SaTCPI '15

National Science Foundation
Secure and Trustworthy Cyberspace
Principal Investigators' Meeting (2015)

January 5–7, 2015 • Arlington, VA

Presented by usenix

Breakout Reports

7 January 2015

Breakout Participants

Nael Abu-Ghazaleh (SUNY at Binghamton)
William Adams (University of Michigan Ann Arbor)

Nicolas Christin (Carnegie-Mellon
Michael Clarkson (George Wash

Michael Gorlick (University of California Irvine)
Manimaran Govindarasu (Iowa State University)

Vern Paxson (International Computer Science Institute)
Chunyi Peng (Ohio State University)
Roberto Perdisci (University of Georgia)
Zachary Peterson (California Polytechnic State University)
Frank Pfenning (Carnegie-Mellon University)
Victor Piotrowski (National Science Foundation)
James Plusquellic (University of New Mexico)
Dmitry Ponomarev (SUNY at Binghamton)
Donald Porter (Stony Brook University)
Atul Prakash (University of Michigan Ann Arbor)
Portia Pusey
Yanjun Qi (University of Virginia)
Daji Qiao (Iowa State University)
Tal Rabin (IBM Thomas J Watson Research Center)
Mariana Raykova (SRI International)
Paul Reber (Northwestern University)
A.L. Narasimha Reddy (Texas Engineering Experiment Station)
Michael Reiter (University of North Carolina at Chapel Hill)
Kui Ren (SUNY at Buffalo)
Leonid Reyzin (Boston University)
Edward Rhyne (DHS S&T)
Golden Richard (University of New Orleans)
Heather Richter Lipford (University of North Carolina at Charlotte)
Thomas Ristenpart (University of Wisconsin-Madison)
William Robertson (Northeastern University)
Keith Ross (New York University)
Michael Rosulek (Oregon State University)
Brent Rowe (University of North Carolina at Chapel Hill)
Jerzy Rozenblit (University of Arizona)
Andrew Ruef (University of Maryland)
Norman Sadeh (Carnegie-Mellon University)
Rei Safavi-Naini (Boston University)
Jared Saia (University of New Mexico)
Lalitha Sankar (Arizona State University)
Fareena Saqib (Florida Institute of Technology)
Stefan Savage (University of California-San Diego)
Patrick Schaumont (Virginia Polytechnic Institute
Karen Schofield-Leca (Internet Society)
Dawn Schrader (Cornell University)
Stephanie Schuckers (West Virginia University R
Joseph Schwartz
Kathryn Seigfried-Spellar (University of Alabama
Ramasubramanian Sekar (Stony Brook University)
Wendy Seltzer (World Wide Web Consortium-MIT
Cyrus Shahabi (University of Southern California)
Deborah Shands (National Science Foundation)
Zhong Shao (Yale University)
Micah Sherr (Georgetown University)
Elaine Shi (University of Maryland College Park)
Zhijie Shi (University of Connecticut)
Dongwan Shin (New Mexico Institute of Mining and Technology)
Thomas Shrimpton (Portland State University)
Jordan Shropshire (University of South Alabama)

Tomas Vagoun
Jaideep Vaidya (Rutgers University Newark)
Rohit Valecha (SUNY Buffalo)
Michael Valenzuela (University of Arizona)
Jacobus Van der Merwe (University of Utah)
Kami Vaniea (Indiana University)
Eugene Vasserman (Kansas State University)
Pramode Verma (University of Oklahoma Norman
Rakesh Verma (University of Houston)
Giovanni Vigna (University of California-Santa Barbara)
Geoffrey Voelker (University of California-San Diego)
Mladen Vouk (North Carolina State University)
R Wachter (National Science Foundation)
David Walker (Princeton University)
Jesse Walker
Gang Wang (University of California-Santa Barbara)
Honggang Wang (University of Massachusetts, Dartmouth)
Hui Wang (Stevens Institute of Technology)
Jingguo Wang (University of Texas at Arlington)
Weichao Wang (University of North Carolina at Charlotte)
XiaoFeng Wang (Indiana University)
Richard Wash (Michigan State University)
Myra Washington (University of New Mexico)
Ronald Watro (BBN)
Sam Weber (Carnegie Mellon University Software Engineering Institute)
Steven Weber (Drexel University)

Joseph Kielman
Sara Kiesler (Carnegie-Mellon University)
Chris Kim (University of Minnesota-Twin Cities)
Andrew Klapper (University of Kentucky Research Foundation)
Alfred Kobsa (University of California-Irvine)
Janusz Konrad (Boston University)

Von Welch (Indiana University)
Susanne Wetzel (Stevens Institute of Technology)
Daniel Wichs (Northeastern University)
Stephen Wicker (University of California-Berkeley)
Marianne Winslett (University of Illinois at Urbana-Champaign)
Emmett Witchel (University of Texas at Austin)
Rebecca Wright (Rutgers University New Brunswick)

Craig Shue (Worcester Polytechnic Institute)
Akbar Siami Namin (Texas Tech University)
Joseph Silverman (Brown University)
Christian Skalka (University of Vermont & State Agricultural College)
Elizabeth Sklar (CUNY Brooklyn College)
Geoffrey Smith (Florida International University)
Dawn Song (University of California-Berkeley)
Anna Squicciarini (Pennsylvania State Univ University Park)
Vincent Sritapan (Department of Homeland Security)
Ankur Srivastava (University of Maryland College Park)
Christopher Stark (National Science Foundation)
David Starobinski (Boston University)
Angelos Stavrou (George Mason University)
Scott Stoller (Stony Brook University)
Koduvayur Subbalakshmi (Stevens Institute of Technology)
Gookwon Suh (Cornell University)
Kun Sun (College of William and Mary)
Yan Sun (University of Rhode Island)
Berk Sunar (Worcester Polytechnic Institute)
S. Shyam Sundar (Pennsylvania State Univ University Park)
Jakub Szefer (Yale University)
Carolyn Talcott (SRI International)

Gang Tan (Lehigh University)
Yuzhe Tang (Syracuse University)
Richard Taylor (University of California-Irvine)
Mohammad Tehranipoor (University of Connecticut)
Rahul Telang (Carnegie-Mellon University)
Stefano Tessaro (University of California-Santa Barbara)
Russell Tessier (University of Massachusetts Amherst)
Bhavani Thuraisingham (University of Texas at Dallas)
Derron Thweatt
Laura Tinnel (SRI International)
Scott Tousley (DHS S&T/CSD)
Samir Tout (Eastern Michigan University)
Ari Trachtenberg (Boston University)
Patrick Traynor (University of Florida)
Ryan Triplett
Nestan Tsiskaridze (University of Iowa)
Akhilesh Tyagi (Iowa State University)
Gary Tyson (Florida State University)
A. Selcuk Uluagac (Georgia Tech Research Corporation)
Sennur Ulukus (University of Maryland College Park)
Shambhu Upadhyaya (SUNY at Buffalo)
Salil Vadhan (Harvard University)

William Martin (Worcester
Douglas Maughan (DHS S&T)
Rio Maulana (Syracuse
Roy Maxion (Carnegie-Mellon University)
David McClure (Duke University)
Damon McCoy (George Mason University)

Breakout 1: Cryptocurrency

Elaine Shi
University of Maryland

“The Rise and Rise of Cryptocurrency”

- Bitcoin came around in 2009.
- Today, traded at $284 per bitcoin.
- Total available bitcoins: billions of dollars.
- Cryptocurrency startups: **551**
- Average evaluation: **$3.9M**
- Numerous altcoins
  - Ethereum, dodgecoin, litecoin, …
- Large online service providers have started accepting Bitcoin payments
  - Expedia, Reddit, and Overstock.com

Usage of cryptocurrency outstrips our understanding

- Various attacks observed, e.g., Mt Gox failure
- Several altcoins flawed designs exploited
- Many research papers showing attacks
  - "Selfish mining"
  - Attacks against anonymity

Therefore, it is imperative to develop a "**science of cryptocurrency**"

What is the "**science of cryptocurrency**"?

1. What are the main scientific challenges?
2. What makes this a science?

– Jeremy Epstein

1 What are the main scientific challenges?

- What makes a cryptocurrency popular? How do we model user incentives?
- How do you design a provably secure cryptocurrency? How do you even define security?
- How do you design a cryptocurrency that accommodates inspection and legal enforcement?
- How can we design technologies to help users protect themselves, e.g., not commit money to a buggy contract?
- Can we have a theoretical characterizations of possible tasks/ applications atop a blockchain-based cryptocurrency?
- How can we formally model adversarial behavior/incentives?

2 What makes this a science?

Demonstrate the generic applicability of an approach beyond a single embodiment of cryptocurrency.

What areas of research are needed for the “science of cryptocurrency”?

- **Computer Science**
  - Cryptography/security, PL, data science, formal methods, hardware, game theory, mechanism design
- **Public policy**
- **Psychology**
- **Economics** and **finance**

How can we bring communities together to make cryptocurrencies better?

Workshops that bring together researchers and the developer community

Cryptocurrency conferences/workshops with PC members from developer communities

Message for NSF

Digital money will be the way of the future: it will enable rich smart contract applications, and enable new markets and eco-systems.

- It is imperative to develop a "science of cryptocurrency"

- Cryptocurrency in the broader form
  - Not just about Bitcoin or a single cryptocurrency.
  - Related to "why this is a science" question

Breakout 2:
Social Networks and Crowdsourcing

Ben Zhao
UC Santa Barbara

The Challenge

- Security work in social networks / crowd systems has been very focused on small set of problems
  - Detection of Sybil (fake) identities
  - Detection of forged content, e.g. Yelp/Amazon reviews
- Challenge:

  Can we formulate clear research challenges in the space for the near- and long-term

1. Leveraging/Managing the Crowd

- The crowd is a powerful resource for good…
  - Can go significantly beyond state of art ML/AI systems
  - e.g. reporting phishing sites (phishtank), Sybil profile detection
  - How to incentivize/how to separate wheat from chaff
  - Can we leverage it to solve harder security problems?
- But also powerful tool for attackers…
  - "Crowdturfing" observed in multiple countries/sites
  - Malicious crowds difficult to distinguish from normal users
    - Can generate "authentic-looking" original content
    - Can launch attacks against ML classifiers
    - Easily bypass existing tools that detect scripts/automation
  - Need to develop robust defenses (adversarial ML?)

