


Mechanism Design for Improving Hardware Security Workshop Report



CCC

Computing Community Consortium
Catalyst



The material is based upon work supported by the National Science Foundation under Grant No. 1734706. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Mechanism Design for Improving Hardware Security Workshop Report

Workshop
August 24-25, 2022

Workshop Organized by

Simha Sethumadhavan, Columbia University / Chip Scan
Tim Sherwood, UC Santa Barbara / Cycuity

With Support From

Maddy Hunter, Computing Community Consortium (CCC)
Haley Griffin, Computing Community Consortium (CCC)
Catherine Gill, Computing Community Consortium (CCC)
Ann Schwartz, Computing Community Consortium (CCC)
Adam Hastings, Columbia University
Deeksha Dangwal, UC Santa Barbara

Simha Sethumadhavan has a significant financial interest in Chip Scan Inc.

Timothy Sherwood has a significant financial interest in Cycuity Inc.



CCC

Computing Community Consortium
Catalyst

Introduction1

What is Mechanism Design?3

Recommendations.....4


Process5

Summary10

Appendix11

Workshop participants12

Agenda14



Introduction

From election security to critical health applications, trustworthy hardware is the bedrock of a modern free and healthy society. Once niche and arcane, the field of hardware security has recently become one of the most pressing issues in cybersecurity. Microarchitectural side channel attacks like Spectre and Meltdown have shown how pervasive, dangerous, and hard-to-fix a hardware attack could be; integrity attacks such as Rowhammer and CLKSCREW show how attackers can practically overwrite user data. As hardware development becomes more like software due to availability of free hardware designs and tools, the prevalence and discovery of these types of design/security problems are likely to accelerate. Similarly, the benefits of hardware-based tokens for multi-factor authentication are well known, but adoption has been sluggish. Especially concerning is that these problems, while well-known and publicized, have generally not been fixed pervasively. Why? The answer, perhaps, is not only a lack of technical solutions that are considered practical, but also a series of market failures such as information asymmetry, prisoners dilemmas, and markets for lemons, which disincentivize those who are able to fix serious security vulnerabilities from doing so.

The workshop on Mechanism Design for Improving Hardware Security brought together experts in hardware and software security, economics, and government policy to explore the opportunity to ask new research questions with a significant potential to improve the design and uptake of hardware security mechanisms. Through a combination of interdisciplinary discussions and active perspective taking exercises, participants considered both traditional technical solutions and new mechanisms to incentivize designers, system integrators, and users to create and maintain the security of their systems. The workshop aimed to address a number of key questions, including: how current policies and market structures disincentivize hardware-oriented security solutions and how these issues can be addressed through technical and policy frameworks; the mechanisms necessary to enforce government mandates on the allocation of resources for security and the process for determining the appropriate level of investment; the potential for hardware innovations to impact software dark economies; and the incentives and education/certification requirements needed to encourage timely patching of hardware bugs.

Ultimately, the goal of the workshop was to identify practical strategies for improving hardware security, broadly defined to include security of the hardware, and hardware supported software and user security, and to consider the broader policy and economic context in which these solutions could be implemented. Three main ideas underpinned all of the discussion.

► **Hardware security is the bedrock on which all other system security properties rely.** Hardware plays an absolutely critical role in the overall security of a system, serving as the foundation upon which software security measures are built. However, building secure hardware is a particularly challenging task, due in part to the complexity of modern hardware systems and the need to account for a wide range of threats. Additionally, hardware vulnerabilities can be difficult to detect and fix, as they often remain hidden until they are exploited by attackers. This makes it important for designers and manufacturers to anticipate and proactively address potential vulnerabilities in the design and production process. Furthermore, the fact that hardware components have a longer lifespan than software means that vulnerabilities that are discovered after deployment can be much more difficult and costly to address. If we are going to figure out how to do security well we are going to need a solid hardware foundation on which to build.


