New Security Paradigms Workshop (NSPW 2010)

September 21–23, 2010
Concord, MA

Summarized by: Matt Bishop, University of California at Davis; Steven Greenwald, Independent Consultant; Michael Locasto, University of Calgary

The 2010 New Security Paradigms Workshop began with a reception and dinner on September 20 and ended at noon on September 23. The highly interactive workshop has its participation limited to about 35 people. NSPW encourages authors "to present ideas that might be considered risky in some other forum, and all participants are charged with providing feedback in a constructive manner. The resulting intensive brainstorming has proven to be an excellent medium for furthering the development of these ideas." For more on NSPW, check out http://www.nspw.org.

Disclaimer: due to severe time constraints, the presenters have not checked the accuracy of these summaries. We take full responsibility (but not the blame) for all mistakes.

PANEL

- ***Why Is There No Science in Cyber Science? (first part)***
  Panelists: Roy Maxion, Tom Longstaff, and John McHugh
  Moderator: Carrie Gates

The theme of this panel was the relationship of science with computer security. All three panelists believed that computer security would benefit from a good dose of scientific rigor.

John McHugh began by polling the audience to see how many believed scientific methodology should be applied to experimental computer security research. Two thought the claim was questionable; everyone else agreed with it. McHugh pointed out that injecting this rigor requires clear statements of hypotheses (which is easy), and then collecting data and performing data analysis in a way that can be reproduced (which is hard). In the discussion that followed, someone pointed out that perhaps the classical model of experimentation in physics is a poor analogy to experimental computer security, because the universe is benign and will not lie to you. Several people suggested that a better model is anthropology, noting that field workers who collect the data and lab analysts who analyze and draw conclusions from those data are completely different. This is like the culture of computer security.

Roy Maxion followed. He pointed out that making computer security research more rigorous could be done incrementally. To emphasize the importance of describing the data collection methodology, he cited a study on using mouse movements to identify users. The researchers built a special apparatus and had 11 people browse. The data they collected enabled them to identify each person. But they did not tell people *what* to browse. That the subjects were browsing different Web sites might have produced the discrepancies that enabled the researchers to uniquely identify each person. In the discussion that followed, someone pointed out that science is not necessarily hypothesis testing, but is really contributing to generalized knowledge—and this means that papers that describe attacks are usually not scientific, as they are not generalizable. Maxion agreed, commenting that problems often arise when one does an experiment using a small sample size and tries to generalize the results to a large population. In response to another comment that the problem is the lack of an argued methodology for phenomena we wish to investigate, Maxion recommended that researchers have a method for doing something, and the method be made transparent by including in papers an experimental methodology section detailing exactly how the thing was done.

Tom Longstaff went last. He discussed the culture of publication, lamenting that there are now so many venues for publication that conferences and workshops often accept weak papers because they need to fill sessions. Further, program committees often take papers that are not good science, and instead of rejecting them on those grounds, try to fix them. He argued that we need to reject such papers outright. He concluded with a challenge to the workshop. He noted that papers in the workshop fall into two categories: speculative papers putting across new ideas, and papers with conclusions and results. He asked whether the former provided ideas that could be rigorously evaluated so that a good scientific paper could be produced, and whether the latter had scientifically sound methodologies and conclusions. In the ensuing discussion, one participant suggested focusing not on "science" but on "theory building," arguing that the methods used in physics are different from those in social science, but theory building applies to all disciplines in that it requires identification of premises and rules for drawing conclusions (whether inductively or deductively). The response was that the panel was discussing experimental science in the way it handles and manipulates data being collected, and so can resemble social science as well as physics. There was also considerable discussion about the relationship between engineering and science, with a comment that security as a discipline does not know what it is trying to do—we want to make things "secure" but do not know how. The response was that it's a bit like working on a perpetual motion machine; we can't get there, but we can continually approach the goal, and we can measure how far we fall short of it. People expressed hope that security could measure how far it falls short.

All three panelists emphasized the importance of including a section on methodology, especially experimental methodology, in all papers so that reviewers and readers can assess the results properly.