2. The Content Curation Tussle

For user-generated content, curation is a necessity

- Yet unclear how transparent providers should be in the process
- e.g. server-side black box vs. user decisions on fully-transparent data

Less Transparency

- Providers have established credibility
- Leverage access to variety of data, more powerful models, robust against Sybils/Turfing
- Simpler process addresses a need to reach broader, non-technical users

More Transparency

- Complex black boxes, e.g. reputations, can be gamed
- Transparency reduces impact of "bandwagon heuristic"
- Providers have incentives mismatch
  - More content → more users → more content …

2. The Content Curation Tussle

For user-generated content, curation is a necessity

Yet unclear how transparent providers should be in the process
e.g. server-side black box vs. user decisions on fully-transparent data

Less Transparency

- Providers have established credibility
- Le
  d
  r
- Simpler process addresses a need to reach broader, non-technical users

More Transparency

- Complex black boxes, e.g. reputations, can be gamed
- Transparency reduces impact of

ore
content …

Is there a solution that addresses both need for transparency and does not exclude less-technical users? Perhaps solutions lie in the HCI space…

3. Educating Users on OSNs

- Many users still unaware of security risks on social networks, or the tools to mitigate them
- Can we develop more effective tools that leverage the social systems themselves?
  – Can we apply tools / lessons from social psychology?
    - Challenge: establishing credibility in absence of visible pedigree
  – Tap into power of first-hand stories, or folk models
  – Can we make stories about cybersecurity *go viral*?

Breakout 3: **Cryptographic Assumptions and the Real World**

Tal Malkin
Columbia University

Matching Crypto Models to the Physical World

- Side Channel Attacks
  - Theoretical leakage and tamper resilience models vs practical attacks and countermeasures
- Theoretical Modeling and Building Secure Crypto over Vulnerable Hardware (e.g., Trojans)
- Underlying Physics: How do we model/ define/ verify what we physically need / have? and what can be done with it? E.g., :
  - Physical assumptions like Wyner wiretap model, noisy key agreement, etc
  - Physical Unclonable Functions (PUF)
  - Understanding Randomness

Basic Crypto Research (for the Real World)

- Cryptographic Complexity Assumptions
  - How do we validate assumptions / avoid working with inappropriate assumptions?
- Foundations of Symmetric Cryptography
  - Better understanding of primitives like block ciphers, hash functions, ROM
  - Weaker assumptions while maintaining efficiency
- Secure MPC
  - Why isn't it used in the real world? (are we solving the wrong problems? Wrong models? Economic considerations?)
- Power-aware cryptography
  - Minimize communication complexity, though computation also relevant.

Employing Crypto in the Real World

- IoT Key Management (e.g, medical, cars,…)
  - Issue: complex usage environment (many parties / life cycle / removing and replacing and adding devices out in the field)

- Proving Security for large systems like TLS
  - Issue: complex system / many cryptographic components

New Dimensions Beyond Current Crypto

- Security problems often due to poor implementation, misuse, and other software engineering issues, not crypto
  - where is the boundary?
- Simplicity of implementation and use
  - Often more important than just efficiency

Can Crypto help? Can we design rigorous models to address these (traditionally non-crypto) issues?

- Questioning Kerckoffs' law / Asymptotic Approach
  - Security by obscurity /increased reverse engineering
  - Better concrete security models / metrics for time/work to break a system

Meta Issues

How to incentivize researchers to do the right thing?

- More interdisciplinary research
  - Help bridge the gap to the "real world"
- More long-term research
  - E.g., work on appropriate, well studied assumptions

Possible problems:

- Do we over publish? (expect fast/many publications, quality less important?)
- Interdisciplinary research difficult (e.g., find common language), may or may not be hard to publish?
  - Suggestion: submit real-world crypto proposals to AITF

- Crypto Education

Breakout 4:
Benchmarks for Security Research

Erez Zadok
Stony Brook University

Security Benchmarking Needs

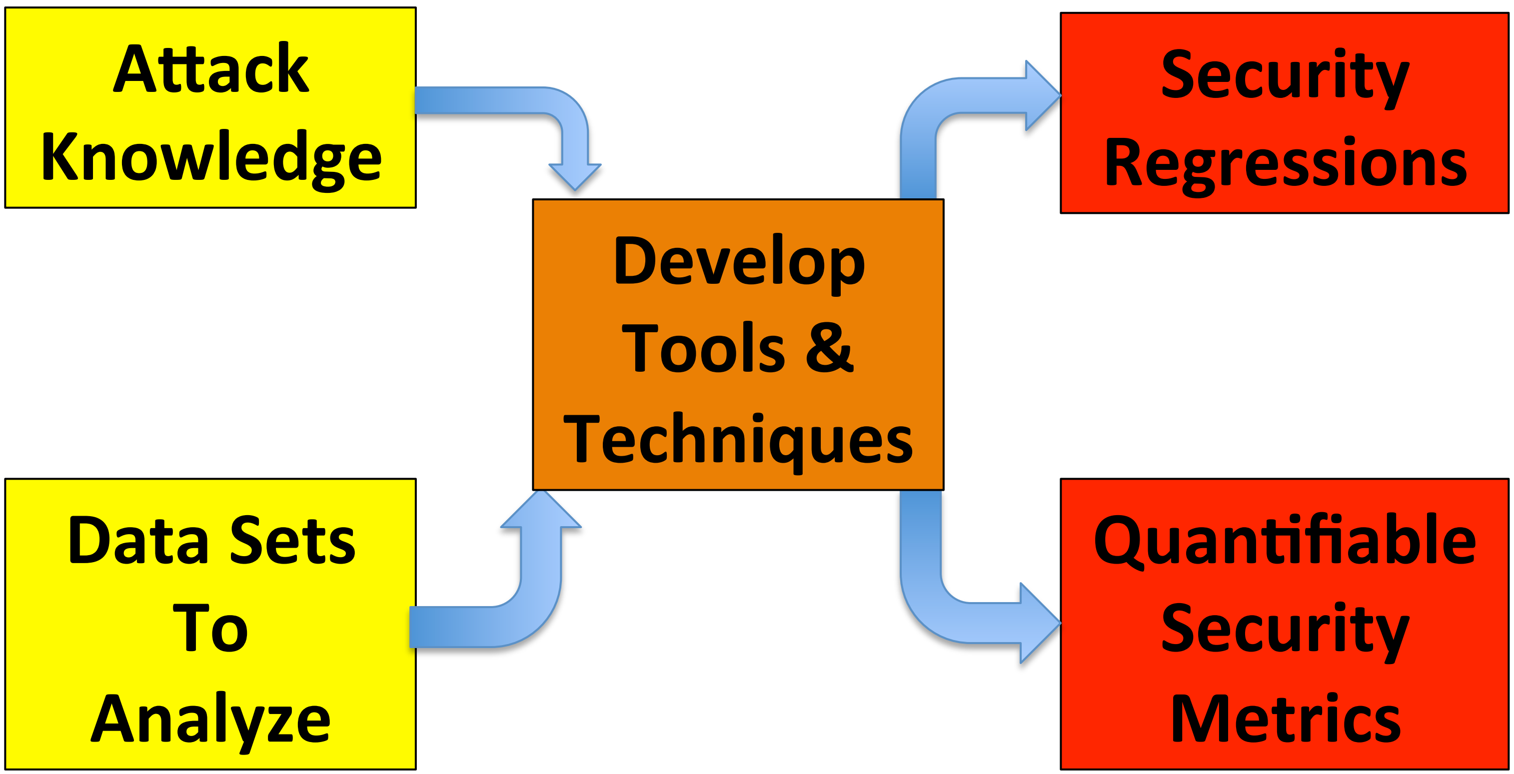

Attack Knowledge

- Need:
  - Understand basic principles
  - Comprehensive list of attacks, updated
  - Companies to disclose attack details and internals
- Understand complex interactions
  - Hardware, software, networks, people

Data Sets to Analyze

- Have:
  - WINE, CAIDA, DNS/Farsight, CRAWDAD
  - Anti-Phishing Working Group (APWG)
- Problems:
  - Old, synthetic, small
  - Overly sanitized: nearly “useless”
- Need:
  - Lots of new data
  - Minimal/configurable anonymization
  - Incentives for companies to share data
    - NSF I/UCRC model?

Security Regressions

- Have:
  - “Red” teams
  - Static code analysis (e.g., Coverity)
- Need:
  - Security vulnerability tools
    - Automated
  - Domain-specific suites
    - e.g., network routing, Web, SQL, etc.
  - Comprehensive, continually updated
  - Community effort, open/free access

Quantifiable Security Metrics

- Have:
  - Metrics for performance, energy
  - Coarse security classifications/regs (EAL1-7, SOX, HIPAA, PCI, …)
- Need metrics such as:
  - TCB size; code complexity metrics, correlate with safety
  - Time needed to break security; time to recover
  - Resources needed to break security (#machines, CPUs, etc.)
  - Number of infected systems; amount of lost data
  - $cost:
    - Price of buying attacks, cost of ransomware
    - Cost of insurance, lost revenue
- Useful combination metrics (cost functions)

Develop Tools & Techniques

- Need:
  - Inventory of existing tools & techniques
  - Identify gaps
  - Timeliness of tools/techniques key
  - Rich set of tools & techniques
  - Apply or “port” existing techniques to new threats
  - Reduce false alarms
  - Collaborate with other fields
    - e.g., ML, Prog. Lang., Verification, Viz. Analytics
    - e.g., Economics, Business, Sociology, Psychology, Medicine

To Funding Agencies

- Benchmarking is bigger Broader Impact than SaTC
- Incentives to develop/release software
- More "Transition to Practice" (TTP)
- Greater access to events (e.g., Black Hat)
- Incentives for community efforts
- Encourage in GPG/CFPs
  - NSF BRAP: Benchmarks of Realistic Scientific Application Performance(?)