► **System security, and thus hardware security, is not solely a technical problem.** To design and implement effective hardware security solutions, we must understand the many factors that influence their development and adoption. These include technology, economics, law, education, equity, and policy, which intersect in complex ways and shape the security landscape. Current policies and market structures may discourage the creation and use of hardware security mechanisms, making it harder to secure systems against various threats. To overcome these challenges, we must conduct research that produces new and more effective hardware security mechanisms. This research should adopt a multi-disciplinary approach, considering the technical, economic, and policy aspects that shape the development and adoption of hardware security solutions. Only by

actively examining how these factors intersect and how they can be harnessed to create widely adopted hardware security mechanisms, we can make progress in this important area.

► **Measures of success must consider the many intersecting implications of the future of system security.** Typically when considering hardware security from a technical perspective one may be considering the stake of a developer of a piece of intellectual property, the manufacturer of a device, or the infrastructure of a particular service provider. Much less often is the societal implications of such technological change considered, and as is pointed out in “The Devastating Consequences of Being Poor in the Digital Age”, the implications of this can be dire and lasting. There is a need for hardware security research that considers a far more rich set of metrics of success including the total security effectiveness, the cost of ownership, the cost of development, the ethical, legal, and environmental implications, the broader national security posture, and the impact on equity and justice.

This is clearly a complex and highly interacting economy of competing interests to navigate and one of the reasons for slow progress on these issues may be the failure of the existing market. Hardware security usually comes at a cost in terms of performance, power, or area; present issues in hardware security can be seen as the result of the players in the game of hardware security finding ways of avoiding paying this cost. Without a counteractive incentive, these costs can easily dominate decision making to the detriment of other critical factors. In fact even this description is far too simple as hardware security decisions exist within a complex multidimensional space where security effectiveness, policy enforceability, total cost of ownership, the increased cost of development, impacts on system efficiency and performance, ethical and legal implications, equity, and usability all interact in complex ways. So how does one begin to think about restructuring such a complex set of interests?

At this workshop, experts in hardware security, software security, economics, government, and policy investigated the potential of using *mechanism design* to balance these complex, competing interests. Mechanism design is a branch of economics that focuses on designing incentives, institutions, and rules to create an environment where rational agents are incentivized to choose strategies that collectively achieve some desired outcome (i.e. improve hardware security and security adoption). Thus in addition to looking at traditional technical solutions, the workshop considered the potential for new research into using mechanisms and designing incentives to drive systems designers, users, and policymakers towards improving the design and uptake of hardware security in today's systems.



What is Mechanism Design?

Mechanism design is a field of economics that stems from game theory: it supposes the existence of a game with rational players who choose strategies to achieve some personal objective. However, whereas game theory supposes the rules of the game to be fixed (and focuses on the players' strategy-making), the focus mechanism design is on a new, special player—the game designer—who has the ability to change the rules of the game itself. Mechanism design is the study of the levers, sticks, and carrots that game designers have available to them to nudge the game players into choosing strategies that achieve some overall desired outcome (giving mechanism design the nickname “reverse game theory”).

This “objectives first” approach has seen widespread usage across many domains. For example, in auction theory, the game designer can choose the rules of the auction so that bidders place truthful bids and that the auction good is guaranteed to be awarded to the bidder to the highest valuation (e.g. in Vickrey auctions); in voting theory, the game designer can choose voting systems that eliminate so-called “strategic” voting (where voters may not vote for their preferred candidate). And in environmental economics and policy, mechanism design has been employed via cap-and-trade policies, which incentivize lower emissions via a marketplace of carbon credits.

At present, there exists a “game” of security, played between product vendors, users, and attackers, but little in the way of a game designer. There is immense potential for mechanism design to solve the pressing issues of hardware insecurity through a game designer who is able to improve outcomes by adjusting incentives and punishments. Framing the problems of hardware security in terms of mechanism design lets us ask: what tools do the “game designers”—aka regulators and policymakers—have available to address these problems? What are the “rules of the game”, and how can they be adjusted to incentivize product vendors into providing security at a personal cost and on behalf of others? What tools are available from the fields of technology, economics, law, education, equity, and policy that can be combined to improve the state of hardware security?

