

UC Davis

UC Davis Previously Published Works

Title

The IEEE Symposium on Security and Privacy, in Retrospect

Permalink

<https://escholarship.org/uc/item/80k0f33m>

Journal

IEEE Security & Privacy, 12(3)

Authors

Neumann, Peter G.

Peisert, Sean

Schaefer, Marv

Publication Date

2014-05-01

The IEEE Symposium on Security and Privacy, in Retrospect

Peter G. Neumann | SRI International

Sean Peisert | Lawrence Berkeley National Laboratory and University of California, Davis

Marvin Schaefer

Tracing the history of computer security and privacy is a mammoth undertaking, somewhat resembling efforts to combine archaeology and ethnology with a compendium of past and foreseen risks—and how different courses of history might have affected those risks in different ways. (For example, the University of Minnesota’s NSF-funded collection of oral histories from influential people in this area is a wonderful effort to capture some this information; https://wiki.umn.edu/CBI_ComputerSecurity/WebHome.)

Tracing the history of the IEEE Symposium on Security and Privacy (SSP), the longest-running computer security research meeting, is considerably easier—and quite relevant to the somewhat shorter history of *IEEE Security & Privacy* magazine. Indeed, a previous article written for the proceedings of the 31st SSP did exactly that,¹ so it seems unnecessary to duplicate it here.

Instead, we focus more on SSP’s evolution and its vital relevance to the research and development communities along its path from a community gathering to premier security research meeting. We highlight some of the technological and engineering paradigms that SSP either stimulated or were reflected in intense discussions that ensued, and also to some extent SSP’s potential impact on the world at large.

Early Days

SSP began in 1980 as the result of Stan Ames and George Davida wanting to hold a meeting with a few practitioners and others interested in security and privacy. That

first gathering attracted 50 people who were all seriously involved in the field in one way or another. It was more like the traditional notion of a workshop, rather than the modern ACM/IEEE/Usenix notion of a workshop as a small conference. Initially with invited papers and panels, this informal setting morphed into calls for papers and then into active discussions of beliefs, apparent progress, and known open problems and challenges. There were few distractions in SSP’s early years at the Claremont Resort (whose front door is in Oakland and back door in Berkeley). Over 31 years, SSP grew in depth, breadth, and organizational structure, with a mix of practical and academic participants, papers, panels, and occasional invited talks. In 2012, with the number of attendees having outgrown the Claremont fire laws, the symposium moved to San Francisco, with more than 450 people attending in 2013, despite restricted travel budgets and related factors. With attendance approaching 500, the symposium outgrew even the St. Francis in San Francisco. Now, it’ll be held in San Jose, California—at least, in 2014 and 2015.

SSP’s early participants genuinely thought they were on track to find solutions to the computer security problem—until reality and justifiable cynicism entered the picture. When worked examples began to be available for study, recognition of the costs of security (efficiency, features, and sufficiency), and “new” discoveries (Shannon, Turing, Dijkstra, and Hoare) deepened the recognition that applications and experimental trends were just as important as theoretical research.

Cryptography was an integral part of the first three SSPs, but perhaps inspired the creation of the International Association for Cryptologic Research (IACR) and its annual Crypto research conference in Santa Barbara. The deep theoretical research in cryptography remains at the Crypto conference, but the field has recently found its way back into SSP, largely through its applications.

There was initially a strong partition between system security and cryptography at SSP. Some of the systems crowd talked about undecidability and halting problems, whereas some of the cryptographers discussed highly theoretical theorems and proofs. After a cryptography session at the 1981 SSP, all the panelists left the room for private discussions and missed the subsequent operating system and formal methods panel. Somewhat in response to that situation, SSP 1982 had a panel that approached operating systems, cryptography, and formal methods in a single session. (Full disclosure: one of the authors of this article was the PC chair that year.)

Historical Insights

In retrospect, one of the most important insights learned from the early meetings might be that the need for holistic thinking must be internalized by each new generation of researchers. Point solutions often tend to ignore the reality that security and privacy are total-system issues, involving not just hardware and software but people and operational environments as well.

In the spirit of holistic thinking, consider the quote that Butler Lampson and Roger Needham attributed to each other and that others seem to attribute to Jim Morris: “If you think that cryptography is the answer to your problem, then you do not understand cryptography and you don’t understand your problem.” Dorothy Denning made a related comment in her 1999 National Computer Security Conference (NCSC) Award acceptance speech, when she noted that “security models and formal methods do not establish security. Systems are hacked outside the models’ assumptions. ... Provable security, even if it were achievable, is not a panacea.” Obvious generalizations of these quotes apply to computer-communication security more broadly.

Recurring Issues

In the wake of SSP’s early workshop-ish beginnings, research results have still remained a primary focus. Specific topics have changed continually over time, to reflect changing threats and needs—for example, much more attention is now paid to “systems in the large,” networks, and various applications, a natural consequence of the fact that isolated computer systems no longer dominate. Particularly relevant here is Bob Morris’s statement before the National Research Council’s Computer Science and Technology Board on 19

What’s in a Name?

We note the serendipitous foresight of Stan Ames and George Davida in calling SSP a symposium and not a conference. In 2013, the US Navy and possibly other government organizations established a policy that personnel are no longer allowed to attend any meeting that has “conference” in its name—based on the potentially erroneous supposition that conferences are inherently boondoggles rather than technical meetings. This could well be overkill, inflexible, and not very fair, but it bodes well for SSP.