- ***E unibus pluram: Massive Scale Software Diversity as a Defense Mechanism***
  Michael Franz

Michael Franz reviewed the process of software creation and distribution. The current practice is to use a "unicompiler" that produces a single object code, so all instances are susceptible to the same attack. He proposed a "multicompiler," which gives different object files to different customers; the same attack would not work on all object files. He proposed doing this diversification in an "app store" to simplify the distribution process, hiding it from both the developer and the user so that their processes need not change. This solves the problem of attackers reverse-engineering patches to find vulnerabilities, because the store would either need to send each customer a custom patch for their version of the application or simply send a new, patched version of the application. He discussed four paradigm shifts that underlie this work: online software delivery, ultra-reliable compilers such as just-in-time compilers, cloud computing, and the economy of scale.

Franz noted that there were two costs involved. The first was the cost of generating the diverse object modules. As a compiler generates object code by choosing from several possibilities, it could simply save all those possibilities and the app store could choose one set to generate the software to send to the customer. Some users would get non-optimal software; when asked about this, Franz said he did not know the impact of this or how much degradation would be involved. Then there is an up-front cost of actually distributing the software. Assuming $0.18 per hour for cloud computing, Franz estimated that each unique version of Firefox, which has 30 million lines of code, would cost $0.09.

In the discussion, someone pointed out that use of the app store to generate the diversity solves many problems but introduces a single point of failure. Someone also asked whether redistribution would need to be banned, and Franz replied that he didn't care: the number of diverse object files would make attacking much harder even in the face of redistribution. A number of people, including Franz, emphasized that this method works better as the scale increases, and if done on a small scale rather than a large one, the benefit would be minimal. Questions were raised about the amount of diversity; this is one of the research areas that must be explored. For example, the amount of diversity would control whether one could generalize from a patch to be applied to a single software instance to an attack that would work across all patches.

- ***On Information Flow for Intrusion Detection: What if Accurate Full-System Dynamic Information Flow Tracking Was Possible?***
  Mohammed I. Al-Salah and Jedidiah R. Crandall

What if we used information flow tracking methods for intrusion detection instead of (or in support of) the currently popular appearance-based or behavior-based methods? If we did, then we should research ways of approximating dynamic information flow tracking as accurately as possible. This also means that we should move to the paradigm of looking at global properties and to dynamic quantitative flow analysis.

Jed Crandall described their use of the dynamic information flow tracking (DIFT) method, which tags/taints data in order to measure the information flow throughout the system. As a first step towards this goal, they created a prototype DIFT system that supports address and control dependence in a general way and measures these specific information flows.

A lively discussion ensued. At one point, Crandall emphasized that, because they have only a prototype system, they did need more accuracy and made a lot of approximations. He viewed their biggest research challenge as how to handle the expansion of taint. Regarding the quantities of information passed, Crandall mentioned that they looked at data provenance and forensics and they now look at threat models. They also want to have a system where they can make statements such as "These data first were cut-and-pasted from Office and saved to a text file and then were saved to a USB memory device"—in such a case everything becomes tainted but just a little bit at a time, so (as an example) labeling it with bits from a tainted file might wind up transmitting something like 0.00005 bits from the contaminated file.

- ***A Stealth Approach to Usable Security: Helping IT Security Managers to Identify Workable Security Solutions***
  Simon Parkin, Aad van Moorsel, Philip Inglesant, and M. Angela Sasse

Simon Parkin presented a stealthy approach to convincing IT security managers and chief information security officers (CISOs) to include usability in their policies. The tension between security and usability often makes people believe that the two are incompatible, so they sacrifice usability for security, but research has shown that, in general, the two can be compatible. CISOs, however, typically do not know how to apply current research in usability. This work engaged CISOs by using their language and methods, testing the ideas by involving them in a user-centered design. They led three CISOs from large organizations through a semi-structured requirements analysis for a password policy. The researchers developed a tool that helped show the impact of various parameters such as length, complexity, change notification, and other aspects of passwords on the productivity of workers, the cost of the control, and the number of breaches of security. They did mock-ups of the tool's output and met with the CISOs to see if they could integrate usability issues into their existing processes.

The discussion was lively. The researchers picked a password policy to begin with; someone asked whether they included organizational processes such as auditing and so forth. This led to some comments that the choice of a user-centered activity approach rather than a user-centered design might be more fruitful, because while CISOs care

about users, they often do not know much about what users are doing. A recurring question concerned the background of the CISOs in the study. In particular, what did the CISOs think was most important? This is based in part on their evaluations, the criteria for which varied widely depending on the business. One CISO said that productivity did not count but that the lack of security breaches did.

