Breakout Participants

Nicolas Christin (Carnegie-Mello
Michael Clarkson (George Wash
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 Car
Thomas Ristenpart (University of Wisconsin-Ma
William Robertson (Northeastern University)
Keith Ross (New York University)
Michael Rosulek (Oregon State University)
Brent Rowe (University of North Carolina at Cha
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 Dieg
Patrick Schaumont (Virginia Polytechnic Institut
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 Universit
Wendy Seltzer (World Wide Web Consortium-M
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)

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

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 Norma
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
Steven Weber (Drexel University)

Joseph Kielman
Sara Kiesler (Carnegie-Mellon University)
Chris Kim (University of M
Andrew Klapper (Universit
Alfred Kobsa (University of
Janusz Konrad (Boston Un
David Ko
Farinaz K
Shriram
Ram Kris
Marwan
Brenden
Sandeep
Sandip K
Stephan
Brent La
Carl Lan
Catherin
Sarah La
Gary Lea

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)

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)

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 (Worce
Douglas Maughan (DHS
Rio Maulana (Syracuse
Roy Maxion (Carnegie
David McClure (Duke
Damon McCoy (George Mason University)

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

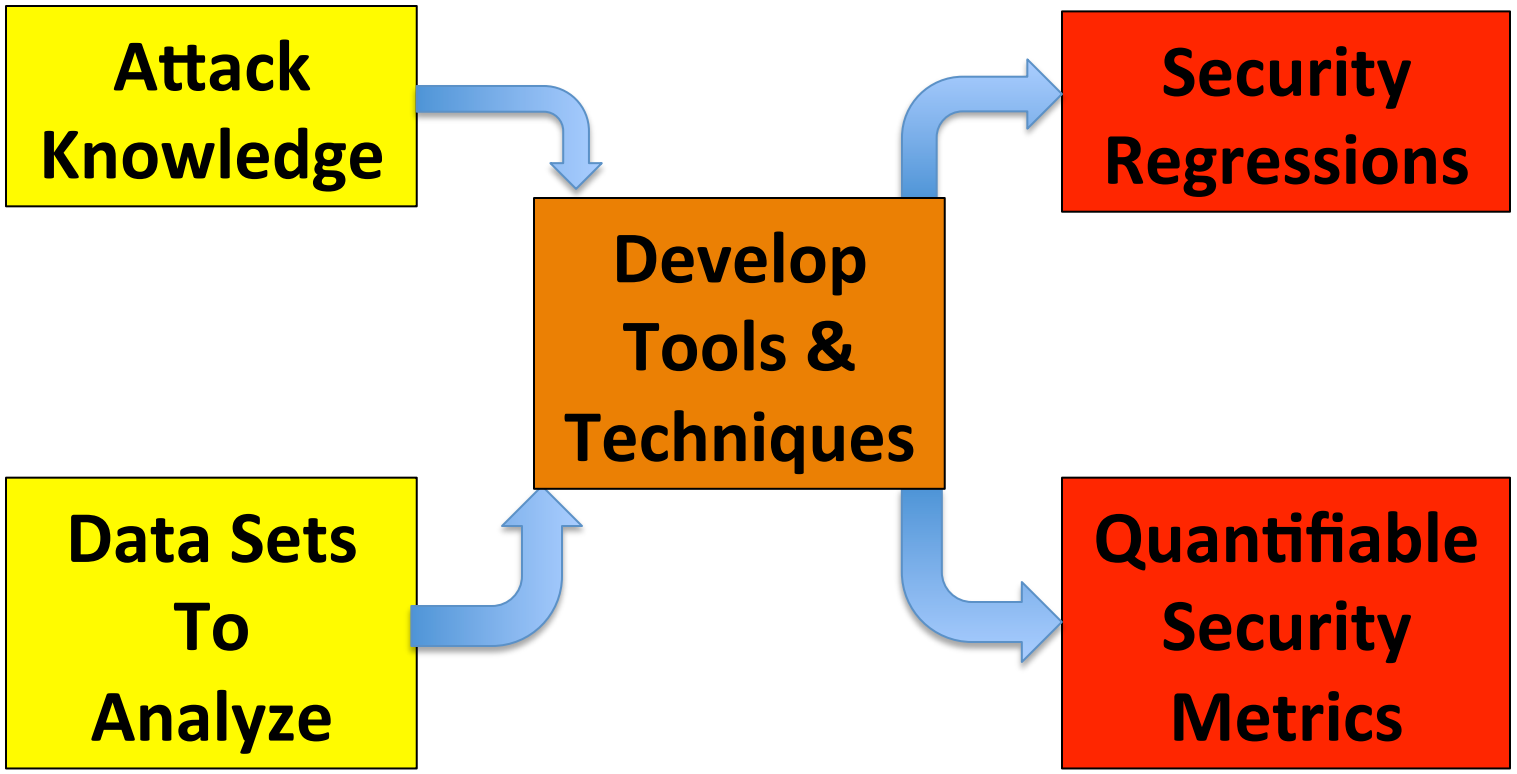

Attack Knowledge

- 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