September 1988: “To a first approximation, every computer in the world is connected with every other computer.” Although K Speierman (then chief scientist of the NSA) and Peter Neumann were on the same panel and echoed Bob’s remarks, many others in the US government seemingly still believed that isolated, secure enclaves could exist, ignoring sneakernets, insider misuse, incomplete deletions, and dependence on untrustworthy third-party software. SSP was certainly a leader in recognizing some of these issues.

There is also an increasing focus at SSP on research into system vulnerabilities and specific attacks on computer-based systems, including electronic voting systems, medical devices, online payment systems, and automobiles. SSP maintains its role as one of the primary meetings for discussing new research related to improving and understanding security and privacy.

Another recurring issue involves how and when to discuss vulnerabilities discovered in real systems, which has led to some contentions about whether or not to divulge serious exploitable flaws. The SSP research community seems to balance this issue fairly carefully, but typically comes down on the side of openness rather than security by obscurity. It also tends to focus on underlying principles and deeper technical issues, as well as noting subtleties such as that the semantics of observed data can have multiple covert meanings, rather than just short-term remedies, such as patching.

Growing Up and Out

As SSP grew, there were concerns from paper submitters, would-be attendees, and committee members; for example, about the limits on papers and attendees, as well as occasional lack of impartiality in the reviewing process—which then led to blind reviewing and more difficult decisions on paper acceptances. At about the same time, a move from featuring mostly academic research to a more balanced inclusion of practical and pragmatic papers took place. Standards for paper acceptance gradually became much more rigorous, although

controversy was usually sought and celebrated. While other major security conferences have started using multiple, parallel sessions, SSP's "single-track" nature along with ample time for informal interactions have been persistently manifest.

Throughout its 35-year history, SSP has—perhaps without intending it—also become arguably *the* flagship meeting place for research in security and privacy. While ACM's Conference on Computer and Communications Security (CCS), the Internet Society's Network and Distributed System Security Symposium (NDSS), and the USENIX Security Symposium are also considered among the "big four" meetings, SSP seems to have unusually high sway with tenure committees (and other judges of academic prestige), along with other top systems-oriented conferences, such as the ACM Symposium on Operating Systems Principles (SOSP). Why this is the case is an open question, of course, but it could be a combination of factors, including the breadth of programs, the strong combination of both theory and practice, the connection to a major professional society, and perhaps its longevity—but also the willingness of R&D leaders to volunteer their expertise.

SSP's single-track nature and the continued increase in the number of papers submitted have, for better or worse, limited the number of papers accepted more than for other so-called Tier-1 security conferences. As such, while other conferences have let their acceptance rates drift up to between 15 and 20 percent in recent years, SSP's has been well below 15 percent for many years now—even high single digits. Although extending SSP to three full days while reducing the length of talks slightly has enabled more papers to be accepted, unless the conference expands to parallel sessions, the single-track nature will continue to keep acceptance rates low because of the dramatic increase in submitted papers. In an ideal world, program chairs would like to accept all top-quality papers submitted and reject the others. With acceptance rates of between 10 to 15 percent, there's some question as to whether the odds of a top-quality paper being accepted come down to a crap shoot or not—that is, acceptance is a sign of quality, but rejection doesn't mean a paper is bad. How this plays out for SSP in coming years remains to be seen.

The increased number of papers submitted and the low selection rate have also had an effect on the reviewing process. As the number of papers submitted has increased the workload, the symposium's selectivity has raised the stakes for choosing the right ones. Now, rather than a single round of reviews and a discussion among the program committee, there are three rounds, where the number of "live" (not yet rejected) papers progressively decreases. Only then is an in-person meeting of the program committee held and final decisions

made about acceptances. This multi-tier reviewing process has made the task of reviewing papers at least somewhat more tractable—but whether it stays tractable remains to be seen. Already, other venues, such as the the Workshop on Learning from Authoritative Security Experiment Results, have started using the concept of structured abstracts—a very short piece identifying the elements thought by the program committee to be most indicative of a paper's suitability for acceptance—as a first-stage filtering process before encouraging a paper to be submitted. Perhaps SSP will find that the sheer number of papers submitted and pressure to accept the right ones to maintain top quality will force it to innovate further in the future as well.

We've noted that due to desires to increase the number of attendees, SSP (euphemistically still called the *Oakland Conference*, although both words are inappropriate, for different reasons—see the sidebar) is moving again this year. Some people argue that SSP has become too competitive, in that it currently has the lowest acceptance rate of all comparable meetings and sometimes rejects potentially worthy papers. Others might argue that its increasing popularity can threaten intimacy, open discussions, single-track sessions, and collegiality. Thus far, this isn't a serious problem with SSP, but watching (for example) the former National Computer Security Conference and the RSA Security Conference morph into trade shows suggests that largeness could be a future risk, at least for a research meeting. For SSP to maintain its relevance and historical significance, its organizers will presumably want to maintain its unique characteristics. ■

Reference

1. P.G. Neumann M. Bishop, S. Peisert, and M. Schaefer, "Reflections on the 30th Anniversary of the IEEE Symposium on Security and Privacy," *Proc. 31st IEEE Symposium on Security and Privacy*, 2010, pp. 3–13.

Peter G. Neumann is Senior Principal Scientist in the SRI International Computer Science Lab and the moderator of the ACM Risks Forum. Contact him at Neumann@csl.sri.com.

Sean Peisert is jointly appointed as a staff scientist at Lawrence Berkeley National Laboratory and as an assistant adjunct professor at the University of California, Davis. Contact him at peisert@cs.ucdavis.edu.

Marv Schaefer is retired after many fruitful years in computer security research. Contact him at bwapast@verizon.net.