This also led to a discussion of assumptions. The study focused on CISOs, but did not examine other parts of the organization, nor did it gather data from anyone other than the CISOs. The presenter responded that gathering data from others was the next planned step in the research, and by talking to the CISOs, they learned which stakeholders they needed to talk to later. They plan to incorporate this information into their tool, which prompted another observation that senior management often wants to hold back information from subordinates, so they may not be able to include it.

The session concluded with a suggestion that some day the tool might be able to replace CISOs entirely. Someone observed that CISOs have two roles: making the policies, and being fired when there is a breach. The conclusion was that, assuming the tool could incorporate artificial intelligence techniques and be generalized to include everything, it might be able to replace the CISO—but this would be very bad for users.

- ***VM-based Security Overkill: A Lament for Applied Systems Security Research***
  Sergey Bratus, Michael Locasto, Ashwin Ramaswamy, and Sean Smith

Michael Locasto and his co-authors challenged the common idea that virtual machines are isolation and introspection panaceas and, in particular, that any credible research into kernel-level modifications requires the use of a virtual machine because one can only monitor software effectively from a lower layer. He argued that emergent complexity in the virtualization environment greatly increases the attack surface. There are pressures from below (such as remote managers, and the need to maintain both the host operating system and the virtual machine operating system), at the interface (the need to create and maintain guest-host APIs so the host can extract data from the processes running on the virtual machine), and from the virtual machine itself (the information flow policy and the machine's use as a resource emulator and controller).

Locasto suggested that the key research challenge in this space requires devising mechanisms that monitor systems at the same privilege level rather than from below. These "self protection" mechanisms still represent an interesting path of research. Indeed, in some scenarios, software must monitor itself.

Someone pointed out the "observer effect," in which the very act of introspection of a virtual machine taints the guest operating system, and asked if there were another way to get the information. But this paper was simply identifying a semantic gap between what is being monitored and observed. Determining the best mechanism to extract the information is a fundamental challenge and one the authors did not solve. Another part of the discussion brought out the fact that, originally, virtual machines were designed for multiplexing separate systems onto one piece of hardware. But when the hypervisor uses the system hardware for paging, for example, it no longer acts as an intermediary between the guest operating system and the host operating system. While this improves performance, it also greatly increases complexity because now new traps and checks must be added to support the security requirements. Thus, there is a gap between what we want and what the hardware provides. This led to some thoughts about a lightweight virtualization mechanism, which the presenter said was a valid approach but not considered, because the research focused on detecting rootkits.

Locasto concluded with the observation that as our dependency on virtual machines increases, they become less trustworthy as they become more trusted.

- ***A Billion Keys, but Few Locks: The Crisis of Web Single Sign-On***
  San-Tsai Sun, Yazan Boshmaf, Kirstie Hawkey and Konstantin Beznosov

San-Tsai Sun explored the lack of acceptance of single sign-on mechanisms on the Web. Their model consisted of three parties: a user, an identity provider (which manages the single sign-on credentials), and a relying party (which should accept and trust the credential). The relying party has no business incentive to rely on the identity provider, because the relying party is responsible for any loss when the identity provider is compromised or unavailable. Further, relying party sites often rely on user data to survive; not obtaining that data, by relying on the identification by the identity provider, may not be acceptable from a business point of view. The user has no incentive to rely on the relying party's use of the credential, because the different user interfaces among all parties that rely on the identity provider is confusing. Perhaps more importantly, the goal of single sign-on is to simplify the identity management process for the user. But most browsers come with password managers, hiding the complexity of identity management. Further, the use of an identity provider creates a single point of failure: if the identity provider is unavailable, the user cannot get her credential. Finally, privacy concerns abound.

The model was praised for its structure and thoroughness. Someone suggested that the authors should integrate other sorts of architectural impacts and error handling, and much of the discussion that followed concerned these points. The work here focused on individuals, not organizations, so a relying party for businesses rather than individuals may have different issues. Someone suggested that, intuitively, the use of single sign-on improves both security and usability, but no research was cited to support this view, and others disagreed. For example, the use of an identity provider enables a denial of service attack against the user;

the identity provider (who in many cases is invisible to the user) can void the user's credential at will. This led to a discussion of federated identity providers, including the comment that relying parties determine which identity providers they trust. Participants noted that several large corporations and, especially, governments were developing their own identity providers; for example, the German government will provide a single sign-on that makes a verified shipping address available to e-merchants, simplifying the process of purchasing items over the Internet. In addition to browsers providing password managers, many Internet service providers do as well. These two observations vitiate the business case for single sign-on because the ISPs and browsers provide the same level of convenience to the user that an identity provider would.