Breakout 5:

Cybersecurity and the Social Sciences

Robert Axelrod
University of Michigan

Advice for Collaboration between Computer Scientists and Social Scientists

1. **Include both sides from the start.**
2. **Explicitly discuss goals and expectations** including publications and fundraising.
3. **Organize brown bags** across departments.
4. Beware that **joint PhD's have limited job prospects**.
5. Avoid **joint appointments** for Assistant Professors.

[No classified material will be shown in this breakout summary]

Breakout 6:
Responding to the NSA Revelations

Wendy Seltzer
W3C/MIT

Responding to the NSA Revelations

- Should our research change post-Snowden?
  - New or expanded topics of research
  - Changing research methods
  - Participation in public discourse

Research: Defending privacy

- Definitions and policy
- Technology and systems
- Institutions

Topics: definitions and policy

- Threat modeling: Identifying and scaling up the adversary
- Contribute to ongoing public discussion, challenge false and misleading statements
  - Demonstrate the importance of context data – it's not "just metadata"
    - Push-back on the third-party doctrine
  - Develop and publicize the more privacy-protective analytic methods we have
    - Shift the burden of proof to the information-gatherers
  - Utility-modeling
    - Small data – what we can learn from it; old-fashioned gumshoe work
- Quantifying privacy harms and risks
  - Quantifying vs. contextual?
  - Does quantifying force particular personal or policy responses? Backlash?
- Incentive alignment.
  - Not storing data might be in a business's interest
  - Industrial privacy; business trade secrecy
- User convenience, role of usability
  - Evaluation of privacy/security
  - Could there be a security label?
  - FDA (gov't) or UL (industry) model?

Topics: technology and systems

- Systems resilient against coercion/legal intervention
  - Eliminating central points of control/infiltration
    - Multi-party access control
  - “Warrant canary” transparency: “we have not yet received a request to turn over data”
    - Jurisdictional diversity?
  - Provable security
  - Secure randomness
  - Search on encrypted data
  - Exfiltration-resilient cryptography
  - Threshold crypto
  - Alternative approaches to crypto
  - Secure Multi-party computation

Topics: Institutions

- Governance: Research on norms of organizations, communication and its break-downs
  - Understanding the interactions between norms, laws, technology
  - How do new mechanisms interact with oversight?
  - Building systems to enable transparent citizen control
- Systems to enable individuals to choose/change privacy parameters (as individuals and as democratic citizens)
  - Make the costs and benefits more transparent
  - Provide meaningful choice
  - Designing good defaults

Methods

- Build in security from the beginning
  - With appropriate threat modeling, risk analysis
- Don't say “stop cryptanalysis”
- Think about protecting research subjects
  - Destroy data that's not needed
  - Secure “dark archiving” of identifying data needed for reproducible research
  - Don't expose subjects to new surveillance risks

Public involvement

- Interaction between research community and gov't agencies in setting security standards
  - Choosing experts
  - Transparent process
- Fund basic research, whatever its political valence.
  - Protection of privacy is in the national interest

Public engagement

- Public dissemination, communication, and translation of research, methodology and results
  - Demonstration of transparency best practices
  - Discussion with policy-makers
  - Interaction with tech companies
  - Participation in standards-setting
- Long-term research response

Breakout 7:

Cybersecurity Experimentation of the Future: Supporting Research for the Real World

David Balenson (SRI International)
Terry Benzel (University of Southern California)
Laura Tinnel (SRI International)

Tomorrow's Cybersecurity Challenges

- Cyberspace is rapidly evolving with nearly every aspect of society moving toward pervasive computing and networking
- Need to move quickly to meet tomorrow's needs
  - Highly specialized cyber-physical systems (CPS)
  - Interdisciplinary experimentation
  - Modeling and reasoning about human behavior
  - Advanced networking architectures (e.g., SDN
- CEF is community-based effort to study current and expected cybersecurity experimentation infrastructure, and to produce a strategic plan and roadmap for developing infrastructure that supports tomorrow's research



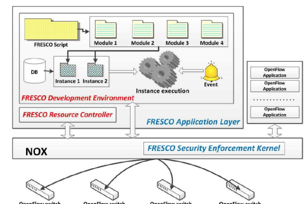

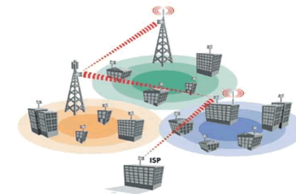

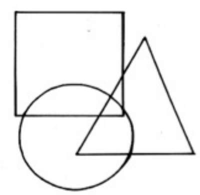
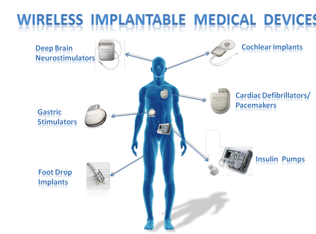

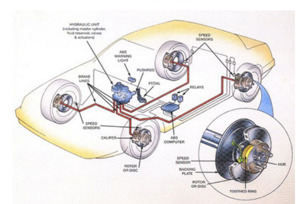


Cybersecurity Experimentation of the Future

Future Experimentation Infrastructure Objectives

- Catalyze and support research
- Advanced experimental research tools, technologies, methodologies and infrastructures
- Broadly available national resources
- Beyond today's state of the art:
  - Multi-discipline, complex, and extreme scale experimentation
  - Emerging research areas specialized cyber-physical systems and cybersecurity-relevant human behavior
- Advances in scientific methodologies, experimental processes, and education
- Strategies for dynamic and flexible experimentation across user communities and infrastructure facilities

Breakout Discussion Highlights

- Experiment metrics, including those mapped to defender objectives
- Support for internal vs. external validity of experiments, context matters – ecological validity
- Capabilities to support reproducibility
- Sharing of data collection and analysis algorithms, benchmarked datasets
- Special considerations for cyber security research
- Can't just provide tools when people don't know how to use them effectively
  - Need to couple with methodologies and education
  - Need case studies to show how the RI can be used

General RI Discussion

- Caveat: can't foresee everything needed in the future
- RI should include benchmarked data
- Can't just provide tools when people don't know how to use them effectively
  - Need to couple with methodologies and education
  - Need case studies to show how the RI can be used
- Support for experiment metrics that are mapped to defender objectives
- Recognize and support for internal vs. external validity of experiments, context matters – ecological validity

Experiment Reproducibility

- How do we describe everything needed in order to reproduce an experiment, especially in complex and/or large scale experiments?
- What level of fidelity must be captured for an experiment to be reproducible?
  - What does and doesn't matter is a research topic itself.
- When documenting an experiment that uses a complex range, need ability to point to location where the detailed info is kept.
- Bundle: data + code + environment

Sharing of Common Algorithms, Data

- Data validity can be impacted by faulty data collection methods
  - Share validated collection methods, algorithms and tools
- Shared datasets are needed to perform apples to apples comparisons between approaches
  - Share datasets for specific research areas (e.g., keystroke dynamics)
- Common analysis algorithms/tools are needed to perform apples to apples comparisons between approaches
  - Share vetted analysis algorithms/tools

Characteristics of Cyber Security

- How is RI for cyber security different from other cyber problems?
  - Must take adaptive adversaries into account – models & ability to automatically generate and validate models
  - Intent (purposeful vs. accidental) may not matter when a failure occurs until we see the behavior change

Conclusion

- Science-based experimentation infrastructure is critical to enabling future cybersecurity research
- Need for revolutionary capabilities for advancing multi-discipline, complex and extreme scale experimentation for emergent cybersecurity research areas
- Lively and helpful discussion that reinforces CEF study outputs and provides guidance on what to highlight and expound upon

- Consider: How would ***you*** contribute to a collaborative effort to build and share this infrastructure?

Breakout 8:

Developing a Principled Security Curriculum

Rebecca Wright
Rutgers University

Guiding Questions

What should a security curriculum cover?

How can we improve how security principles are taught?

Who are you teaching and **what** do they need to learn?

- Need different kinds of programs – different audiences coming in, different pathways going out.
  - Concentrations or tracks in different majors (CS, IS, etc.), stand-alone cybersecurity major
- Potential interest in different kinds of career paths.
- Different principles suitable for different groups.
- Some philosophical questions still unresolved:
  - Is practicing offense necessary for understanding defense, or is offense its own specialized skill?
- Pragmatic concerns and constraints
  - Overfilled curricula, long pre-requisite sequences, students of varying backgrounds, etc.

Many Existing Useful Resources

- **NIST NICE Framework**
- National Academies Report: *Professionalizing the Nation's Cybersecurity Workforce*
- **NSA/DHS Academic Centers of Excellence:** now divided to cyber defense and cyber operations (smaller program, specialized on offense). Includes existing knowledge units.
- Military academies developing **"Cyber Science"** as a starting point separate from CS.

  Working group of about 60 people (mostly in cybersecurity) working with ABET to develop an ABET-accredited program.
- Various courses, including some with materials or entire course available freely online.
- Many more…

Principles, **Practice**, and **Mindset**

- **Scientific** principles, **engineering** principles, and **social science** principles, among others.
- Effective to combine principles with **practical activities** and **examples** that illustrate the principles, build interest, and encourage engagement.
- In the context of a broad education (vs. training for specific skills), focus in a discipline can serve as a way to **develop a mindset**, a **culture**, and a **body of shared knowledge**. (Should also ensure teaching of problem solving, communication, and critical thinking.)
- We could do a better job of explaining the differences between different kinds of programs to potential students: what background do you need to succeed in this programs? what kinds of career or further educational pathways are natural from this program? what kinds of interests are a good fit for this program? [But beware being too narrow and scaring people off.]

