The Importance of Measurement and Decision Making to a Science of Security

ACM Symposium and Bootcamp on the Science of Security
Champaign/Urbana, IL – April 22, 2015

Prof. Patrick McDaniel

Alternate Title: Protection Amulets, Drunk Driving, Cubits, and the Magic Eight-Ball

ACM Symposium and Bootcamp on the Science of Security
Champaign/Urbana, IL – April 22, 2015

Prof. Patrick McDaniel

Cyber Security Collaborative Research Alliance

PENNSTATE

Technical Approach:

Trans-disciplinary; Emphasis on understanding human attackers-defenders-users; Experimentation to validate models

Impact: Create fundamental understanding of cyber science encompassing risk, agility, detection and the underlying human dynamics

PI Expertise: Cyber-security, systems, theory, human factors, psychology, networking

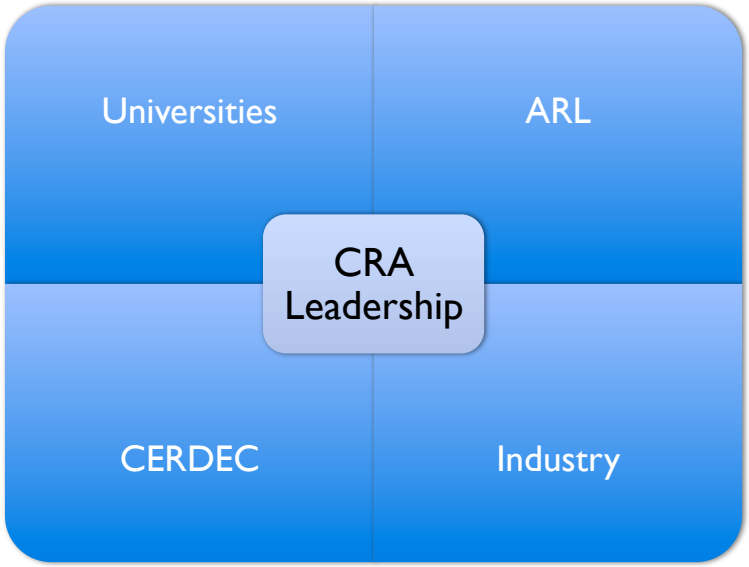

Teaming:

Collaborative teams co-led by PIs from government, academic and industry partner organizations

Accelerate transition to practice via close partnering with SMEs at ARL and CERDEC

PENNSTATE 1855 ARL INDIANA Carnegie Mellon University UCDAVIS UNIVERSITY OF CALIFORNIA ★CERDEC US ARMY – RDECOM APPLIED COMMUNICATION SCIENCES UC RIVERSIDE UNIVERSITY OF CALIFORNIA

CRA Vision

- Goal: develop a *rigorous and fundamental science of cyber-security* that

  (a) detects the risks and attacks present in an environment

  (b) understand/predict users, defenders, and attacker action

  (c) alter the environment to securely achieve maximal mission success rates at the lowest resource cost.

- Outcome: dictate and control the evolution of cyber-missions and adversarial actions