The discussion concluded with an ancient observation that provides a basis for much of the work: who benefits? What benefit does the single sign-on provide for all?

- ***To Boldly Go Where Invention Isn't Secure: Applying Security Entrepreneurship to Secure Systems Design***
  Shamal Faily and Ivan Flechais

Shamal Faily described the goal of this work as applying models and principles to create, organize, and manage security design elements in such a way as to improve system security. The authors compared three different models of innovation, and contrasted security architects with security entrepreneurs in each. Their models of innovation were incremental vs. radical, component vs. architectural, and static vs. dynamic. They pointed out that the environment shapes architects but entrepreneurs shape the environment. The architect builds things intended to work in the world, and hence must take the environment into account. Further, the security entrepreneur was independent, whereas an architect typically worked for someone and hence was dependent on that person or corporation. An entrepreneur in the audience pointed out that it was easier to have a boss to handle much of the business process work (such as laws, patents, etc.). But the security entrepreneur is free to make her own decisions.

Someone suggested that those who were risk-averse were architects, and those who were not were entrepreneurs. Faily replied that the situation was not that clear-cut; everyone has some aversion to risk. The question was one of risk management. Someone else observed that the difference between innovation and entrepreneurship was that "research is transferring money into knowledge, and entrepreneurship is transferring knowledge into money." Also, the diffusion of the innovation was critical to the success of the entrepreneur. Someone noted that the traits normally deemed good were attributed to entrepreneurs, and ones normally deemed bad were attributed to architects. Faily said that these were simply observations, and the researchers took no position on whether the traits were good or bad; they simply were observed. In response, another participant noted that architects have done things to change the world by building well-built projects. The rejoinder, from yet another person, was that architects don't cause anything to be built; entrepreneurs pull together the strands that enable the architects to build. Further, architecture is easy enough so that architects are unnecessary. A large portion of the audience loudly disagreed.

Faily also elaborated on the art of chindogu to prototype security controls. Essentially, chindogu is the invention of a gadget that solves a problem but introduces so many new problems that it has no utility; the example they used was a baby mop (see the paper for the picture). The security example was a "forget-me-not" digital certificate, stored on a dongle that is attached to the picture of a loved one. The theory is that one would not lose the picture and, hence, the dongle containing the certificate. The authors argued that the chindogu can be used to bridge the gap between open innovation (in which ideas and paths to market are generated) and security design.

As a result, predicting the impact of a new security control from different perspectives must be combined with creativity in order to implement innovative security controls. Further, theories from entrepreneurship can apply to security innovation with minimal changes.

- ***Would a 'Cyber Warrior' Protect Us? Exploring Trade-Offs Between Attack and Defense of Information Systems***
  Tyler Moore, Allan Friedman, and Ariel Procaccia

Tyler Moore applied game theory to a simplified version of the computer security problem. The US Cyber Command has among its missions the defense of US cyberspace while exploiting vulnerabilities of its adversaries. This means that the same actors are both attackers and defenders. The bases of the games used were twofold. First, they assume a zero-sum two-player game in which each player has incomplete information. Second, they assume that making vulnerability information public helps everyone (in this context, both players) to defend their systems, and hiding (or "stockpiling") vulnerability information helps only the stockpiling player attack the opponent. They developed two games.

The Simplified Stockpile Game examines the trade-off between stockpiling and defending (protecting society by fixing vulnerabilities). It has player 1 choosing to stockpile only. It has two parameters, one that represents the player's relative technical ability and the other, the social harm of undisclosed vulnerabilities. When the social cost is 0, both players will stockpile. As the social cost increases, the less technically sophisticated player will compensate for their lack by defending (making information public). The Cyber Hawk Game changes the social harm parameter to be one of aggressiveness, specifically the probability that a player will attack after discovering a new vulnerability. Hence it focuses on the costs and benefits of being aware of vulnerabilities, rather than the results of a conflict. The results of this game show that if technical sophistication is equal,

both players will attack; if not, the technically dominant one will be more aggressive.