Mechanism Design and “Dirty” Air

Mechanism design is an applicable tool in situations where a game designer has the ability to choose or modify the rules of a game and wishes to achieve some overall outcome. A recent and real-world example can be seen in the world of Formula One auto racing. In previous years, Formula One has suffered from the so-called “dirty air” problem, where racecars would leave massive amounts of turbulent air in their wake. This turbulent “dirty” air made it very difficult for trailing drivers to overtake the cars in front of them, even if the trailing car was faster. As a result, drivers would rarely choose to overtake the cars in front of them, making the races boring for fans. Not wanting to lose viewership, the FIA (the governing body for FI) employed a form of mechanism design to solve the problem. The FIA (the game designer) could not force the drivers (the players) to attempt to overtake each other more (i.e. could not force the players to change their strategy). However, the FIA *could*—and *did*—change the race car regulations (i.e. the game design) to encourage more overtaking: For the 2022 season, the FIA required automakers to design their cars’ aerodynamics in such a way that pushes the “dirty” air upward, leaving a pocket of “clean” air for the cars behind, thus making it easier for trailing cars to overtake the ones in front of them. As a result of the 2022 car changes, there were 30% more overtakes during the 2022 racing season, and the season was described as “one of the all-time great seasons in Formula One”.

Recommendations At-a-Glance



01

Foster Diverse Educational, Professional, and Industrial Communities in Hardware Security




02

Lay the Scientific Foundations for Work that Combines Incentives and Technology



03

Make Security Accountable and Explainable



04

Co-Develop Emerging Technologies with the Understanding of their Hardware Security Ramifications



05

Prioritize the Human Impact of Hardware Security

Process

Prior to the Workshop

In January of 2022, there was a virtual pre-workshop orientation. Hosted by coorganizers Timothy Sherwood and Simha Sethumadhavan, the orientation included a presentation on potential topic areas, the mission of the workshop and important questions the workshop was seeking to address in addition to a live Q&A session. Orientation attendees were then invited to join a slack channel where important conversations could continue and connections could be made. Attendees were also told that the organizers would be collecting white papers to help guide the workshop agenda and select attendees.

Shortly after the orientation, the call for white papers went out on the CCC blog, the CCC website and to various mailing lists. Interested parties were requested to write no more than two pages investigating ways to improve the design and uptake of hardware security mechanisms. The call for white papers included some example topics and questions for the white papers, but participants were encouraged to write whatever they found to be important that pertained to the mission of the workshop. Topics of interest for the position paper included, but were not limited to:

- ▶ How do current policies and market structures disincentivize hardware oriented security solutions? How do we fix this: what technical and policy frameworks are necessary to make progress in this area?
- ▶ What are the mechanisms necessary to enforce a government mandate that says that X% of the performance or cost should be set aside for security? What mechanisms

are necessary to determine X? How often should X be determined? Is there a quantitative approach for the organization to use up this security budget? How would this be enforced on user systems? Are there alternate government mandates that are actionable and can be supported technically?

- ▶ Is there an equitable way to proportion the benefits of security and impacts of security attacks? What hardware support, if any, is necessary to facilitate this process?
- ▶ How do we establish a chain of responsibility for malicious and negligent action while also maintaining privacy?
- ▶ Are the mechanisms for hardware security different from those required for privacy?
- ▶ How can hardware innovations fundamentally impact software dark economies? What incentives are necessary to patch hardware bugs in a timely manner?
- ▶ What education/certification requirements are necessary for increasing the awareness and application of hardware security solutions?
- ▶ Are there parallels to software certification requirements for hardware? What would these assurance/certification requirements look like?

White papers were due in April, four months prior to the workshop. Twelve papers were submitted. The rest of the participants were invited to ensure a distribution across both industry and academia, small and large organizations, and expertise across hardware and security.

At the Workshop

Day 1

The workshop was structured with a mix of talks, group discussions and breakouts. Day 1 had two 30 minute presentations, one on Incentivizing Cybersecurity and the other titled Reflections on Assurance. Each presentation was followed by a 45 minute group discussion where participants had the opportunity to ask questions and discuss as a larger group. The participants then discussed how topics apply to real world scenarios, which would help frame the breakout discussion on day 2. The workshop then adjourned for dinner.

