

Systems research, education and industry in Europe

Marc Shapiro, Yolande Berbers, Willy Zwaenepoel, Tim Harris

SIGOPS European Chapter (EuroSys)

15 June 2005

This text is submitted to the European Commission by EuroSys. EuroSys is the European Chapter of ACM SIGOPS, the scientific society in “Computer Systems” [5].

1 The field of “Computing Systems”

The “Systems” area of computer science is the study of software platforms supporting a large spectrum of applications. It includes operating systems, distributed systems, real-time systems, systems aspects of databases, language runtimes, embedded systems, computer networks etc., and intersects a large number of application areas.

The general approach is to develop high-level yet efficient abstractions in order to provide properties such as simplicity, performance, reliability and security. For example, topics addressed at SOSPP 2003, the premier conference in the area, included [6]:

- Executing untrusted plug-ins: essential for a large number of applications.
- The Google file system: an example of removing complexity from applications by pushing new abstractions into the OS.
- Distributed authentication: obviously extremely important for security.
- Distributed debugging: a difficult area, and of increasing importance for Grids and Web Services.
- Peer-to-peer storage and Peer-to-peer streaming: increasing economical impact.
- Mobile multimedia.
- Virtual machine monitors: a technique to decrease hardware costs and increase security and dependability; a European paper.
- Preventing system crashes.

A particularly telling example of Systems principles at work is the Burrows-Abadi-Needham logic of authentication. This formal study was published first at SOSPP and in TOCS [1, 2]. It has since become the underpinning of much current research in security.

1.1 Economic importance of the area

Any successful application rests on a good system, hence the very high technical and economic importance of the area. Witness the major companies that operate in this area, for instance Akamai, Amazon, Apple, Google, Groove, IBM, Iona, Jaluna, Microsoft, Mandrake, RedHat, Sun, Suse, Symbian or VMWare.

Many of these are directly descended from recent academic research, e.g., VMWare from Stanford and Akamai from MIT. However, despite occasional notable exceptions¹, it is obvious from this list that Europe is largely missing out.

Systems also play a vital role in the telecommunications and embedded systems, a field where Europe is competitive. Examples are Siemens, Infineon, ST-Microelectronics, Nokia, Airbus, etc. Another area is the open-software community, which was largely kick-started by Linux. Europeans are very active in this community, even though it is hard to measure their impact with any precision.

1.2 The Systems community

Like many scientific areas, Systems work is best defined, not by semi-arbitrary technical boundaries, but by the community. The community defines itself around the major conferences (SOSP and OSDI, events such as HotOS and the SIGOPS European Workshop), and its scientific societies (ACM SIGOPS, Usenix and IEEE TCOS). As we shall see later the political centre of gravity is in the USA, around an informal but closely-knit network.

The top conference and journals such as SOSP, OSDI and TOCS are extremely selective and have very high standards. Systems constitute an area of intense scientific study, with a particular mix of theory and experimentation: pure performance without principles is not of interest, but conversely neither are ideas that appear elegant but cannot be implemented efficiently. What is publishable are the principles that underly successful engineering yielding measurable end-to-end performance, reliability and security.

A crude but compelling measure of the academic and economic impact of a publication venue is its position on the CiteSeer impact list [3], which lists 1221 venues. The top ones are (in descending order): *OSDI*, USITS, *PLDI*, SIGCOMM, MOBICOM, ASPLOS, *USENIX Annual Tech Conf.*, *TOCS*, *SIGGRAPH*, *JAIR* and *SOSP*. *Eight* of the top eleven are systems-related: either pure Systems venues (in italics) or having strong Systems component (underlined). The premier systems venues, OSDI (Operating Systems Design and Implementation), TOCS (ACM Transactions on Computer Systems) and SOSP (Symposium on Operating Systems Principles) are *first*, *eighth* and *eleventh* respectively.