The discussion focused on the underlying assumptions of the game. First, the question of what a "vulnerability" was engendered considerable debate, ending when the presenter said the definition was axiomatic, but the vulnerabilities of interest had to be shared by the players and be able to be exploited by at least one. It was also noted that the assumption of disclosure of vulnerability information and patches benefiting the community was questionable because not everyone could fix their system before being attacked; this was a simplifying assumption. Second was the use of a zero-sum game to model this situation. A zero-sum game is one in which losses and gains are balanced, and it is not clear that holds here. For example, in a military situation the military is also concerned about budget and other constraints not dealt with on the battlefield. The third, and most important, challenge was the need to add other players (at least one) and consider non-attribution and asymmetry. Some players may have excellent attack infrastructures, but little infrastructure to defend. A third party may attack one player and make the attack appear to come from the other player, causing the two non-aggressive parties to fight. The concern expressed was that the game might create stability points that do not hold in real life, leading to policymakers making decisions based on a simplified model.

The authors intend to expand this model to make it better match real situations. The first step is probably to increase the number of players. The game is nevertheless useful in trying to understand the trade-offs between attacking and defending.

- ***On-line Privacy and Consent: A Dialogue, Not a Monologue***
  Lizzie Coles-Kemp and Elahe Kani-Zbihi

Lizzie Coles-Kemp explained that their goal was to better understand the types of online privacy dialogues that service providers and users want, in order to understand how their data will be used and shared. The study synthesized two layers. Privacy negotiation is supported as part of the physical and information-processing layers, and how it works is fixed because the system providing the negotiation is understood. But at the communal, cognitive layer, the negotiation is treated as a "black box" as part of the process of managing the user, and how it works is variable because we cannot predict human behavior. For such an open system, we cannot determine a full set of variables to constrain, so we can at best postulate partially grounded theories, and not predict behavior. So the question is, what scientific routes are open to the researchers?

Coles-Kemp described three types of dialogue systems. The first provides information about the service provider, users, and peer groups. The second raises issues about how the service provider interacts with third parties. The third deals only with service provider behavior. A particularly interesting finding was that none of the service providers felt they should negotiate the level of privacy they provided with users; they felt this is done through the legislative process. Other findings were that current privacy dialogue systems performed poorly in terms of providing the users with information, because the users avoided reading privacy and user agreements. At the design level, the dialogues need to address how the users can communicate with the service provider and give feedback for improving the dialogue. At the social level, the dialogues must address the conditions and methods under which the provider will disclose information, taking into account the privacy rules from the law, culture, and commonly held beliefs and norms, and must provide sufficient transparency that the user understands how all this will work.

One participant asked whether the services being studied were private or governmental. The presenter said they were communal services, covering everything from garbage collection to land planning to supervision of social work and provision of payments. The key point is that the services were all about relationships and trust with key workers. As these services move online, those relationships and trust change. Someone else mentioned that one can trade privacy for both benefits and trust, and asked if they had looked at the connection between privacy and trust in depth. The response was that inculcation of trust was a building block, and they looked at how that makes citizens less vulnerable. Going online changes the dynamic from that involved in face-to-face meetings, and the issue is how to compensate for the changes in the dynamic. A third question was the truth of the provided information; several people said that if they did not understand why the online provider needed information, they would simply make something up. The presenter said this phenomenon was quite common among younger people they had surveyed, but others called and asked why the information was needed, or simply disengaged. Finally, there was some dispute about whether privacy notices were a true "dialogue," because they typically specify the terms and the user either accepts them and gets the services, or declines them and does not get the services.

- ***A Risk Management Process for Consumers: The Next Step in Information Security***
  André van Cleeff

The premise of André van Cleeff's paper is that users should have a personal risk management strategy for protecting their privacy online. The paper, however, discusses the complexity that arises when attempting to realize such a system in a tool for end users. As one attendee noted, social disclosure of private information also has benefits that such a tool would need to account for. One major issue that the system faces is how to enumerate the different categories of risk, describe risk details, and assign probabilities of events. Another observed that systems like some online health information systems had to have the number of privacy controls reduced simply because user studies indicated that

people became overwhelmed with controls. One benefit of this paper is to help refute the illusion that someone else is doing risk management for you online; as one attendee pointed out, assuming that delegation of risk management works is a moral hazard.