Breakout 9:

User Authentication

Nicolas Christin
Carnegie Mellon University

Passwords & authentication

- Simple, cross-platform, one-size-fits-all for human-to-machine authentication
  - We'll probably still talk about passwords in a few years
- Historically, poor usability of alternatives (e.g., biometrics)

- **This may be changing**
  - Commoditization of usable biometric systems (e.g., iPhone touch ID)
  - Increased importance of machine-to-machine authentication (Internet of Things)
    - RFIDs/NFC tokens are now extremely cheap to produce and are increasingly deployed (you're using one to open your room)
  - Single-sign on systems (e.g., Google/FB accounts) are increasingly used for credential delegation
  - Multi-factor authentication

Future research directions in user authentication (1/2)

- **Privacy-preserving** authentication
  - Group signatures / pseudo-identities for large systems (e.g., transportation networks)
    - Research question example: how to scale group signatures (expensive to verify) so that they can accommodate very large networks (e.g., automobile networks)
    - Potential communication overhead to disseminate pseudo-identities
- **Reconciling threat models** with deployed primitives
  - e.g., "authenticating" to the newspaper
  - Segmentation of authentication primitives
- Potential **arms race**
  - Well known in biometrics (research on spoofing)
  - Is there an end to this arms race – can it be proven?

Future research directions in user authentication (2/2)

- Incentives to **decouple identification** from **authentication**
  - Identity providers/SSO systems – avoiding core root of trust (multiparty computation?)
  - How to decouple? Preserving privacy vs. long-term "reputation"
  - How much trust are users willing to give to authentication providers?
    - E.g., failure to accept the German National ID card
- **Metrics** to evaluate authentication
  - Going beyond false negative/false positive rates
  - Scope of the threat model, adoption rate, usability / lightweight, cost, failure implications
- Deployment of **forward secrecy**
  - Technology probably already exists but needs to be deployed to a much larger extent

Breakout 10:

An End to (Silly) Vulnerabilities

Matthew Might
University of Utah
matt.might.net
@mattmight

Research

Education

Incentives

silly vulnerability. *n*.

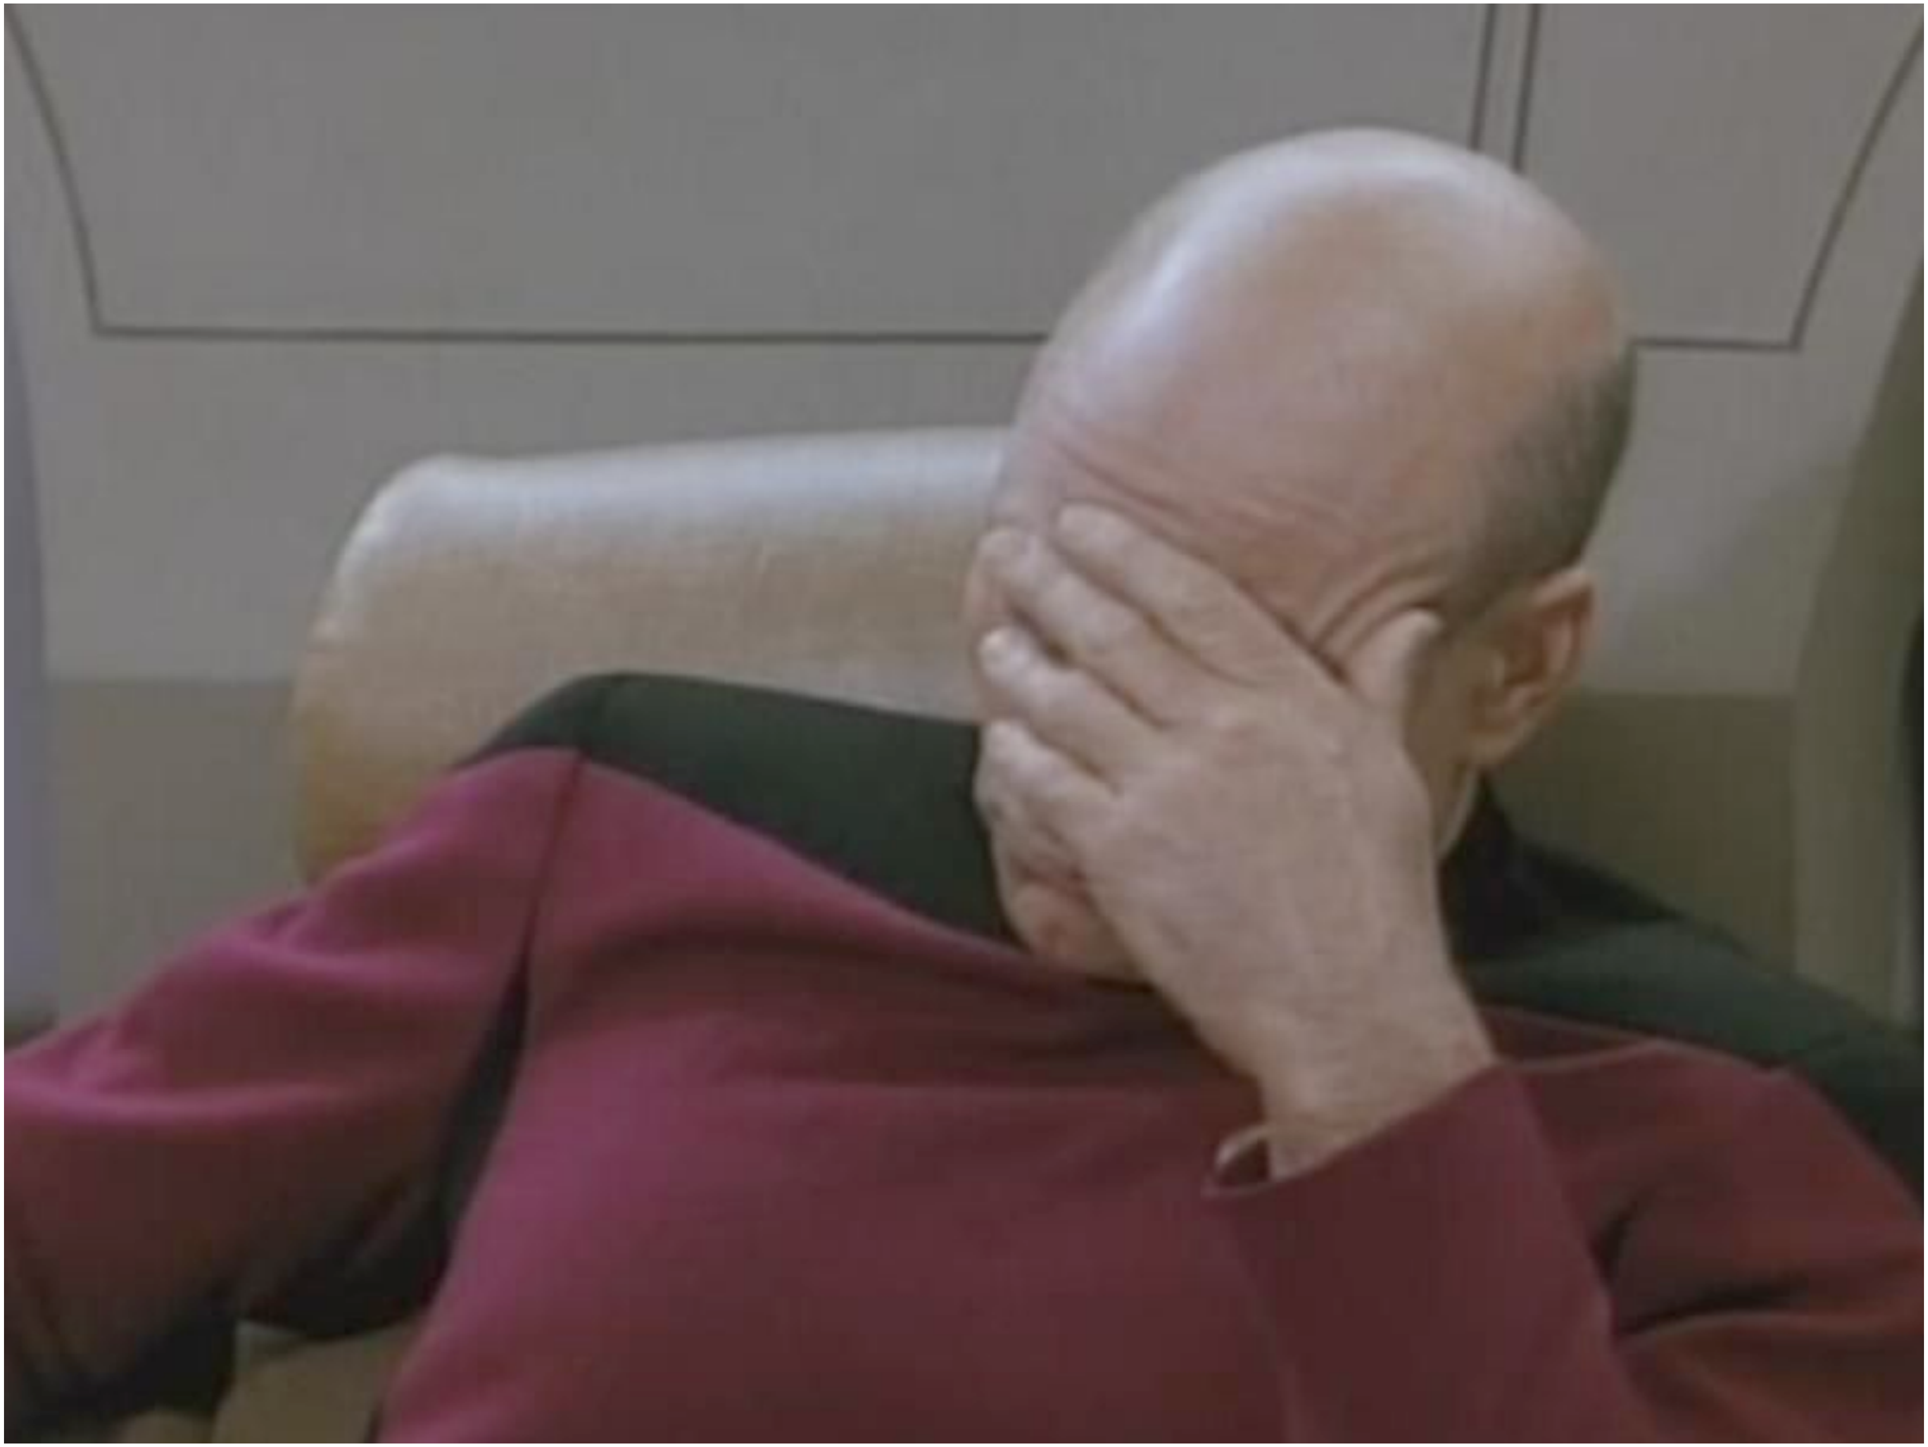

All vulnerabilities are silly!