As we will see later (Section 2) the European publication rate in these venues is disappointing. The focus of events like SOSP/OSDI is narrow but unpredictable. The explanation is that the centre of gravity of the community is in the USA. There is an informal but very active network of visits, exchanges, NSF

¹Jaluna and Iona are older descendents of research (INRIA and Trinity College Dublin respectively).

workshops, DARPA meetings, coffee-break encounters at numerous conferences that Europeans rarely have the chance to attend. Unfortunately, this means that the current “hot topics” are defined by informal consensus, and Europeans are outside of this crucial decision loop.

The main scientific society for systems work, ACM SIGOPS, has been proactive for a number of years in supporting European research. For instance, a biennial SIGOPS European Workshop has been held since 1985, and the SIGOPS vice-chair is traditionally a European. SIGOPS is organising SOSP 2005 in Europe for the second time, in Brighton, 23 Oct (attendance of 250–350 expected). Unfortunately the first time (St Malo 1997) did little to increase the impact of European systems research or increase our self-confidence.

2 Status of European research in Computer Systems

The concept of an “operating system” was invented at Cambridge University in the early days of computers. Vital concepts such as virtual memory, capabilities, local area networks, network file systems, arrays of PCs, etc., were invented in Europe. But after the 60s the focus of the computer industry moved to the USA. Now, despite a great potential, European systems research is underperforming and has low visibility. We conclude to a number of necessary improvements, which we have started implementing.

2.1 European strengths

Paradoxically, Europe’s somewhat isolated status can be an advantage, as European researchers are not boxed in to the prevailing thinking as in the USA. 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:

- MARS: a real-time operating system based on the time-triggered paradigm; had a high industrial impact on the European embedded-systems industry.
- The L4 micro-kernel, a different approach to high-performance operating system kernels.
- Xen: an original angle to virtual machine monitors that eschews full hardware virtualisation for a simpler, more efficient, elegant hardware-software emulation.
- Camille: an OS for embedded devices (smart cards) that leverages language and program proof techniques to make program code very small and very secure.
- LFS: a file system designed from the top down, starting from a logic theory.
- Pastry: a novel substrate for building scalable distributed systems

Of course, the World-Wide Web has its origins in Switzerland and the Linux and Minix operating systems in Finland and the Netherlands.

Appendix B provides a sample of some of the institutions involved in top Systems research in Europe.

There is a reservoir of untapped ideas, that have real industrial potential. Although complex, there are many sources of European and national funding. Finally, Europe has a vibrant open-source community.

2.2 European weaknesses

An important problem is the poor recognition in European academia and funding organisations, for computer science in general, and for experimental/systems research in particular. Systems work is often considered “just hacking” and is not recognised as a field of scientific interest. The unique characteristics of systems work is not taken into account in recruitment and promotion.

Contrast with USA’s top universities, such as MIT, Harvard, CMU, Stanford, or Berkeley, which all value their systems researchers highly. As a result, European researchers leave for the US, or do other things (e.g., theory), or publish peripheral results.

Another 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. This is what it takes to come up with a good, original idea, to prove its soundness, to package it as a usable abstraction, to evaluate its performance, to implement and measure it in the context of real applications, while also publishing. The three-year limit makes it impossible to complete the cycle; often the system is incomplete and/or unpublished when the student leaves.

Variable, complex funding 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 good publication: the infrequency of the major conferences,² 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 2 or 3 years. Instead, groups resort to publication in small, local or peripheral venues.

Finally, our community is small, dispersed, and out of the decision loop. There are many good systems groups across Europe but they do not interact as closely as the US groups.

The result is clear in Figure 1: a dramatic underperformance of European publications in the main venues of the area. We are out of the decision loop, both in industry and academic. This undermines Europe’s capability of technological independence.

²The major venues, SOSP, OSDI, HotOS, etc., meet only once every two years.




Figure 1: Relative number of papers from Europe at SOSP and OSDI and in TOCS

3 EuroSys: A European scientific society and conference in Systems

Clearly, something needs to be done.