Attendees vigorously discussed whether humans were good at performing risk management at all, but largely concluded that humans *must* do this because no one else is available to do it for us. Attendees noted that the large literature on risk management suggests that no matter how much additional information people see, they still make poor decisions, and that people misperceive risk, particularly in online scenarios. One attendee noted that "risk management" is an overloaded and abused term and suggested a clever alternative definition: "risk management is planning for your next defeat." Another attendee noted that, as in the theme of the paper, if you personalize risk management, it becomes "fear management."

- ***Ontological Semantic Technology for Detecting Insider Threat and Social Engineering***
  Victor Raskin, Julia M. Taylor, and Christian F. Hempelmann

Julia Taylor noted that people sometimes give off a signal by the way they say things or the information they omit in descriptions of otherwise normal events. Casual conversation or conversation in the same context (e.g., with a colleague) often results in relaxed conversation, and people insert or describe new information or modify existing defaults. Another example occurs with social engineering: amateurs often attempt to mimic all or most of the domain defaults in an attempt to establish credibility—words and actions that a real expert would never do. The paper makes the central observation that such deviations from normal speaking or writing patterns convey information that may be of use in detecting certain types of insider and social engineering attacks. Taylor reported on the use of a mature ontology-based technology for understanding the relative semantics of pieces of a sentence (i.e., the main task is not in parsing the sentence itself). Questions by the audience helped clarify that one potential application of such analysis was detection of insider threats, particularly to support further research and analysis once an initial suspicion is formed (and the analyst thereby has legal access to spoken and text corpora of the research subject).

The subsequent conversation focused on clarifying the value proposition of the system. Taylor stated that the system as such was already a mature technology and the product of more than a decade of research. As Taylor went on to say, the modality of the talk is that they have the resources to interpret sentences, but the paper was about additional capabilities to understand what is *unsaid*. Taylor noted that people unintentionally say things that reveal their habits, values, and defaults. Unless subjects are paying attention, they do not notice such disclosure: the point of the system is to do this. Taylor pointed out, in response to a question about what test data set the authors were using, that the point of the paper was to establish the model and theory, and that rigorously verifying the results was a challenging task. As a result, the system is most applicable in scenarios where the corpus is derived from an ongoing session of legal wiretapping or observation of the suspect's communications, although the authors specifically disclaimed any intent to use the system for providing digital evidence strong enough to stand up in court. Several attendees raised concerns about the eventual use of such evidence in a courtroom.

One attendee noted that an interesting source of data to analyze might be derived from an FBI operation set up to infiltrate a specific drug market with online communications; from the perspective of the drug market, the FBI agents were insider threats.

- ***The Pervasive Trust Foundation for Security in Next Generation Networks***
  Leszek Lilien, Adawia Al-Alawnch, and Lotfi Ben Othmane

Lilien began by noting that trust is a pervasive concept in social interactions. This paper examines the logical foundations for incorporating measurements of trust in new computing systems and presents a case study suggesting how one might accomplish this in the context of designing the next-generation Internet. The paper suggested the design of a "pervasive trust foundation," or PTF. Noting that there are degrees of trust and that trust is usually asymmetric and bi-directional (although one direction might be implicit), the paper suggested modeling such relationships to provide security. One major contribution of the paper is to describe how various security services (SS) might be supported or constructed through the use of an underlying "Trust in the Large" (TIL) subsystem. The paper then considers obstacles to finding an efficient PTF/TIL implementation and suggests approaches for decreasing the performance burden of checking security properties by taking advantage of the TIL layer.

Discussion included a wide variety of interesting issues, as trust seems to be a cross-cutting issue in the information security field. One attendee pointed out that the definition of trust offered in the paper might need to be augmented with the concept of trust sourcing: building trust requires a well-founded evidentiary process to establish a trust basis. One major theme of the discussion was how usability of a software artifact leads to a simple mental model, which then leads to or creates trust in the operation of the software artifact. It was noted that such a "contract" is a substitute for security in the sense of measuring real properties and constraints on execution. The discussion also focused on how security typically implies some objectively measurable value; trust is usually a subjective notion—in many situations, appropriate security measures should be determined by the size of the threat, not the level of confidence users have. Finally, one interesting technical issue that was raised was how to deal with efficiently revoking trust relationships (given the assumption of a pervasive trust foundation).