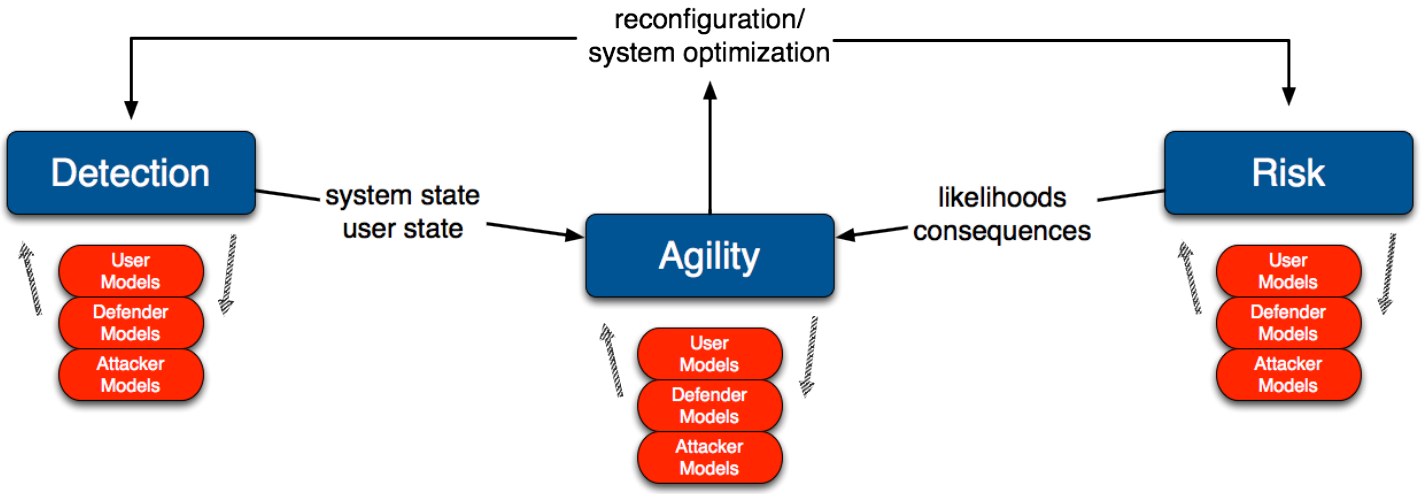

Protection from Auto Accidents

"Amulet for destroying unwelcome forces and providing protection from accidents." ($7.50 US)

Why don't we wear amulets?

- ~~**Because they don't work.**~~

No, because we don't have any way of knowing that they do work?

" … we are alarmed at the kind of subversive untruths that vendor 'spin doctors' are using to draw well-intentioned customers to their doors." (Paul Vixie – April 20, 2015)

Science of auto-accidents

- A science of predicting and diagnosing auto accidents
  - Took 50-75 years to evolve
  - Applied fundamental science from different disciplines
    - Physics, engineering, psychology, physiology, …

- Science predicts under what conditions accidents are more/ most likely to occur …
- … and importantly develop models for optimal conditions to avoid accidents.

- Top reasons for accidents*:
  1. Inattention (e.g., texting)
  2. Speeding/Reckless Driving
  3. Driver Fatigue
  4. Drunk Driving
  5. Auto Defects
  6. Weather

*National Motor Vehicle Crash Causation Survey, National Highway Traffic Safety Administration, July 2008

A Pessimistic View …

PENNSTATE 1855

- Many outside of the security community want a science to predict whether a specific system will be compromised ... which is *impossible*.
- In all likelihood it is *probably impossible* to ascertain whether any general-purpose computing systems is compromise-able.
- What can we hope to accomplish?
  - Probabilistically identify where compromise is likely.
  - Identify modifications to the system/environment that will reduce the likelihood of compromise.

Why not?

PENNSTATE 1855

- One could reasonably argue that our expectations for a science of security are misplaced ..

  Current computer science, engineering, physics and mathematical models *can't model even a tiny fraction of the complexity* of modern systems and environments.

  The same with automotive safety engineers: modeling cars, traffic, electrical and kinetic systems, weather, drivers, … is impossible.

And more …

- And this is where security must depart from traditional computer science …

  The very tools we use to reduce our problem domains to make analysis tractable (abstract modeling) invalidate the result …

  … the intellectually pleasing micro-worlds we create formally ultimately divorce of us from hard problems of understanding how the different layers of complex systems interact.

And more …

- And this is where security must depart from traditional computer science …

Disclaimer: I am not suggesting that continuing work on micro-world security (crypto, formal models, security-typing, etc.) is not fruitful or necessary, but it is unlikely to lead to the pervasively secure environments many hope to achieve (any time soon).

hard problems of understanding how the different layers of complex systems interact.

Now what?

