

Distributed Wisdom: What's Wrong with Us?

Marc Shapiro, *INRIA Rocquencourt and LIP6*

As an active researcher in distributed systems since 1980, I've seen a lot of good work in Europe—interesting, useful, provocative, novel, out of the box, rigorous, and well thought out. Yet it seems to have little impact: it doesn't make major publications, our colleagues elsewhere haven't heard of it, and industry ignores it.

Well, let's do something! A number of us have created a professional society for the European systems community, EuroSys (<http://www.eurosys.org>), which set up a white-paper committee to analyze and make proposals. This article derives from the committee's discussions, filtered through my own perception (and my biases as a French academic who has worked in an industrial lab in England). In the interest of brevity, I'll skip the discussion and jump straight to my conclusions; I encourage you to read the full paper, which is balanced and comprehensive.¹

European weaknesses

When it comes to publication, we're way behind our US colleagues. Look at **figure 1** (collated by John Wilkes and Yolande Berbers), which compares the number of European articles published at the major systems venues to the total articles published there, and make your own conclusions. Perhaps you publish elsewhere? But the Symposium on Operating Systems Design & Implementation, *ACM Transactions on Computer Systems*, and the ACM Symposium on Operating Systems Principles rank 1st, 8th, and 11th on the CiteSeer impact list (<http://citeseer.ist.psu.edu/impact.html>), whereas the ACM Sigops European Workshop is 216th, the International Conference on Distributed Computing Systems is 217th, the International Symposium on Fault-Tolerant Computing is 270th, and the European Conference on Parallel Processing is 491st. (The International Conference on Dependable

Systems and Networks isn't even listed.) Of course, publication isn't the only metric—but where are the Akamais, the Googles, and the Microsofts of Europe?




Figure 1. Comparing the number of papers published by European authors and by other authors at three major systems venues: the Symposium on Operating Systems Design & Implementation, ACM Transactions on Computer Systems, and the ACM Symposium on Operating Systems Principles.

Part of the problem also relates to the administrative constraints on European PhD programs. You've probably seen the pattern many times: a student rushes to finish a PhD and barely manages to submit it before his or her stipend runs out. The paper gets rejected because it doesn't have a clear focus, the experiments are unconvincing, and it doesn't advance the state of the art. Does this sound too familiar? A systems PhD is a major undertaking, and it's incomplete without at least some theory, substantial software development, and lots of measurements in realistic conditions. This is a lot to ask for, yet in many places, after three

years, you're done. Contrast that with a US PhD, which takes five years to complete on average.

Also, I don't know whether it's the same for you, but I find that attracting good people to my research facility is increasingly difficult. Interns prefer industry, a French PhD grant is unattractive, and junior faculty barely earn enough to pay the rent. Career prospects aren't much better, and faculty are overburdened with teaching and administrative duties. The odds of attracting a first-class researcher from elsewhere are practically nil. Looking at the next level up, we see that Europe underinvests in fundamental research in systems. To compete, our academic institutions must encourage and reward excellence, but this conflicts with their egalitarian culture. Our groups are still too isolated from one another, from the core US community, and from industry.

Some of these problems aren't under our control; they depend on university, national, or European politics. Still, we can improve things, both by working within the system and by changing it.

Working within the system

Our first job is to help our PhD students do better research, make an impact, and get good jobs. To beat the three-year limit, the white paper recommends starting top students on research earlier, at the master's level, and keeping them. Encourage your students to publish, early and often. Help them broaden their interests: expose them to new ideas, send them to seminars outside their topics, and have them travel and collaborate.

An incredibly successful device, which we don't use enough, is the doctoral internship. Students spend three months elsewhere, often in an industrial lab, working on a new project, with new people. Far from being a waste of time, this boosts students' creativity and morale—and helps them build a CV. (The EuroSys Web site has a special area devoted to internships, <http://www.eurosys.org/interns>.) When they graduate, send them off on a post-doc. (Yes, there's also a ">page for job and post-doc offers, <http://www.eurosys.org/jobs>.)

Do your colleagues and your hierarchy understand systems? Systems aren't "just hacking." Systems are a scientific area with its own methods and evaluation criteria. Explain why working code is essential and counts as a scientific result, just like a paper; explain the importance of rigorous measurement and evaluation; explain why we publish in conferences rather than journals; and explain why we publish less than researchers in other areas do. Encourage diversity: join hiring committees and make sure your institution always prefers

outside candidates. Generate excitement: the white paper suggests a number of challenging applications—for instance, pervasive democracy, a large-scale medical information system, communication and information management in a disaster area, global-change monitoring, eternal personal and shared memory, or a global federation of libraries. These applications underscore the hard scientific challenges of distributed-systems research: security, reliability, scale and diversity, data in the network, management complexity, experimental facilities, and testbeds and simulations, to list a few.

Changing the system

A more long-term task is to work for fundamental changes to our own universities and to funding agencies. At a recent EuroSys workshop in Lisbon (<http://www.gsd.inesc-id.pt/~pjp/sigops-senior-workshop-2005>), the participants generally agreed that our community's senior members must speak out and reach the decision makers. Up to now, EU funding has targeted large, multi-institutional projects with predefined objectives. While this has boosted European research immensely, such projects are necessarily conservative. We need more funding to come up with bold, innovative ideas and to stay ahead as our area changes so quickly.

Furthermore, funding hasn't really focused on systems challenges. We need more diverse funding: for small-team, innovative, high-risk research; for basic, sustained research attacking the deep scientific challenges in systems; for long-term infrastructure projects to build experimental facilities; for technology transfer funding; and so on. (And PhD funding should allow internships and extend beyond three years.)

Careers must become more attractive. A few European institutions, such as EPFL (École Polytechnique Fédérale de Lausanne, Switzerland), ETH (the Swiss Federal Institute of Technology), and MPI-SWS (the Max Planck Institute for Software Systems, in Germany), are competing with major US universities on their own terms. They have managed to hire several high-profile academics from the US. To be sure, we can't expect every university to compete at the same level. But if an institution wants to belong to the top tier, it must have the tools to attract, motivate, encourage, and promote the best researchers.

Conclusion

Systems are a fundamental component of informatics, and a healthy research community is essential for future innovation. Many sectors of the European economy will benefit. Systems

research in Europe faces many challenges on the way to becoming competitive internationally. Researchers have already been helping to change the system from within by creating the EuroSys professional society, holding workshops, launching the EuroSys conference series, networking and exchanging information, and so on. You too can contribute, by generating excitement for systems projects and helping your students gain breadth and depth. EuroSys has started working on changing the system itself, with the white paper and by holding policy discussions with European Commission officers. Again, we need you too to carry on this fight, with decision makers in your university, nationally, and in Brussels.

Reference

1. P. Druschel et al., "Fostering Systems Research in Europe," <http://www.eurosys.org/whitepaper-2006>, white paper, EuroSys, Apr. 2006.



Marc Shapiro is a researcher at INRIA, the French National Institute for Computer Science and Control Science, and is part of Regal, a joint group between INRIA and LIP6 (the Computer Science laboratory of Université Paris-VI Pierre-et-Marie-Curie). Contact him via www-sor.inria.fr/~shapiro.

Related Links

- DS Online's Operating Systems Community, <cms:/dsonline/topics/os/index.xml>
- "Fundamental Research Challenges in Real-Time Distributed Computing", <http://doi.ieeecomputersociety.org/10.1109/FTDCS.2004.1316586>

Cite this article: Marc Shapiro, "What's Wrong with Us?" *IEEE Distributed Systems Online*, vol. 7, no. 5, 2006, art. no. 0506-o5003.