- ***This Is Your Data on Drugs: Lessons Computer Security Can Learn from the Drug War***
  David Molnar, Serge Egelman, and Nicolas Christin

David Molnar proposed using observational data to draw conclusions about computer security. The controversial issue is what data is useful, what data is biased, and how we compensate for that bias in order to we might draw sound conclusions from the data. The authors draw a parallel between the observational data accumulated over the years in the drug wars; there is a body of data (the STRIDE data, a time series of price and purity of the drugs seized) from drug busts, and this data is used for research. It is controversial, in that the data was not intended for use in research, and so the biases inherent in it must be noted and compensated for. Further, the data comes from those who have been caught and does not reflect the price differences between those and others who do not get caught. The authors then ask what we can learn from the drug wars. Their hypothesis is that applying lessons from the analysis of the STRIDE (and other) data will improve observational research in computer security.

Molnar noted that papers drawing conclusions from observational data in computer security have a shaky basis; indeed, one well-known paper on purchasing services to crack CAPTCHAs was known to be flawed because it drew unwarranted conclusions. A participant asked if this demonstrated naïveté on the part of the authors or reviewers or both, and the presenter responded that that is the reason he brought up the controversy about the STRIDE data. A second participant said that the authors of the CAPTCHA paper probably felt the data, although biased, was better than no data; the retort, from another member of the audience, was that it would be worse, because if you are using data you know to be flawed, you must reveal the flaws and the possible consequences of using that data. This led to a discussion of what kind of conclusions the experimenter can draw, and the consensus (such as it ever is at NSPW) was that the authors must be clear about the methods they used and demonstrate that the conclusions are reasonable given the data (biases and all). It was also pointed out that errors may be caused by unknown factors as well as malevolence and incompetence; for example, when measuring the gravitational constant, a series of experiments used different methods, and lots of the research was to analyze others' methodologies to find out why the results disagreed. This is actually supportive of the experimenters, because it helps them refine both their methodology and their understanding of the experimental process. It reveals things that were not known before.

The discussion then moved back to the analogy of data from drug busts and data from computer security. In their talk, the presenters had observed a huge price dispersion in drug prices (the figure cited was a $6000 difference in the price of a kilogram of cocaine in New York and in Boston). Someone noted that the price dispersion for drugs was based on physical separation and asked if this was also true for computer security data such as stolen credit cards or attacking CAPTCHAs. The presenter pointed out that the problem was not physical separation but logical separation, specifically the different groups with different rules of access creating different markets—and it was necessary to study these differences in order to understand how they affected price. The discussion concluded with some thoughts on experimenter bias, the well-known ways other fields compensate for that, and the question of how to do the same in computer security experiments.

- ***Relationships and Data Sanitization: A Study in Scarlet***
  Matt Bishop, Justin Cummins, Sean Peisert, Anhad Singh, Bhume Bhumiratana, Deborah Agarwal, Deborah Frincke, and Michael Hogarth

Matt Bishop focused on an interesting path in data sanitization. Although the field has had a lot of examination (from privacy-preserving databases to sanitizing network traces), the authors have been examining a new paradigm for dealing with a major challenge in this space: the availability of external knowledge to an attacker wishing to reverse the sanitization effect or otherwise infer some knowledge from the sanitized data trace. One key insight is that the authors treat sanitization as a problem of *risk*, not certainty. They assume that (1) relationships used by attackers are unknown to the sanitizer, (2) effective sanitization might not exist, and (3) inferences might not be correct, but incorrect conclusions are potentially damaging. As a consequence, they construct a framework for asking, "What relationships enable desanitization?" and "How likely is it that the existence of these relationships is discoverable?"

The framework Bishop presented uses an ontology to help reveal the conflicts between a privacy (or sanitization) policy and an analysis policy. The threat model informs the structure and content of the privacy policy, and security requirements (or whatever domain requirements exist for the specific analysis being performed) inform the structure and content of the analysis policy. The system logically performs analysis on both the raw data set and a sanitized version of that data set. The system compares the results and produces a set of conflicts arising from the two policies. The system enables an expert to help resolve these conflicts. The overall purpose of the work is to help inform users and organizations interested in data sharing about the risks specific to their activity. Bishop also observed that one big practical problem is how to create a "consumer-friendly" assurance argument.