2014


KEEP
CALM
AND
goto fail;
goto fail;

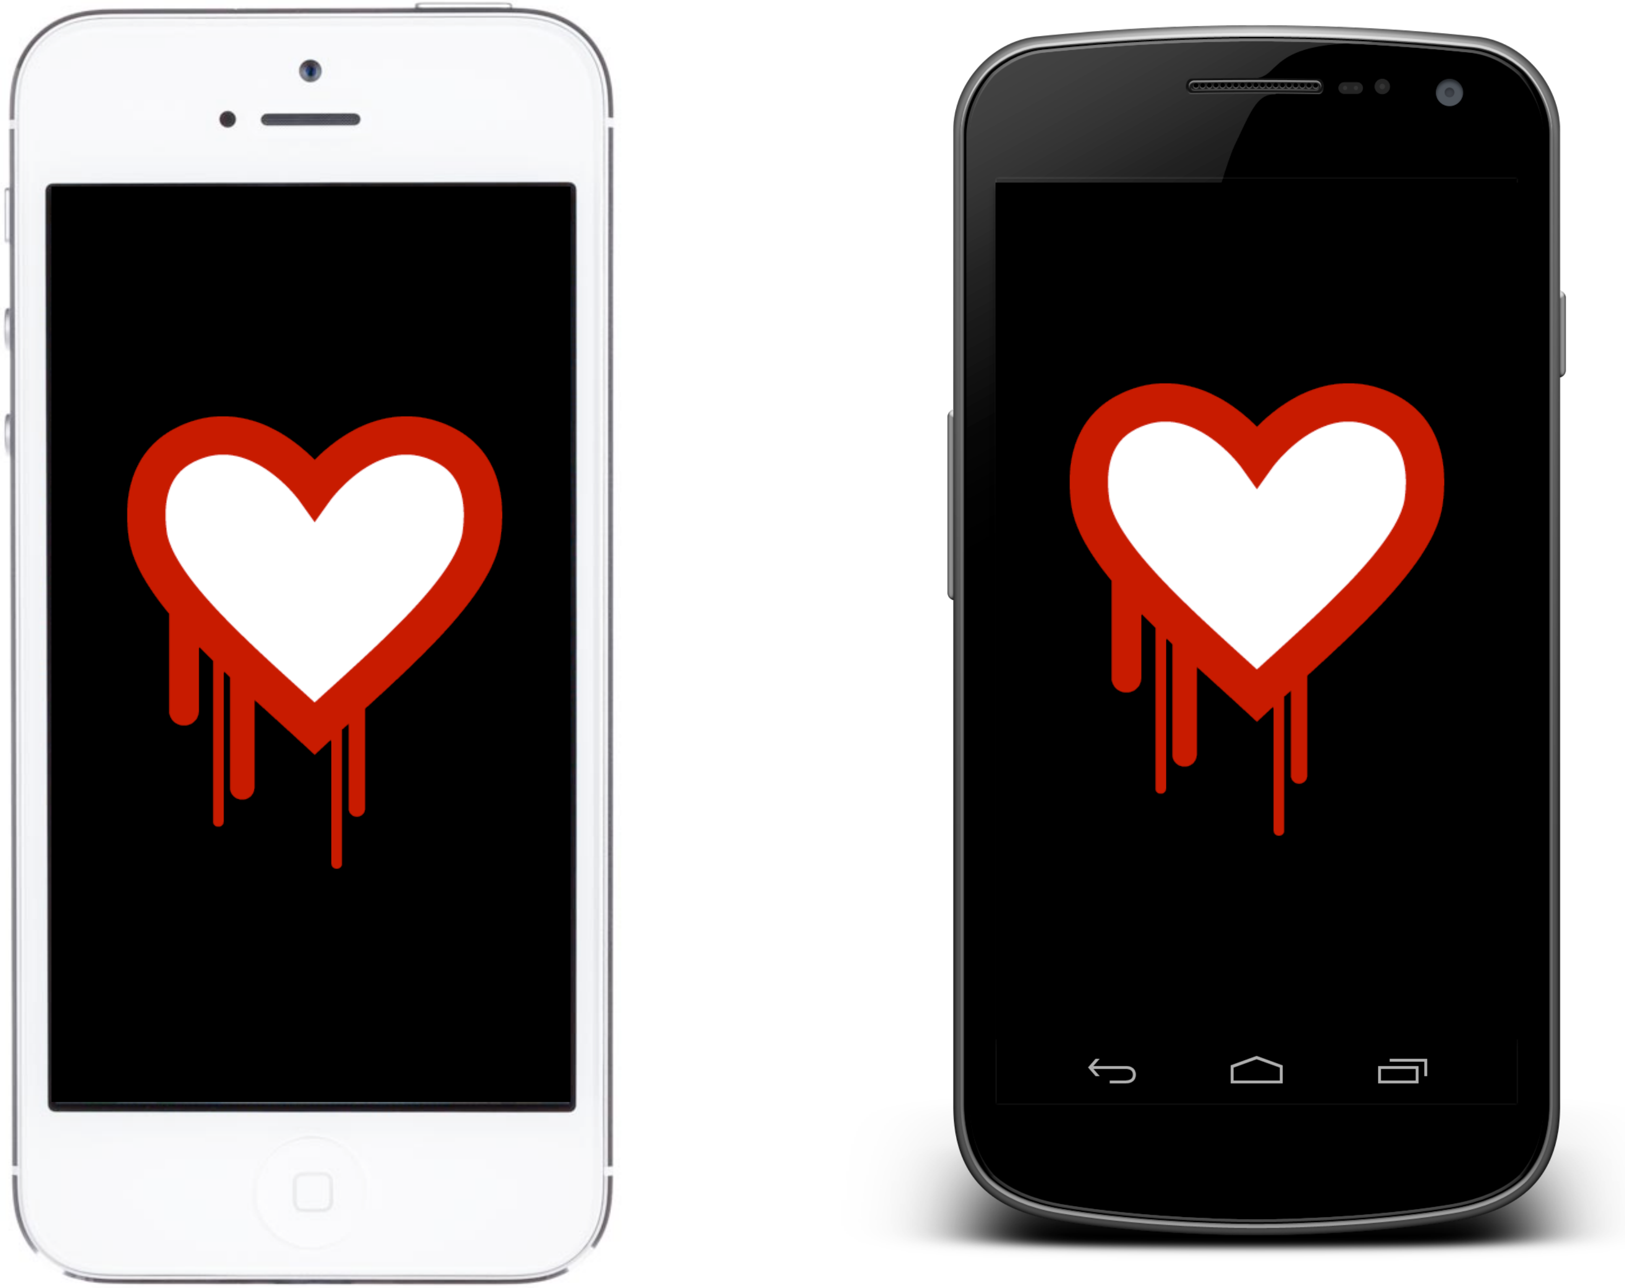


bash
Shellshock


goto fail;
goto fail;


I
SSL

$1 billion

Proposed Resolution

No further advances in research and education are necessary.

It's up to you, industry.

No further advances in research and education are necessary.

It’s up to you, users.

