal contact.

9.4 Participate in the academic network

Sending staff to attend conferences is an important way to acquire new technical ideas, and signals that the company values the continuing education of its employees. In the Systems area, conference presentations are usually given by the doctoral students who will soon be on the job market; attending a conference is an opportunity to spot promising future employees.

Volunteer work, such as serving on programme committees or organizing committees, is another extremely important role and one where industry can make a visible and valuable contribution.

As European universities implement the Bologna Agenda, mid-career individuals will be able to return to university for advanced (post-graduate) training. The IT sector needs to work with academia to create appropriate programmes for technical leaders, similar to MBAs in prestige, breadth and depth, and impact.

9.5 Provide sabbaticals to proven performers

One way to increase the flow of information from companies to industry is to give senior personnel the opportunity to spend a sabbatical of 3 or 6 months in a research lab. While keeping critical information confidential, they will be able to describe problems in sufficient detail to inform academic research.

10 Conclusion

Systems are a fundamental component of informatics, and a healthy research community is essential for future innovation and for the success of those sectors of the European economy that are either based on IT or directly depend on it. Large swathes of the European economy will benefit. Europe has a number of excellent Systems research groups, but their impact is too low. Compared to the US, they are under-recognised academically, their publication impact is low, they do not receive sufficient recognition from national and European funding agencies, and they do not benefit from a strong industrial base. The European Systems community recognises its own shortcomings and is actively taking the initiative. This paper appeals to academic administrations, governments, funding agencies, and industries to back up these efforts. Measures to be taken include making research careers in Systems more attractive, recognising the specific metrics for impact in the System area, and providing funding for long-term, fundamental, publishable research in Systems.

References

[1] ACM Sigarch, Sigplan and Sigops. ASPLOS: Int. Conf. on Architectural Support for Prog. Languages and Operating Systems, October 2006. http://www.princeton.edu/~asplos06/.

[2] ACM Special Interest Group on Operating Systems. SOSP: ACM Symposium on Operating Systems Principles. http://www.sosp.org.

[3] M. Burrows, M. Abadi, and R. Needham. A logic of authentication. In *Symposium on Operating Systems Principles (SOSP 12)*, pages 1–13, Litchfield Park Arizona USA, December 1989. ACM.

[4] M. Burrows, M. Abadi, and R. Needham. A logic of authentication. *ACM Transactions on Computer Systems*, 8(1):18–37, February 1990.

[5] EuroSys. EuroSys, european professional society in Computer Systems. European chapter of ACM Sigops. http://www.eurosys.org/.

[6] EuroSys. Senior workshop: "Promoting excellence in European Systems research". Lisboa, Portugal, July 2005. http://www.gsd.inesc-id.pt/~pjpf/sigops-senior-workshop-2005/.

[7] J. Hayes. *The Complete Problem Solver*. Lawrence Erlbaum Associates, Hillsdale, NJ, USA, 1989.

[8] INRIA. Esterel: a synchronous reactive programming language. http://www-sop.inria.fr/esterel.org/, 2005.

[9] Institut für Technische Informatik. MARS, Real-Time Systems Group. http://www.vmars.tuwien.ac.at/.

[10] Laboratoire d'Informatique Fondamentale de Lille, RD2P Research Group. Camille NG, an exokernel for small devices. http://www.lifl.fr/RD2P/CAMILLE/Index.

[11] RAD Laboratory. RAD Lab: Reliable Adaptive Distributed Systems Laboratory, December 2005. http://radlab.cs.berkeley.edu/.

[12] Linux online. http://www.linux.org/.

[13] Bertrand Meyer and Willy Zwaenepoel. European computer science takes its fate in its own hands. *Communications of the ACM*, To appear 2006. http://se.ethz.ch/~meyer/publications/acm/ecss-full.pdf.

[14] Microsoft Research Cambridge. Pastry, a substrate for peer-to-peer applications. http://research.microsoft.com/~antr/Pastry/.

[15] The Minix-3 Operating System. http://www.minix3.org/.

[16] National Science Foundation. The GENI Initiative. http://www.nsf.gov/cise/geni/.

[17] National Science Foundation. NSF Workshop on Distributed Systems. MIT Stata Center, Cambridge, MA, USA, September 2005. http://pdos.csail.mit.edu/~kaashoek/nsf/.

[18] Yoann Padioleau and Olivier Ridoux. A logic file system. In *USENIX Annual Tech. Conf.*, pages 99–112, San Antonio, Texas, USA, June 2003. Usenix. http://www.usenix.org/events/usenix03/tech/padioleau.html.

[19] Computer Science and Telecommunications Board, editors. *Academic Careers for Experimental Computer Scientists and Engineers*. The National Academies Press, 1994. ISBN: 0-309-04931-8, Library of Congress Catalog 93-084437.

[20] UK Computing Research Committee. Grand Challenges for Computing Research. http://www.ukcrc.org.uk/grand_challenges/index.cfm.

[21] University of Cambridge Computer Laboratory. The Xen virtual machine monitor. http://www.cl.cam.ac.uk/Research/SRG/netos/xen/.

[22] University of Dresden. The L4 μ-kernel family. http://os.inf.tu-dresden.de/L4/overview.html.

[23] Usenix. HotOS: Workshop on Hot Topics in Operating Systems. http://www.usenix.org/events/bytopic/hotos.html.

[24] Usenix. OSDI: Symposium on Operating Systems Design and Implementation. http://www.usenix.org/events/bytopic/osdi.html.