Day 2

At the beginning of day 2, the organizers gave a quick recap of the previous day's discussion topics, and then presented a mock scenario in which the US government had just passed a new law to no longer purchase any new hardware with known vulnerabilities. The group then broke into six smaller groups, and each was assigned the role of a different stakeholder, these being: regulators, hardware companies, software companies, attackers/bad actors, system integrators, and end users. The goal for each group was to maximize their value or profit. The breakout session was an hour long and followed by an hour long report back where each group shared what they had discussed. After a break, there was an hour-long Industry Panel. The panel was led by a moderator

who asked the panelists questions about various stages of development, explored barriers to adoption, important real-world constraints often under-considered in research, and lessons learned. The panel session was followed by breakouts similar to day 1. Each group was given a different set of challenge problems to talk through as different stakeholders.

After lunch the breakouts reported back to the larger group. Each breakout group had a leader who shared a 2-4 minute discussion summary with the larger group.

The workshop then turned to focus on generating ideas for the workshop report in order to consider major takeaways while topics and discussions were still fresh in everyone's minds. There was an open brainstorming session for participants to discuss as an entire group the repeated themes brought up during the workshop. As people shared their key takeaways, the ideas were captured for all to see.

Finally, the last session of the workshop was a government panel, with representatives from the Office of Naval Research (ONR), the National Institute of Standards and Technology (NIST), the National Reconnaissance Office (NRO), the National Science Foundation (NSF) and the Defense Advanced Research Projects Agency (DARPA). This panel discussed common pitfalls and weaknesses of hardware and software systems and advice for junior researchers trying to obtain their first research grant, among other topics.

Recommendations

(I) Foster Diverse Educational, Professional, and Industrial Communities in Hardware Security

To address the challenges of hardware security and ensure that it is a priority for designers and manufacturers, it will be helpful to build a robust and diverse community of people working on these issues. This could include providing opportunities for individuals to pursue careers in hardware security and creating an “on ramp” for people to get involved in the field. Additionally, it may be helpful to collate and present unbiased perspectives on hardware security issues from a variety of stakeholders, including different nations, industry players, and end users.

One way to encourage collaboration among the community to combat hardware problems is increasing the prevalence of open source hardware. A potential issue with this is that larger projects may have better resources and processes in place to ensure security, while smaller projects may struggle to keep up with the maintenance and bug fixing required to maintain secure systems. This could lead to a situation where smaller open source hardware projects are more vulnerable to security threats. To address this issue, it may be helpful to provide incentives or support for hardware researchers to prioritize security in their work. This could include providing education and training on security best practices, as well as funding or other resources to support the maintenance and improvement of security in open source hardware projects. By addressing these challenges and providing support for hardware researchers to prioritize security, it may be possible to help ensure that open source hardware projects are secure, well-maintained and trustworthy.

Continuing in the vein of collaboration, there appears to be a significant amount of interest in building a community-driven approach to addressing hardware security issues, with a focus on bringing people together to share solutions and best practices. An important aspect of this is the process around vulnerability disclosure and patching for hardware and systems, which at present lacks clear standards, guidelines, and expectations. Of special concern is the reported cases of product vendors who, after being informed of a vulnerability by third-party

researchers, delay or embargo the public disclosure and patch of the vulnerability for as long as possible, presumably because the vendors find some economic incentive against doing so. Without clear responsibilities and guidelines, product vendors may not be properly incentivized to patch vulnerabilities in a timely manner. Solutions to this dilemma may require the establishment of standards for vulnerability disclosure and patching for all invested stakeholders, including not just the product vendors themselves but also relevant parties from policy, law, and regulation. Setting a clear understanding of expectations for addressing hardware vulnerabilities can help encourage a culture of transparency and accountability.

Open secure exemplar designs, or open source designs that prioritize security, can also be an important resource for the hardware security community. These designs can serve as examples of best practices and can help to educate and inspire others to prioritize security in their own work. In addition to the value of open secure exemplar designs, it is also important to consider the role of formal and informal education in improving hardware security. Providing education and training on security best practices can help to ensure that hardware designers and manufacturers have the knowledge and skills necessary to create secure systems.

Attendees also discussed how security analysis should be conducted for systems that are not yet available and what would be considered credible threats to protect against these systems. Current approaches to formulating threat models were deemed reasonable but could be improved by collaborating with industries working on these technologies. Some workshop attendees felt that support for such cooperation could be beneficial.