PENNSTATE 1855

- Practical security is a nearly infinite Gordian Knot of non-stationary, invisible, and often unknowable states and forces …

- … in short, practical security is messy (un-model-able).

Security is an optimization.

PENNSTATE 1855

- Definition: agility is a *reasoned* modification to the system in response to a functional, performance or security need.
  - Idea: the system would better address the needs of the environment in some other configuration (or not).
  - Reasoning: requires a means of reasoning about the probable outcomes of a transition to new state

- Decision-making: a modern science of security should evaluates the environment to determine what changes to make to optimize outcomes (positive and negative).

A science of security has to begin with measurement.

Security State Transitions

- System consists of states (S)

$$S = s_1, s_2, \ldots, s_k$$

- Agility maneuvers (A) produce state transitions

$$s_1 \rightarrow^A s_2$$

- Every state has a utility function (U)

$$U : S \rightarrow \mathcal{R}$$

- A state transition makes sense if:

$$U(s_1) \leq U(s_2)$$

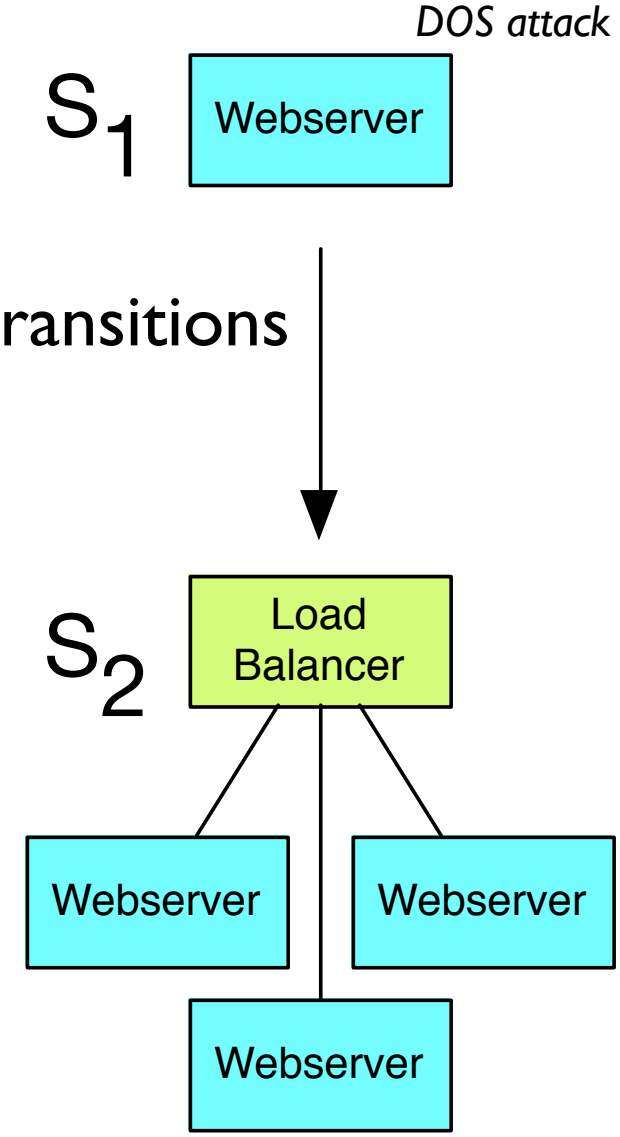

Security as a decision problem

- Find the maneuver (s) that best optimizes the utility of the system:

$$s = max(U(s_i)), s_i \in S$$

- Utility is not only about security ...