ΔResearch

Static analysis

Spectrum of silliness

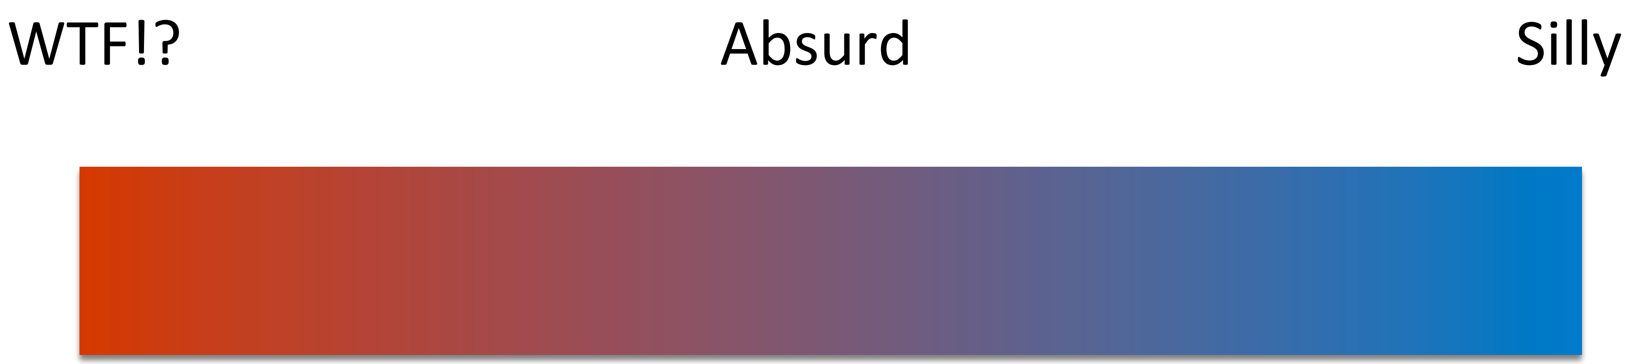

Spectrum of silliness

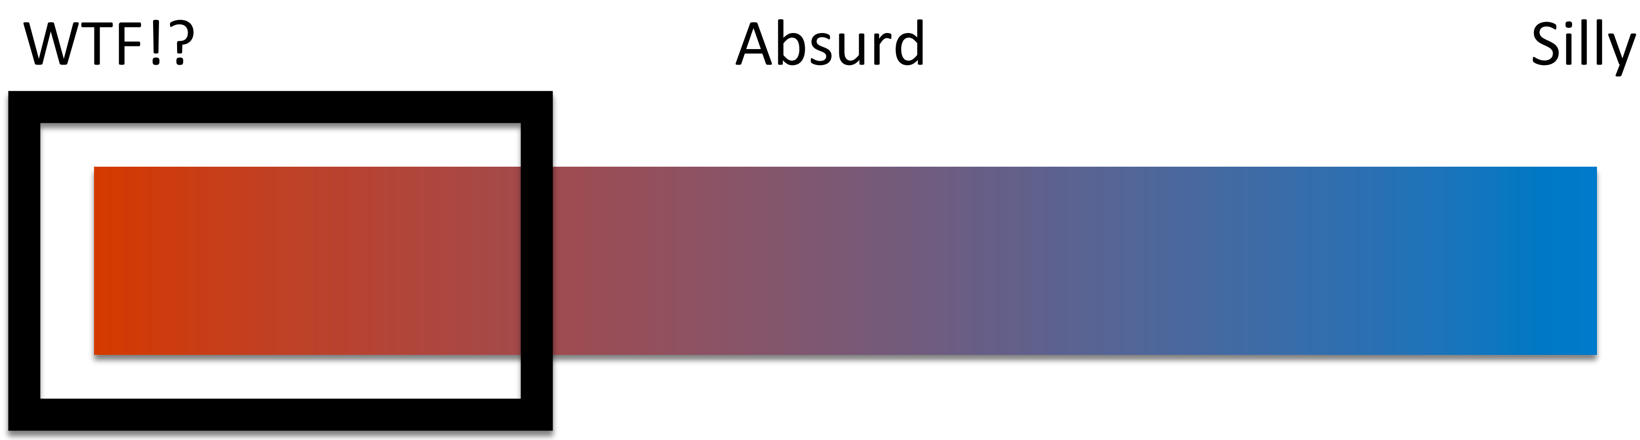

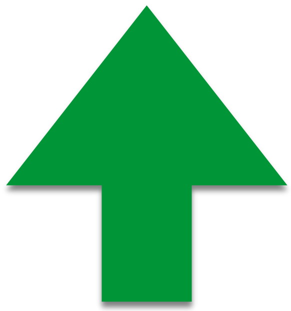

Usability

Scalability

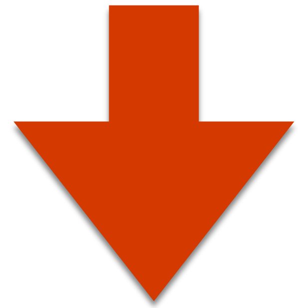

False neg.

False pos.

coverity®

FORTIFY®
SOFTWARE

VERACODE

PARASOFT®

GRAMMATECH

klocwork®

Formal methods

Scalability

Domain expert

Cost

Languages

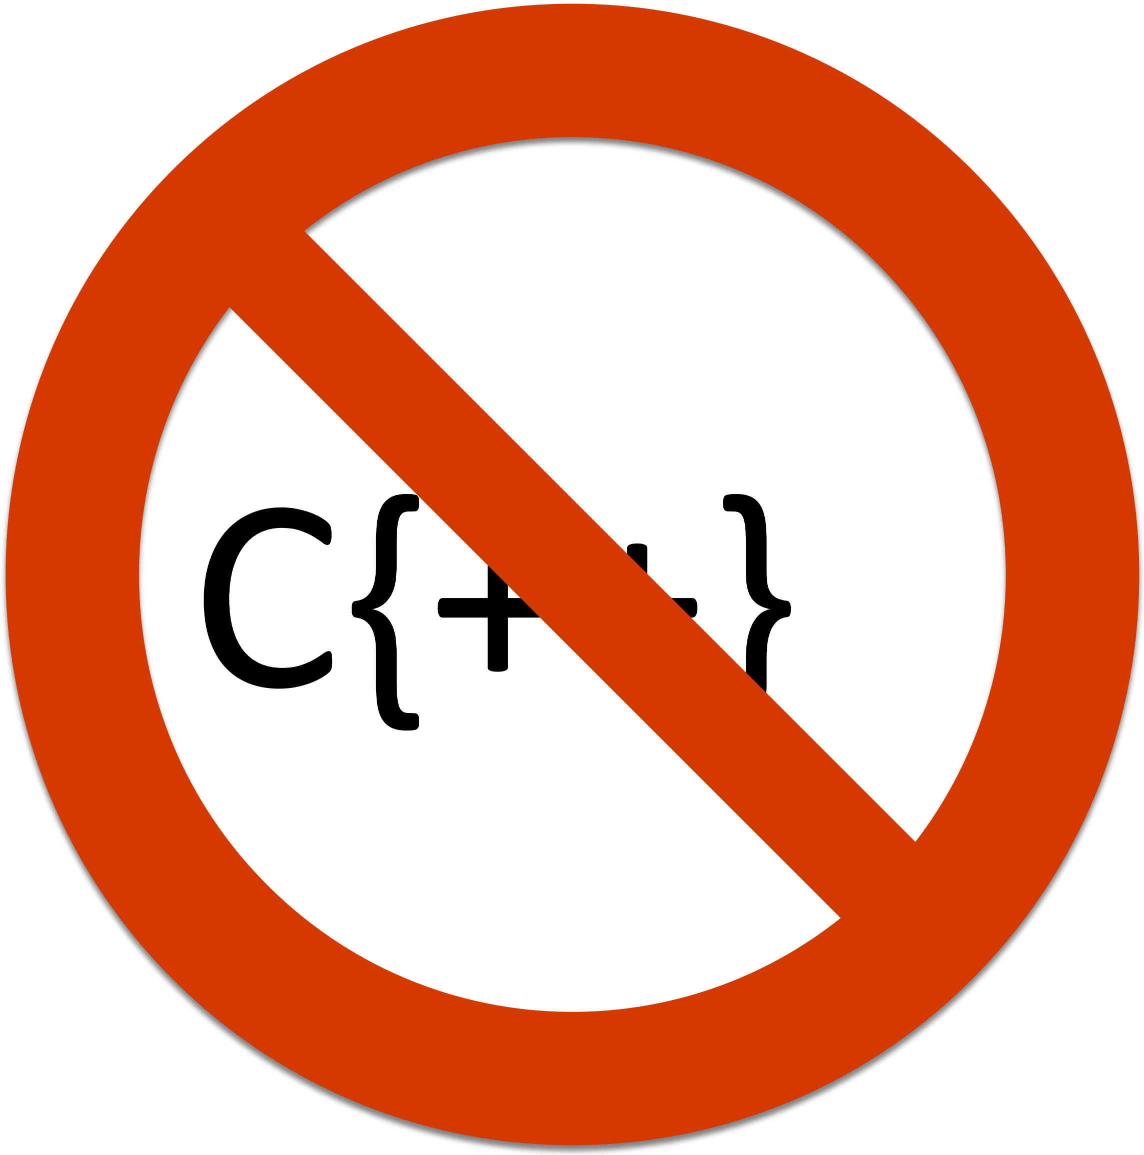
C{++}

ΔEducation

Cross-cutting & Standalone

Security from the start

ΔIncentives

Cyber Ralph Nader

Civil liability for software

Much less vulnerabilities.

Much less software.

Thanks!

Breakout 11:
Human Factors

Damon McCoy
George Mason University

Cyber Insurance

- Deal with security problem by purchasing insurance
- Problem is there is insufficient data to model risk
- “actuary tables” for cyber security would be useful
- Understanding distribution of payouts

Incentivizing Users

- Maybe we could pay users $5 dollars to do X and improve their security
- Problem is we don't know what X should be
- Need better understanding of what effects security outcomes

Teachable Moments

- Warning notices that explain why purchasing from spam is harmful
  - Display at the moment the user is about to visit merchant site

- Does notification work encourage remediation
  - What can be done to improve the effectiveness?

Breakout 12: **Architecture**

Ruby Lee (Princeton University)
Gookwon (Ed) Suh (Cornell University)

Starting questions

1) What are the best opportunities today for architecture-focused security research?
2) What problems in hardware, software and network security can best be addressed by architectural changes or new architecture?
3) How should smartphone, IoT and cloud computing servers be designed to improve cyber security?
4) How should researchers in different domains collaborate with architecture researchers on security problems?
5) What are the application domains where "architecture support for security" can make the most impact?
6) What are the challenges and opportunities in designing and building hardware architecture that we can trust?

Discussion Topic and Direction

- What are the best opportunities for architecture-focused security research?
- The term “architecture” was broadly defined
  - HW, SW, network architecture
- The discussion was focused more on opportunities for *hardware* architecture to enhance security
- HW has both strengths and weaknesses
  - Strengths: 1) real-time, 2) difficult to bypass, 3) difficult to tamper with, 4) performance, energy efficiency
  - Weaknesses: 1) sematic gap, 2) difficult to fix
  - What are the right set of hardware security primitives?

Architecture Research Needs

- Hardware to guarantee critical security and privacy properties even when software layers are compromised, especially for safety-critical applications
- Threat models and security requirements for emerging application spaces such as smartphone, cloud, IoT, CPS, etc.
  - Rethink existing hardware security architecture
- Hardware design methodology and assurance
  - Improve both security and performance
  - Tools and metrics to verify the security of hardware-software designs
  - Tools and platform support to build custom secure architecture
- Facilitate tight interdisciplinary collaborations
  - HW architecture and security communities
- Common infrastructure for security architecture research
  - Open-source SoC HW, security benchmarks, and attack suites

More Research Directions

- How to secure complex heterogeneous SoCs?
  - Many processing elements, untrusted IPs
- How to provide end-to-end security including humans and communications
  - Secure I/O and user interfaces
- How to leverage parallel resources in many-core processors for security?
- What's the implications of emerging nanotechnologies for security? How do we leverage them for security?
- How to authenticate hardware?

Breakout 13:
Cloud Security

Srini Devadas
MIT

Questions

- What does it mean for a cloud to be secure?
- How do we resolve conflicts between security, availability, user convenience and performance?
- How do we minimize the Trusted Computing Base (TCB) of a secure cloud?

Interesting Research Directions (by no means complete!)