We strongly recommend growing and broadening the community of people engaged in hardware security, including practitioners, researchers, and policy professionals through innovative engagement and collaboration models.

(2) Lay the Scientific Foundations for Work that Combines Incentives and Technology

The workshop participants agreed that scientific foundations are necessary and important for measurable progress in the area of mechanism design. There was a clear preference for hypothesis-driven research with quantitative, repeatable experiments, preferably carried out at scale. While some of these elements are performed by computer scientists today to support their new ideas, mechanism design not only includes technical solutions but also economic factors and incentive structures with serious real-world consequences. As such it was felt that there is much more that needs to be done to establish the scientific foundations of this work.

Some of the major challenges in this area are: 1) how do we establish and validate who the stakeholders are in mechanism design? 2) What is the objective function for measuring progress? In multidimensional mechanism design, it is possible to measure individual elements (e.g., the value of goods held by an individual goes down), but there are not good metrics available to capture others (e.g., individuals feel more secure or feel that a mechanism is more culturally sensitive than others), and in some cases, metrics might be feasible but do not currently exist (e.g., the risk of a new attack vector emerging in the next 5 years). The recommendation from the workshop attendees was to support research activity in the development of these metrics including user behavior studies. There is also the research question of how to combine these different dimensions of progress, and if this is even desirable to do so. 3) Studying the real world impact of multi-agent mechanisms in a realistic manner requires modeling and simulating many actions of each agent. How do we ensure that the community has access to computing resources necessary for such modeling?

Another important aspect of such a modeling study is an assessment of risks and costs. Research is needed to determine the risk of human factors such as resistance to adoption and effectiveness of use, and methodologies for estimating costs such as the real and opportunity financial costs, the human cost in terms of anguish and injustice, and systemic cost such as lack of trust that impedes societal progress.

In terms of mechanisms to study, the workshop attendees suggested that work should take into account historical context so we don't repeat mistakes: mistakes both in terms of policies that are not effective but also understanding broader impacts of those changes on society. Another recommendation was to study the optimal combination of liabilities and regulations.

Finally, there is a need for innovative and effective new ways to understand the complex network of interactions, and how technological changes would make impacts both up and down stream. This work must be done in concert with computer science and engineering.

We strongly recommend supporting creative new research on scientific foundations of mechanism design to address hardware and systems security in a holistic manner.

(3) Make Security Accountable and Explainable

An important aspect of mechanism design is accountable and explainable techniques that can act as a powerful feedback loop for mechanisms to self correct. Workshop attendees strongly favored designing systems so that they are not only secure but explainably so. They recommended the need for research on technically rooted mechanisms that provide evidence to the user that their system is secure. The goal is to provide actionable information in terms of incentives, costs and risks to the end user. How can such "proofs" of security be designed to be able to recognize that a certain security element has been correctly designed, correctly implemented, correctly installed and configured, and correctly operated? And how can these proofs be provided in a continuous and engaging manner to the users? When multiple parties are involved in designing the security element, technical mechanisms for tracking and enforcing security privacy requirements are desirable, and research is needed on this aspect. Once security breaches do occur, automatic mechanisms for attribution were desirable. Like with other areas, workshop attendees suggested that novel security techniques that provide meaningful evidence and usefulness should be favored over mechanisms

that have been unsuccessful at increasing the security posture (such as documenting secure practices as proof of security).

We strongly recommend supporting research on development of composable, explainable, and accountable security technologies and mechanisms and policies to incentivize adoption of such technologies.

(4) Co-Develop Emerging Technologies with the Understanding of their Hardware Security Ramifications

The simultaneous need for improved cost, usability, energy efficiency, and security and privacy will make hardware innovations a centerpiece of emerging systems. In this context, it is essential to study possible protection and privacy ramifications early in the design stage, as designs evolve and after deployment. Given that novel attacks drive security improvements, workshop attendees emphasized the importance of researching the threats posed by next-generation malware for these emerging systems.