Our small contribution is to establish a European scientific society in Computer Systems and to create recognised European events in the field. The goal is to exchange and coordinate, and to create a true European community in Systems, while at the same time having impact within the larger international Systems community.

Obviously, this is only a first step. The development of a strong Systems community in Europe, and its industrial and academic impact, depends a lot on support by European, national and university administrations. In Appendix A we point to the example of what is being done at EPFL.

3.1 EuroSys: European chapter of ACM SIGOPS

Since ACM SIGOPS is the recognised scientific society, and as an amplification of their previous efforts towards Europe, we are establishing a European chapter of SIGOPS. This has the full support of the SIGOPS officers.

The scope of the Chapter is to support Systems research and development carried out in Europe, to provide a European perspective on research and development results on the design, specification, realization, behavior, and use of Computer Systems. It will be operated exclusively for educational and scientific purposes, to promote increased knowledge of and greater interest in the science, design, and development of modern Computer Systems, and a means of communication between persons having an interest in Systems.

We aim to include in our membership the “hard core” of senior Systems researchers in Europe who are recognised through their publication, stature or seniority. We estimate this represents 30 to 50 people.

Our flagship event is a yearly conference (see Section 3.2). Other events include a senior workshop, a doctoral workshop and presence at SOSP 2005 (see Section 3.3). An election is planned for October 2005, to replace the current interim officers.

3.2 EuroSys conference

We need to create our own community, which sets its own priorities but is well linked to international research. Therefore we are establishing a new conference, which aims for SIGOPS’ standard of excellence and impact, but differs in character from SOSP/OSDI. It is an annual event, with a broad scope, located in Europe, yet international.

The EuroSys Call for Papers [4] explicitly encourages submissions on the full spectrum of Systems topics. It targets topics that are of special interest to Systems researchers in Europe or are under-represented in the traditional systems conferences. A session of the conference is reserved for early ideas; another to experience papers, reflecting on the use or the lack of use of earlier concepts, including negative results.

This structure aims to ensure that a diversity of topics emerge rather the narrow focus that the SOSP/OSDI programmes sometimes take. The CFP targets all areas of operating systems and distributed systems, as well as Systems aspects of:

- Programming language support,
- Parallel and concurrent computing,
- Mobile and pervasive computing,
- Embedded computers and tiny devices,
- Novel uses of information technology,
- Dependable computing,
- Management, measurement, monitoring,
- Databases,
- Distributed algorithms,
- Middleware,
- Clusters and grids,
- Novel user interfaces,
- Real-time computing,
- and
- Security.

The first EuroSys Conference takes place in April 2006 in Leuven, Belgium. We expect an attendance of 150. The Program Committee is truly international: six are from US, seven from Europe, one from Israel, one from India.

3.3 Other events

Other activities planned for the first year include the following:

- A Senior Workshop: “Promoting excellence in European Systems Research.” Lisbon, 12-13 July 2005. Attendance of 40 expected. Senior scientists present their directions. Keynote by Werner Vogels, Amazon. Organized by Ant Rowstron, UK, and Paulo Ferreira, Portugal. The list of invited institutions is provided in Appendix B.
- European Doctoral Systems Symposium, Brighton, 23 Oct 2005. 30 Ph.D. students present their work. Session on “transferable skills” e.g., presentation skills. Supported by Symbian. Organized by Tim Harris, UK.

- Presence at SOSP 2005, Brighton (UK), 24-27 October. We hope that, thanks to the chapter's presence, this next SOSP in Europe will do more to promote European research than the previous experiment (St. Malo, 1997).

4 Conclusion

Computer Systems are a cross-cutting discipline of high importance: mastering Systems is essential for building successful applications. Systems are strategic. Systems have real economic and practical impact, both in the international research community and in commercial exploitation. Europe has high potential. Computer Systems constitute an important area of basic research.

However, Europe has been much less more successful than the US at both industrial and academic exploitation. Something must be done. Our own contribution is to create a network of European systems professionals, in order to structure the community and to help improve its level of excellence. Beyond this, academic institutions need to take the teaching and research in systems more seriously, and to take its special characteristics into account. National and European funding organisations must give Systems a high priority, commensurate with the world-wide economic impact of Computer Systems, in order to catch up with the US.