- Track dissemination and processing of private data
  - present to user in an intuitive way
- Efficient Verifiable computation
- Obfuscated computation (to protect program as well as data)
- Hybrid of cryptographic and systems approaches to cloud security
- Security across users in a cloud
- Enhance the security of commercial offerings, e.g., Intel SGX
- Resolving the conflict between obfuscated computation and protecting cloud from obfuscated malicious code

Community-Building Challenge

Clean-Slate design of a secure public cloud

- In two different settings: infrastructure as a service and platform as a service
- Different TCBs and threat models
- Clean-slate secure processor designs
  - Verified and untrusted hypervisor
  - Untrusted OS
- Exemplar software stack and applications

Breakout 14:

Machine Learning

Mingyan Liu
University of Michigan

Machine Learning Applied to Cyber Security: Risks, Opportunities & Future Directions

- The necessity and use of domain expertise
  - Choosing the right domain with the right scope, framing the right problem
  - Beware of overuse and superficial use

- Adversarial ML
  - Robust against manipulation intended to evade ML-based detection
  - Caution against speculative threat models

Machine Learning Applied to Cyber Security: Risks, Opportunities & Future Directions

- Impact of ML on privacy
  - ML techniques help us infer and detect as defenders
  - The same capability in the hands of attackers exacerbates privacy issues

- Focusing on explanation in addition to pursuing performance
  - An opportunity for both the ML and security communities

- Collecting and maintaining high quality data
  - Lack of ground truth
  - Highly dynamic environment

Breakout 15:
App Markets

Ninghui Li (Purdue University)
Somesh Jha (University of Wisconsin)

Challenges

- **Users:** Regular users need to make security-critical decisions
  - How to reduce reliance on users for security while serve diverse individual needs?
- **Extensible resources:**
  - Sensors that are close to users
  - OS lacks ability to protect new types of resources
- **Analysis:** imprecision of analysis and of definition of malicious behavior
- **Fragmentation of app markets**

Ecosystem and App Market

- Needs governance structure, incentives for app markets to promote security
- Create a ecosystem that creates incentives for using less permissions/personal info
- Create economic liability for posting malware
- Need more robust reputation systems for both apps and reviewers/reviews, to detect malware as well as malicious promotion
- Division of responsibility between market and client devices

Towards Better Apps

- “Hygiene rules” for appropriate use of personal information in app
  - Perhaps with certification and verifiable
  - New programming language helping this?
  - Crypto help balance need for code analysis/ verification and prevention of reverse engineering
- More flexible permission model
  - Context-aware, time-limited grant
  - Hide complexity from users
- Can new hardware features help?

Breakout 16:

Securing the Web for Everyone

Roxana Geambasu
Columbia University

Breakout 17:
Cyber-Physical Systems

Stephane Lafortune
University of Michigan

Breakout 17: Securing CPS (1/4)

- 20 participants from academia, industry, government
- Cyber-Physical vs Cyber vs Internet of Things: where to draw the lines?
  - All CPS have sensors and actuators
  - Control (feedback) loops
  - Physical variables: laws of physics, inertia, time
  - Physical consequences of improper behavior: safety, graceful degradation, recovery

Breakout 17: Securing CPS (2/4)

- Find aspects that have analogs in cyber systems
  - Draw parallels with Network Security
- Find aspects that do not have analogs in cyber systems and have research value
  - Both defender and attacker are limited by the laws of physics
- Control theory, real-time and embedded systems
  - Model of physical process; well-defined specifications
  - But: Attacker is not “just” a “disturbance”: adversarial models
  - Role of humans in-the-loop (more or less?)

Breakout 17: Securing CPS (3/4)

- Attacker may be trying to inflict damage or to acquire IP
  - Authentication of components is a critical issue
- Intrusion Detection, Isolation, Recovery
  - Exploit sensor redundancy and physical model
- Importance of timeliness
- Diversity of systems
  - From: Critical infrastructure: power/water/communications/transportation
  - To: Interconnected (bio-)medical devices

Breakout 17: Securing CPS (4/4)

- Security is still an after-thought, even now. What can we do as academics?
  - Need a taxonomy of potential vulnerabilities
  - Vulnerability assessment; quantify impact
  - What-if analyses
  - Identify similarities (with cyber systems) and distinguishing features
  - Scalability of solutions proposed
- Privacy in CPS: domain specific
  - Whose privacy: user, operator, suppliers?

Breakout 18:

Cybersecurity Competitions

Portia Pusey
Edrportia@google.com
Cybersecurity Competition Federation

Opportunities

Technologists to partner with **Competition Developers**

- Test and learn new technologies
- Solve real world problem
- Data sets

Competition Developers and/or **Technologists** to collaborate with **Researchers** in social, behavioral, and economic sciences

- Bake measurement into competition development
- Recommend predictive instruments
- Identify outcomes for players and stakeholders
- Benchmark current characteristics of competitors and competitions
- Produce instruments and tools to evaluate/assess outcomes for within and between competition comparisons

Competition Developers to support **Educators**

- Performance-based assessments for performance outcomes
- Used challenges/puzzles/walkthroughs become instructional materials and labs

Shameless Plugs

NSF Cyber Education/Competition Activities

IseRink.org

Competition environment & virtual laboratory: networking, cyber security, and penetration testing

HandsOnSecurity.org

Materials for teaching cybersecurity

CyberFed.org

A community to communicate, promote and advocate for cybersecurity competitions and related activities

USENIX 2015 '3GSE

Lunch

These slides, and some extras not shown, will be posted on conference site.

Extra Slides
(for posting, not presenting)

SATC PI Meeting 2015

Breakout 4

Benchmarking for Security Research

Erez Zadok (Stony Brook University)

Opening Presentation Slides

Problem

- How to quantify security accurately?
- How to compare security systems fairly?
- What research needs to be sponsored?
- What is benchmarking?
  - Metrics?
  - Test suites for validation?
    - More attainable

What can we Measure Today?

- Evaluate single metrics easily:
  - Performance: e.g., ops/sec
  - Energy: e.g., joules
- Some metrics are harder to evaluate:
  - Reliability(?)
- Challenging to combine metrics:
  - Ops per joule-second, energy-delay
  - How meaningful?

Measuring Security is Hard

- Lots of regulations: SOX, HIPAA, PCI, etc.
  - Qualified guidelines, not easily quantifiable
- Evaluation Assurance Levels: EAL1-EAL7
  - A coarse classification
- How to measure a negative?
  - The absence of a rarely(?) occurring problem
- Take a cue from insurance industry?
  - Risk assessment

Metrics? (part 1)

- Prevention:
  - “How much effort/resources your adversary willing to put in?” -Blaze c. 90s
- Speed:
  - How many “mips” you need to breach a system within time T?
- How many infected computers?
- How much data is lost?
- How much time to recover?

Metrics? (part 2)

- Dollars? Complex cost functions?
  - Need to involve economists
- Risk: how much $$$ invested vs. $$$ lost in case of breach
  - Insurance: pay premium, get payoff in case of disaster
  - Today: we pay for security service/software, but no "payoff" in case of breach
    - There is often quantifiable $$$ lost due to breach
    - How much $$$ ransomware asks vs. paid?
- Is the metric linear or perhaps a power low?
  - Do we need a Richter-like log scale

Metrics? (part 3)

- Social engineering:
  - How many gallons of water[boarding] ☺

Raw Notes Taken During Breakout

Test Suites

- Easier to develop?
- Is a 'red-team' a test suite?
- Security s/w vs. "internet" security?
  - E.g., BGP hijacking
- How to update suites for future attacks?
- Some tools exist, but may not cover all attacks
  - E.g., Coverity, formal verifiers
- Need an inventory of existing tools vs. domains
  - Then identify gaps

Test suites 2

- Many papers exist describing problems
  - Software for these papers?
- Level of security may depend on environment
  - Programming language and system deployed on
- Are suites to verify security, or provide metrics?
- Tools for security testing (regressions)
- Tools for security metrics

Test suites 3

- Before we can develop tools, need to know principles and agree on them
  - Number of implemented principles
  - List of attacks
  - Lack of data to analyze, due to privacy
    - Companies won't tell you their internals
- Some attacks are particular to hardware/sw
  - Need to simulate for newer environments
  - Before you invest too much in new h/w+s/w

Test suite 4

- Lack of automation in test suites
- Misaligned with "research agendas"
  - Incentive to publish the first attack
  - Follow on work/implementation lacking
  - Grad students need to graduate
  - Need a community effort?
- How to "port" attacks to new environments
  - And prove they "work"

Test suites 5

- Metric: TCB size?
- Code complexity metrics?
  – Correlate with code security?
- Verification: tests against known models
  – Security: try to verify the absence of problems
- Problems in common libraries
- Where do we learn about attacks?
  – Black Hat charges $$$$

Test suites 6

- Some business provide insurance
  - Risk analysis: extreme value analysis?
  - Who's the attacker and their capabilities?
- Metrics customized for specific areas
- ML
  - Combine ML with (adversarial) game theory
  - To better deal with 0-day attacks
  - Need to reduce false alarms

Test suites 7

- Evaluate the price of buying attacks
  - E.g., hypervisor attacks cost a lot
- Incentives to develop software for attacks
  - How timely does it need to be to be useful
  - How to make research more valuable in long run
- How to automate and scale attacks
- Common data sets and tools that "everyone" uses?

Test Suites 8

- Predict: network data
  - Real, not synthetic data
  - How much to sanitize the data so it's still useful
- WINE (Symantec)
  - Conduct study in "protected" environments
  - We want "custom" data sets
- CAIDA data set, networking - free
- DNS data set by Farsight? Paid
- CRAWDAD data set
- Incentives for companies to share data and see others'
  - I/UCRC model?