In AI/ML, support for explainability, model reconstruction/extraction prevention, and on-the-edge training will be vital to realize these systems' benefits responsibly. Hardware support is expected to be essential to support these features. What threat models should these features be designed to withstand? What types of attacks can emerge for these systems? While there are myriad security considerations around AI/ML, this is an area where hardware is changing very quickly and significant attention to the interplay of that hardware and security is required.

The security of new models and substrates for computation, such as quantum computers and computing genomics, were also important topics discussed. Such technologies are always, at least initially, expected to be far and few, cost billions of dollars, and thus are likely to be utilized as a shared resource. How do we prevent malicious actors from running programs that destroy or impair these computers? Can we statically analyze these systems and certify their trustworthiness? Are

there new types of malware that are possible on such machines? Research is needed to understand the scope of vulnerabilities latent in these radically different means of performing computation.

While most of the discussion pertained to emerging systems, attendees also cautioned that a spectrum of attacks are possible as technologies receive adoption. For example, practical trust and security issues surrounding disaggregated computing and chiplets were also discussed and the need for further investigation in this area was emphasized. One specific area of interest was how security vulnerabilities arise when individually secure components are composed into a more extensive system in the context of these systems. Of course, attacks on purely traditional systems should never be ignored as the security of emerging platforms often rely heavily on the classical (legacy) systems that they will utilize.

We strongly recommend the support research that identifies vulnerabilities unique to emerging platforms and develop new classes of attack to motivate security and privacy improvements prior to mainstream adoption.

(5) Prioritize the Human Impact of Hardware Security

Hardware security is a pressing issue that affects everyone, from tech industry professionals to the general public. Unfortunately, hardware vulnerabilities such as Spectre and Meltdown can have widespread and long-lasting consequences, including information leakage and data tampering. These, and other, vulnerabilities have a disproportionate impact across the nation and we must fully recognize the technological, social, and systemic issues we face as a society are deeply intertwined. Camille Stewart Gloster, now Deputy National Cyber Director, Technology & Ecosystem Security, noted in her 2020 post "Systemic Racism Is a Cybersecurity Threat":

For years there have been well documented discussions about the need to expand gender and racial diversity in cybersecurity. People have argued that if we address social and systemic issues

separately, we will get the technology right. However, the social and the technological are mutually constitutive. Bringing in new points of view is crucial to cybersecurity, but we also have to change the systems in which technology is embedded and review technology against the backdrop of larger systemic issues to reduce vulnerabilities.

Understanding the cultural nuances of technology use and access is integral to building policies and technical solutions that secure systems and serve people. Understanding the differences within our communities and lived experiences domestically and abroad will help build resilience into mobile voting and any other policy and technical cybersecurity solution we seek to implement.

Understanding the ways in which technology is actually exercised by individuals is essential to achieving the goal of a usable, practical, and effective security. Too often we attempt to simplify such problems by pretending that the most important hardware security truths are to be found outside of any lived experience, or perhaps even worse, delude ourselves into thinking that we currently have the representation in the field to be able to say anything meaningful about that experience. Even outside of the critical question of the justice of the current situation, given the interdependence inherent to security, we know that any disproportionate vulnerability will serve as a foothold that motivated attackers will leverage against the whole.

There are many open research questions and opportunities in finding new ways to incentivize the creation and maintenance of secure hardware systems, both through technical solutions and policy frameworks. How can we ensure that hardware security solutions are providing a shared infrastructure helpful to all members of society and avoid even further compounding existing inequities through new forms of “security poverty”. What are the impacts of policy shifts and hardware security advancements not just on the economy as a whole, but on the daily experiences of the individuals that comprise it? What are the practical means by which impacts and tradeoffs can be quantified or even just qualitatively better understood across these important dimensions? Answering these

questions and exploring these opportunities will require truly interdisciplinary approaches that bring together experts not only in hardware and software security but economics, government policy, social sciences, and beyond.

We strongly recommend supporting the exploration, understanding, and measurement of the holistic impact of hardware security that is mindful of the differential impact such technologies can have, either directly or indirectly.

Summary

The need for hardware security has never been greater. Yet despite its importance and urgency, there is a lack of community focus on studying and understanding the incentives that cause hardware insecurity, let alone efforts to actually address and rectify the problems. As new technologies emerge, cyber attacks evolve, and human reliance on technology increases, we must recognize and rethink the underpinnings of hardware insecurity in our efforts to prioritize security issues and move towards safer and more efficient security systems.