The discussion largely focused on attempts by the audience to understand the semantics of sharing, (de)sanitization, and policy construction. The audience was also curious about the role of the ontology. Bishop asserted that the ontology was useful in helping reconcile the fields in the privacy and analysis policies (as these may come from different domains and use different terms and descriptions). Bishop also noted that one big win with using an ontol-

ogy over a simpler mapping of field relationships was that several tools already exist for manipulating ontologies that make such comparisons much easier than writing a system from scratch. Attendees noted references for several types of desanitization attacks, including inserting marker data and whether it was a risk that an attacker might reverse-engineer the ontology. One observer noted that this work would be valuable because many times in research, a large data collection and sanitization effort is undertaken that is wasted when researchers realize that they can't use or analyze this data in a meaningful way.

PANEL

- ***Why Is There No Science in Cyber Science?***
  Panelists: Richard Ford, Carrie Gates, Lizzie Coles-Kemp

The last presentation was a continuation of the opening panel, but with a surprise: three new panelists presented their thoughts on cyber science.

Richard Ford began by arguing that science is a friend, not an enemy; it is how we actually produce knowledge. Change is incremental and slow, but can begin now—for example, by rejecting papers that do not demonstrate good science. Doing better science does not mean that we will be any less productive; while it is harder, the results are much more long-lived than non-scientific results. So, we must ask ourselves: do we really want to understand our world, or just get published?

Lizzie Coles-Kemp followed. In every artifact, there is a physical object and a social object, and we need to respect this duality. How we produce knowledge about each dimension has its traditions. Some papers presented over the past three days straddled the physical and social worlds. What are good scientific methods in each of these paradigms? How might we use them in each, and how might we take this forward?

Carrie Gates went last. She argued that there was an industry perspective involved. Research should make a difference; indeed, only useful research is good, whether or not it is scientific. Time is of the essence because if the work takes too long, a competitor will grab market share and the company waiting for the science, or for the results, to be scientifically valid will lose; in other words, science takes too long. Incremental improvements are good enough; indeed, the perception of improvement is good enough, even if in reality there was no improvement.

Almost everyone in the workshop indicated that they wished to speak. Someone pointed out antivirus as an example of the need for non-incremental research. Current antivirus software is very poor, due to the near-term focus and incremental approach used to improve the software. If we applied more science to it, we might obtain better results. Gates asked why researchers had not given industry something better than the current mechanisms, and the response was that better ways were known, but the market has yet to adopt them. Someone else suggested that one needs science to produce generalizable results; the response was that one should not conflate product development with fundamental research. Another comment was that industry might not want scientific results that they can use; they focus primarily on whether users need their product, not on whether the product does exactly what the sales force claims it does. The question of how to determine whether something is useful arose, with a participant noting that what may seem utterly useless (the example used was Reimannian geometry) may turn out to be extremely useful (in the example, it was found to be the actual geometry of Einstein's theory of space-time). The panel agreed with this point, arguing that the requirement that results be useful immediately is killing the field.

Someone asked for ideas on how we might change the culture of computer security to be more scientific. Suggestions included not rejecting papers that had claims not supported by science, but instead working with authors (possibly through a shepherding process) to ensure that the claims are appropriate to the work done; requiring a methodology section describing how the experimental work was done; and releasing code and data whenever feasible. With respect to this last point, someone else said that when the research involved the use of proprietary code, releasing the code may not be possible and so if code cannot be released, the results should not be automatically rejected. The panel reminded everyone that as reviewers, we have considerable power to change the culture, and we should use it.

Each panelist then said a few words. Gates argued that the best way to get industry to value and use research is to embed researchers in the different industries: this would communicate the research results in a way that could be incorporated into products. Ford commented that we should change the culture in small steps, and think about how best to communicate the needed changes to others. Coles-Kemp concluded that more venues such as this workshop would raise awareness of the problem and ways to change the culture.

NEXT WORKSHOP

The next NSPW workshop will be held at the Marconi Conference Center in Marin County, California, from the evening of September 12, 2011, through lunch on September 15, 2011. Sean Peisert will be the general chair and Richard Ford will serve as vice-chair. Carrie Gates and Cormac Herlihy will lead the program committee. Submissions will be due by April 4, 2011. Details on how to submit papers will be posted to the Web site http://www.nspw.org in the near future.