Broader Impacts

- Dev. Tools is big BI (NSF)
- NSF "benchmarking" program: mention
- Updated NSF GPG to encourage tools
  - For more than SaTC
- Digital privacy can protect parts of data sets

Proposed 4-minute Summary (Wednesday 2015-01-07 @ 11:00am)

SATC PI Meeting 2015

Breakout 4

Benchmarking for Security Research

A Summary

Erez Zadok (Stony Brook University)

Security Benchmarking Needs

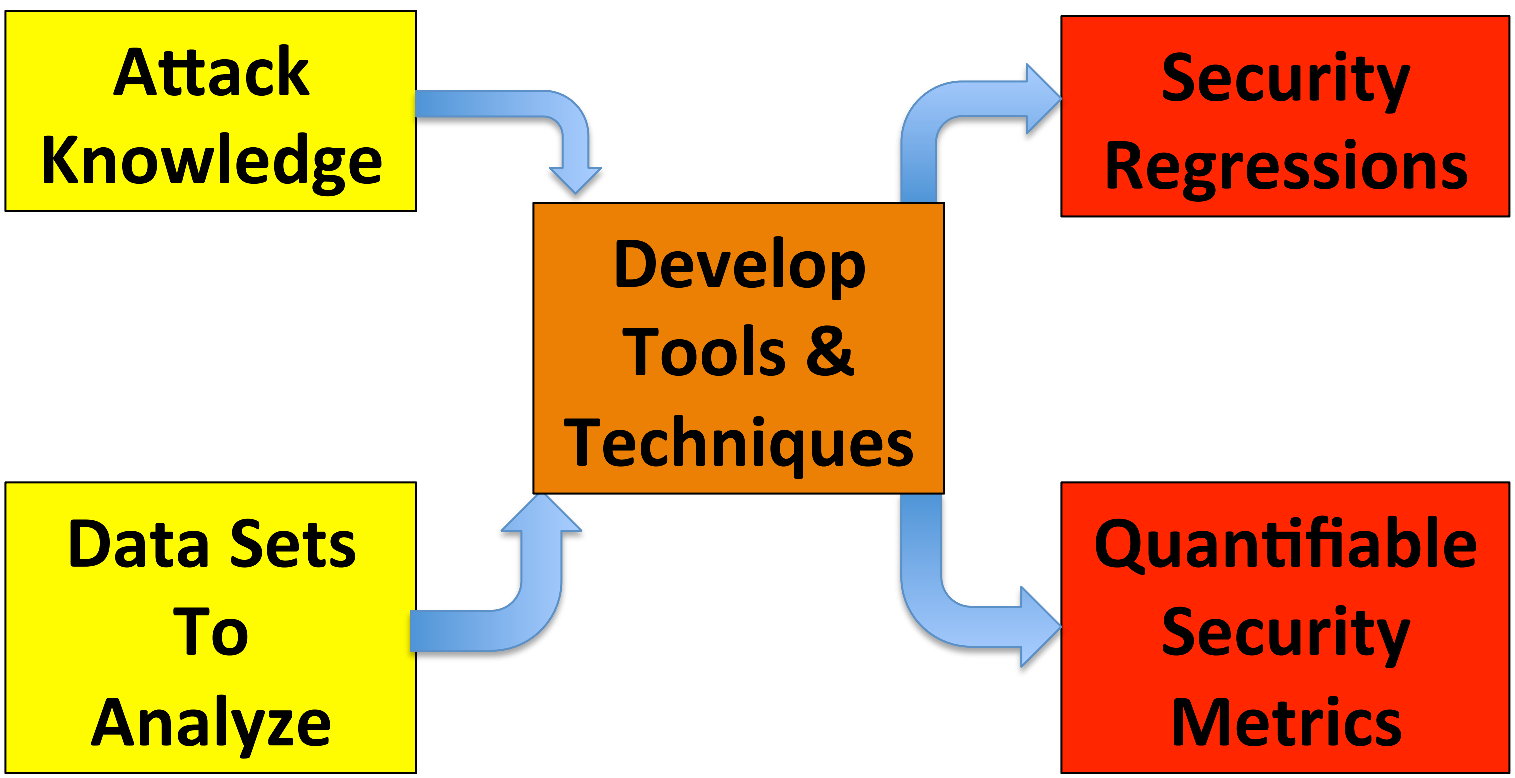

Attack Knowledge

- Need:
  - Understand basic principles
  - Comprehensive list of attacks, updated
  - Companies to disclose attack details and internals
- Understand complex interactions
  - Hardware, software, networks, people

Data Sets to Analyze

- Have:
  - WINE, CAIDA, DNS/Farsight, CRAWDAD
  - Anti-Phishing Working Group (APWG)
- Problems:
  - Old, synthetic, small
  - Overly sanitized: nearly "useless"
- Need:
  - Lots of new data
  - Minimal/configurable anonymization
  - Incentives for companies to share data
    - NSF I/UCRC model?

Security Regressions

- Have:
  - "Red" teams
  - Static code analysis (e.g., Coverity)
- Need:
  - Security vulnerability tools
    - Automated
  - Domain-specific suites
    - e.g., network routing, Web, SQL, etc.
  - Comprehensive, continually updated
  - Community effort, open/free access

Quantifiable Security Metrics

- Have:
  - Metrics for performance, energy
  - Coarse security classifications/regs (e.g., EAL1-7, SOX, HIPAA, PCI)
- Problems: Hard to compare tools/techniques meaningfully
- Need metrics such as:
  - TCB size; code complexity metrics, correlate with safety
  - Time needed to break security; time to recover
  - Resources needed to break security (#machines, CPUs, etc.)
  - Number of infected systems; amount of lost data
  - $cost:
    - Price of buying attacks, cost of ransomware
    - Cost of insurance, lost revenue
- Useful combination metrics (cost functions)

Develop Tools & Techniques

- Need:
  - Inventory of existing tools & techniques
  - Identify gaps
  - Timeliness of tools/techniques key
  - Rich set of tools & techniques
  - Apply or "port" existing techniques to new threats
  - Reduce false alarms
  - Collaborate with other fields
    - e.g., ML, Prog. Lang., Verification, Viz. Analytics
    - e.g., Economics, Business, Sociology, Psychology, Medicine

To Funding Agencies

- Benchmarking is bigger Broader Impact than SaTC
- Incentives to develop/release software
- More “Transition to Practice” (TTP)
- Greater access to events (e.g., Black Hat)
- Incentives for community efforts
- Encourage in GPG/CFPs
  - NSF BRAP: Benchmarks of Realistic Scientific Application Performance(?)

Breakout Group Report
#15 App Market

Discussion Leads:
Somesh Jha (Wisconsin)
Ninghui Li (Purdue)

Members of Group

- Craig Shue (WPI)
- Heng Yin (Syracuse)
- Gary T. Leavens (U. Central Florida)
- R. Sekar (Stonybrook)
- Guofei Gu (Texas A&M)
- Yan Chen (Northwestern)
- Richard Taylor (UC Irvine)
- Gang Wang (UCSB)
- Mengjun Xie (U. Arkansas Little Rock)
- Ari Trachtenberg (Boston U)
- Ron Watro (BBN)
- Yan Sun (U. Rhode Island)

Existing Work Group Members Found Interesting

- Taintdroid (Penn State)
- Baseband attack (Weinman)
- Sparta (Ernst)
- Malware genome project (Jiang, NC State)
- CHEX (Lu & NECLab)
- EpiCC
- AppSealer
- User-driven access control (U. Washington)

Challenge: Users

- Regular users need to make security-critical decisions, e.g., downloading apps
- Need to understand what users really want in terms of security/privacy
  - Perhaps a moving target
- How to reduce reliance on users for security while serve diverse individual needs?
- Needs models of security that users can understand
  - E.g., switching between multiple modes.

Challenges in Analysis

- Fragmentation of Android systems
  - Tens of thousands of variants, often updated
  - Defense mechanisms difficult to be work across platforms
- Inaccuracy from program analysis
- Difficult to determine whether behavior is malicious, depending on user expectation
- Security problems may be due to third-party ads that come with apps. More systematic approach to deal with ads management and security

Challenges: Extensible Resources

- Current mobile platform security model is broken at multiple levels
  - OS level, lack ability to protect new types of resources that are added to mobile platforms
  - User level, needs context-depend decisions from users; current system unable to effectively obtain such decisions
- Large variety of sensors that are close to users
  - More private/personal information
  - Potential for leakage and for enhancing security

Permission Model

- Two current models: Android is installation-time; iOS is usage time (ask once)
- Needs more flexible permission model.
  - Context-aware, time-limited grant of permission
- Need to communicate security/risk information to users in the right way, and asks right questions that they can answer
- Need to balance more powerful control at lower level without exposing the complexity to users.

Ecosystem

- Needs governance structure for app markets to promote security
- Create a ecosystem that creates incentives for using less permission, e.g., enable searching for apps without certain permissions
- Economic incentive/liability for malicious apps
  - How about developers need to post bond to put apps on market?
  - Can attribution be done in a legally valid way?

App Market Design

- iOS uses centralized app market, meaning one set of tools for analyzing apps, creating central point of failure.
- Android has more centralized market.
- Which model is better for security?
- Need more robust reputation systems for both apps and reviewers/reviews, to detect malware as well as malicious promotion

Market and Users

- What is the right division of responsibility for security/privacy between the app store and the client side?
  - App store does static analysis. Client side follow up.
  - Client sends apps to cloud for analysis.
- Use crowdsourcing to collect information about app and communicate to users.
  - How to have a device provide useful feedback regarding an app without compromising privacy?

Developer Involvement

- What constraints can be placed on developers for tradeoff of security, openness?
- Since it is hard to prove maliciousness, perhaps instead "hygiene rules" for good practices for using personal information.
- "Certified Good Behavior" apps?
  - Ways to specify hygiene rules that give required expressive power; e.g., once obtaining location, don't hold it;
  - Certification can be verified

Developer Involvement (continued)

- Are users willing to pay extra for such certified apps? Perhaps government can play a role in creating such a market?
- Would another programming language/paradigm help verifying hygiene rules?
- Developers have incentive to prevent reverse engineering, obfuscate compiled programs
  - Can crypto help balance prevention of reverse engineering and ability to verify (by market place who has the right key)?

Misc Topics

- Defense against baseband attack
  - Low-level library code needs to be vetted
- Cellular botnets for denial of service attacks against cell phone infrastructure
  - Attacks on home registration registrar
- Benchmark for attack and defense research
- Can new hardware features help improve security upstream?
  - Can help attribution, information flow tracking
  - Some are needed by Samsung KNOX

Applicability to Other Platforms

- Can knowledge/lessons learned here extend to other situations?
- Yes !?
  - Desktop computing
  - Software-defined networking
  - Internet as things