In order to realise the high potential, FP7 should consider Systems as a specific thematic priority, and should emphasise the systems component. Systems researchers are already hard at work on exciting, novel systems architectures to solve industry's important challenges such as, for instance:

- Security.
- The storage, management and availability of large amounts of information.
- Mobile and ubiquitous computing.
- Self-managing/self-adapting/autonomous collections of Computers and applications.
- Embedded systems.
- Computer architecture/system/language co-design.
- Peer-to-peer and sensor networks.
- The exploitation of massive networks, computations and data sets.
- Fault-tolerant and latency-tolerant computing.
- Grid computing.
- Languages and tools for the design and management of massive distributed applications.

This is just a preliminary outline; we look forward to cooperation between our community and the EC officers in the preparation of FP7 and in the preparation and evaluation of specific calls or projects.

References

- [1] M. Burrows, M. Abadi, and R. Needham. A logic of authentication. In *Proceedings of the 12th ACM Symposium on Operating Systems Principles*, pages 1–13, Litchfield Park Arizona USA, dec 1989. ACM.
- [2] M. Burrows, M. Abadi, and R. Needham. A logic of authentication. *ACM Transactions on Computer Systems*, 8(1):18–37, February 1990.
- [3] CiteSeer. Estimated impact of publication venues in Computer Science. <http://citeseer.ist.psu.edu/impact.html>.
- [4] EuroSys. EuroSys 2006 call for papers. <http://www.cs.kuleuven.ac.be/conference/EuroSys2006/callforpapers.html>.
- [5] EuroSys. EuroSys, the European professional society for Systems – European chapter of SIGOPS. <http://www.eurosys.org/>.
- [6] Larry Peterson, editor. *19th Symposium on Operating Systems Principles*, Lake George, NY, USA, October 2003. ACM Sigops, ACM. <http://www.cs.rochester.edu/sosp2003/papers.shtml>.

Appendices

A Moving forward: The example of EPFL

The EPFL School of Computer and Communication Sciences wants to be in a position to be competitive with the top ten US schools in Computer Science. One particular focus area is Computer Systems. A first step is to make Computer Science, especially experimental CS, a discipline recognized by our academic colleagues with equal status to other disciplines, say Physics or Biology. A tenure-track program has been established, in order to encourage young researchers to do good work. The school has a proactive program to recruit top-notch people, towards both Faculty and Students. Salaries and teaching load are competitive with US. The results are some excellent faculty recruits:

- Monika Henzinger (Google): Web Algorithms
- Giovanni de Micheli (Stanford): System design tools
- Tom Henzinger (UC Berkeley): Reliable software

B Institutions represented at Senior Workshop, Lisbon, July 2005

- Cambridge University, United Kingdom.
- École Polytechnique Fédérale de Lausanne (EPFL), Switzerland.
- Eidgenössische Technische Hochschule Zürich (ETHZ), Switzerland.
- École des Mines de Nantes, France.

- Institut National de Recherche en Informatique et Automatique (INRIA), France.
- Instituto de Engenharia de Sistemas e Computadores (INESC), Portugal.
- Intel Research, United Kingdom.
- Katholieke Universiteit (KU) Leuven, Belgium.
- Laboratoire d'informatique de Paris 6 (LIP6), France.
- Laboratoire d'Informatique Fondamentale de Lille (LIFL), France.
- Microsoft Research, United Kingdom.
- Technical University Dresden, Germany.
- Università della Svizzera Italiana (USI), Lugano, Switzerland.
- University College London, United Kingdom.
- University of Bologna, Italy.
- University of Erlangen, Germany.
- University of Glasgow, United Kingdom.
- University of Karlsruhe, Germany.
- University of Lisboa, Portugal.
- University of Troms, Norway.
- Vrije Universiteit, the Netherlands.