Measuring “utility”

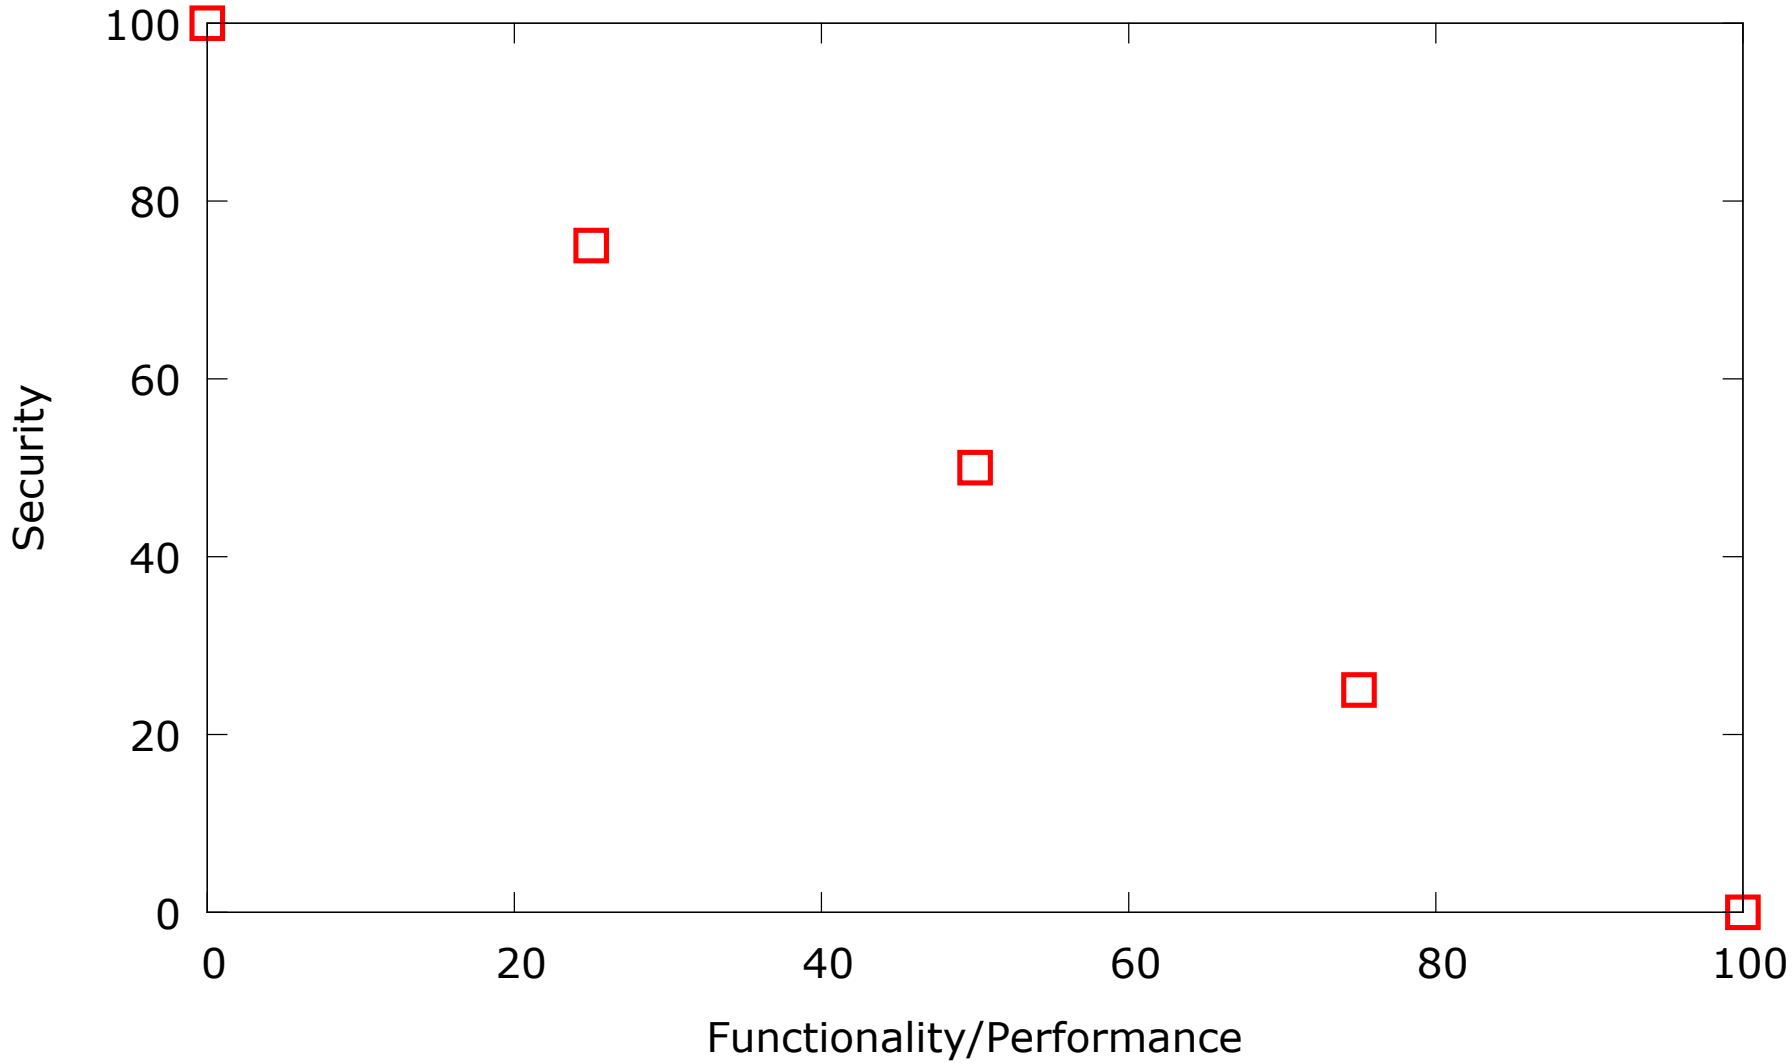

Measuring “utility”