The recommendations provided in this report represent just a starting point for the research and efforts that the community should be putting towards security. While we hope that the recommendations laid out in this report promote thought and discussion in the community, our greater goal is to spur action with the urgency that these problems demand. With nearly every human being vulnerable to cyberattacks, the need has never been greater for research, policymakers, economists, and the community at large to bring cybersecurity concerns to the forefront of their agenda.

Appendix

Call for Proposals

"At this workshop, participants will investigate ways to improve the design and uptake of hardware security mechanisms. In addition to looking at traditional technical solutions, the workshop will also consider new mechanisms to incentivize designers, system integrators, and users to create and maintain security of their systems. The workshop will bring together hardware and software security experts and economists and experts in devising and implementing governmental policies.

We seek short white papers to help create the agenda for the workshop and select attendees. White papers are due April 10th, you can submit them [here](#).

For participation in this workshop, we request white papers of no more than two pages. Topics of interest include, but are not limited to:

- ▶ How do current policies and market structures disincentive hardware oriented security solutions? How do we fix this: what technical and policy frameworks are necessary to make progress in this area?
- ▶ What are the mechanisms necessary to enforce a government mandate that says that X% of the performance or cost should be set aside for security? What mechanisms are necessary to determine X? How often should X be determined? Is there a quantitative approach for the organization to use up this security budget? How would this be enforced on user systems? Are there alternate government mandates that are actionable and can be supported technically?

- ▶ Is there an equitable way to proportion the benefits of security and impacts of security attacks? What hardware support, if any, is necessary to facilitate this process?
- ▶ How do we establish a chain of responsibility for malicious and negligent action while also maintaining privacy?
- ▶ How can hardware innovations (e.g. U2F tokens) fundamentally impact software dark economies?
- ▶ What incentives are necessary to patch hardware bugs in a timely manner?
- ▶ What education/certification requirements are necessary for increasing the awareness and application of hardware security solutions?
- ▶ Are there parallels to software certification requirements for hardware? What would these assurance/certification requirements look like?

Workshop organizers [Simha Sethumadhavan](#) (Columbia University) and [Tim Sherwood](#) (University of California Santa Barbara) held an orientation webinar on Thursday, January 13th, 2022 to outline the goals of the workshop and expand on what they are looking for in the white papers. A recap of the orientation can be found on the [resources tab](#)."

Workshop participants

In-Person

First Name	Last Name	Affiliation
Todd	Austin	University of Michigan
Elisa	Bertino	Purdue University
Jeremy	Blackstone	Howard University
Ryan	Craven	Office of Naval Research
Deeksha	Dangwal	Unaffiliated
Chris	Fletcher	University of Illinois Urbana-Champaign
Kevin	Fu	University of Michigan
Cat	Gill	Computing Research Association
Haley	Griffin	Computing Research Association
Adam	Hastings	Columbia University
Brian	LaMacchia	Microsoft Research
Wenke	Lee	Georgia Institute of Technology
Tamara	Lehman	University of Colorado Boulder
Steve	Lipner	SAFECode
Sharad	Malik	Princeton
Vivek	Menon	Government
Prashant	Nair	University of British Columbia
Jason	Oberg	Cycuity
Ahmad	Patooghy	North Carolina A&T State University
Gang	Qu	National Science Foundation
Jeyavijayan	Rajendran	Texas A&M
Sanjay	Rekhi	National Inst. of Standards and Tech
Paul	Rosenzweig	George Washington University
Patrick	Schaumont	Worcester Polytechnic Institute
Ann	Schwartz	Computing Research Association
Nader	Sehatbakhsh	University of California, Los Angeles
Simha	Sethumadhavan	Columbia University / Chip Scan
Timothy	Sherwood	UC Santa Barbara / Cycuity
Andrea	Stith	Northeastern University
Cynthia	Sturton	University of North Carolina
Josep	Torellas	University of Illinois Urbana-Champaign
Phil	Vachon	Bloomberg LP
Moshe	Vardi	Rice / CCC
Ashish	Venkat	University of Virginia
Guru	Venkataramani	George Washington University
Claire	Vishik	Intel
Lok	Yan	DARPA
Mengia	Yan	Massachusetts Institute of Technology

Virtual

First Name	Last Name	Affiliation
Nadya	Bliss	Arizona State University
Maddy	Hunter	Computing Research Association
Anitha	Gollamudi	Yale
James	Mickens	Harvard

Agenda

August 24, 2022 (Wednesday)

12:30 PM	Rapid Covid Tests Available Cecchi Foyer
01:00 PM	Welcome Reception with Lunch Available Whisky Terrace
02:00 PM	Welcome and Opening Remarks Cecchi Ballroom
02:30 PM	Incentivizing Cybersecurity: Paul Rosenzweig Cecchi Ballroom <i>Abstract:</i> All technological development is bottomed, in the end, on human behavior. So the key to good cybersecurity is to incentivize humans. The question is how? And the answer lies in the economics of cybersecurity. It is, mostly, a private domain with lots of externalities. Economic theory tells us that we can mitigate those externalities with taxes, subsidies or regulation. But those solutions come with their own problems. In the end, we face the challenge of an economic control structure from the horse and buggy era that needs to deal with technological developments that occur at Tesla speed.
03:00 PM	Group Discussion: Incentivizing Cybersecurity Cecchi Ballroom
03:45 PM	Reflections on Assurance: Steve Lipner Cecchi Ballroom <i>Abstract:</i> This brief presentation will introduce the problem of assurance of cybersecurity and review some of the history that led the software industry to make assurance a priority. It will then review approaches to creating and scaling processes for improving the assurance of real-world software products and services. The key finding after more than twenty years' experience is that software security assurance is similar to other attributes of product quality and is a responsibility of developers and subject to continuous improvement based on root-cause analysis of discovered problems. The presentation will discuss the issues raised by a need for certification and of product security, and wrap up with some thoughts on hardware security and the workshop questions and topics.
04:15 PM	Group Discussion: Bringing About Change Cecchi Ballroom
04:45 PM	BREAK Cecchi Foyer
05:00 PM	Group Discussion: Bringing it to Hardware Cecchi Ballroom
05:30 PM	Breakouts Cecchi Ballroom/Boardroom/Corning
06:30 PM	Report Back Cecchi Ballroom
07:30 PM	Dinner Kingbird Terrace

August 25, 2022 (Thursday)

08:00 AM	Breakfast Whisky Terrace
09:00 AM	Recap Day 1 Cecchi Ballroom
09:15 AM	Regulating Security: Success and Pitfalls: Kevin Fu Cecchi Ballroom
09:45 AM	Group Discussion: Regulations Cecchi Ballroom
10:15 AM	BREAK Cecchi Foyer
10:30 AM	Industry Panel: Complexities in Productizing Security and What Incentives Would Really Work? (Claire Vishik, Phil Vachon, Jason Oberg) Cecchi Ballroom This panel will feature a discussion on industry participants at various stages of development and will explore barriers to adoption, important real-world constraints often under-considered in research, and lessons learned.
11:30 AM	Breakouts Cecchi Ballroom/Boardroom/Corning
12:15 PM	LUNCH Whisky Terrace
01:15 PM	Breakouts Report Back Cecchi Ballroom
01:45 PM	Report Writing Cecchi Ballroom
02:45 PM	Group Discussion Cecchi Ballroom
03:15 PM	BREAK Cecchi Foyer
03:30 PM	Writing Breakouts Cecchi Ballroom/Boardroom/Corning
04:00 PM	Government Updates: Ryan Craven (ONR), Sanjay Rekhi (NIST), Vivek Menon (NRO), Gang Qu (NSF), Lok Yan (Darpa) Cecchi Ballroom
04:45 PM	Wrap up and Next Steps Cecchi Ballroom
05:00 PM	Workshop Ends Cecchi Ballroom



CCC

Computing Community Consortium
Catalyst

1828 L Street, NW, Suite 800
Washington, DC 20036
P: 202 234 2111 F: 202 667 1066
www.cra.org cccinfo@cra.org