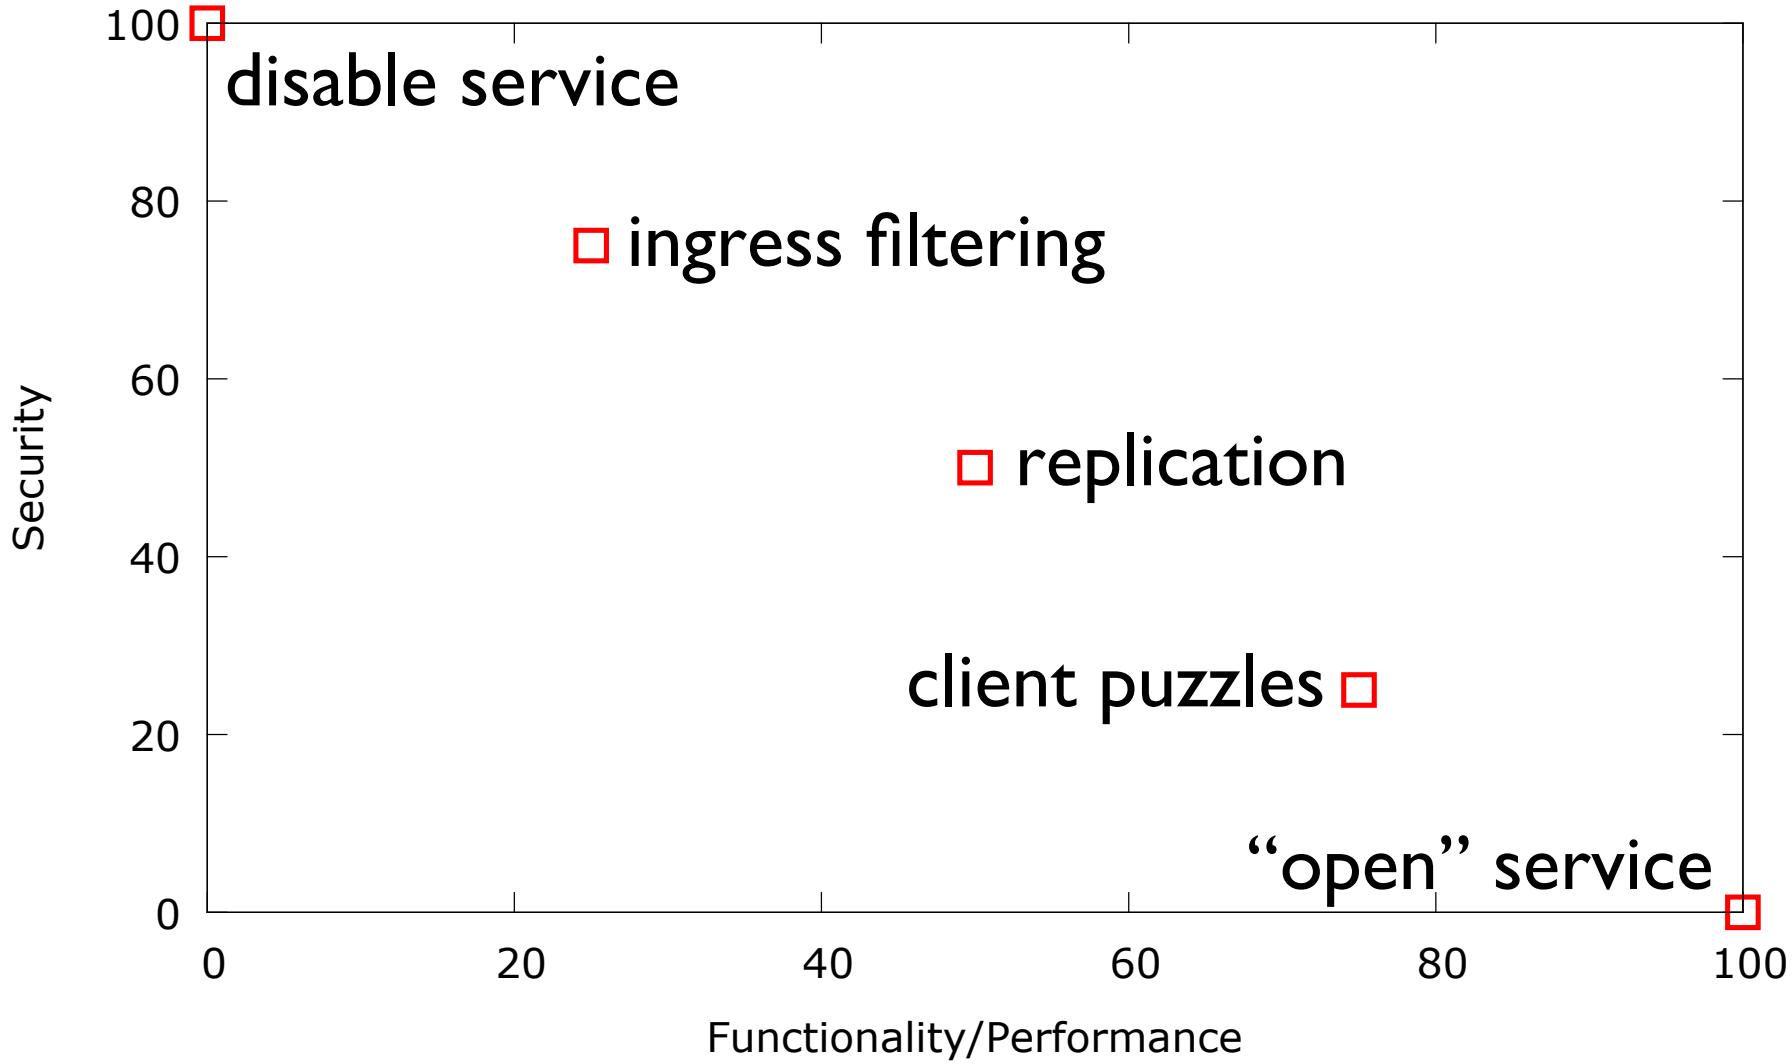

Measuring "utility"

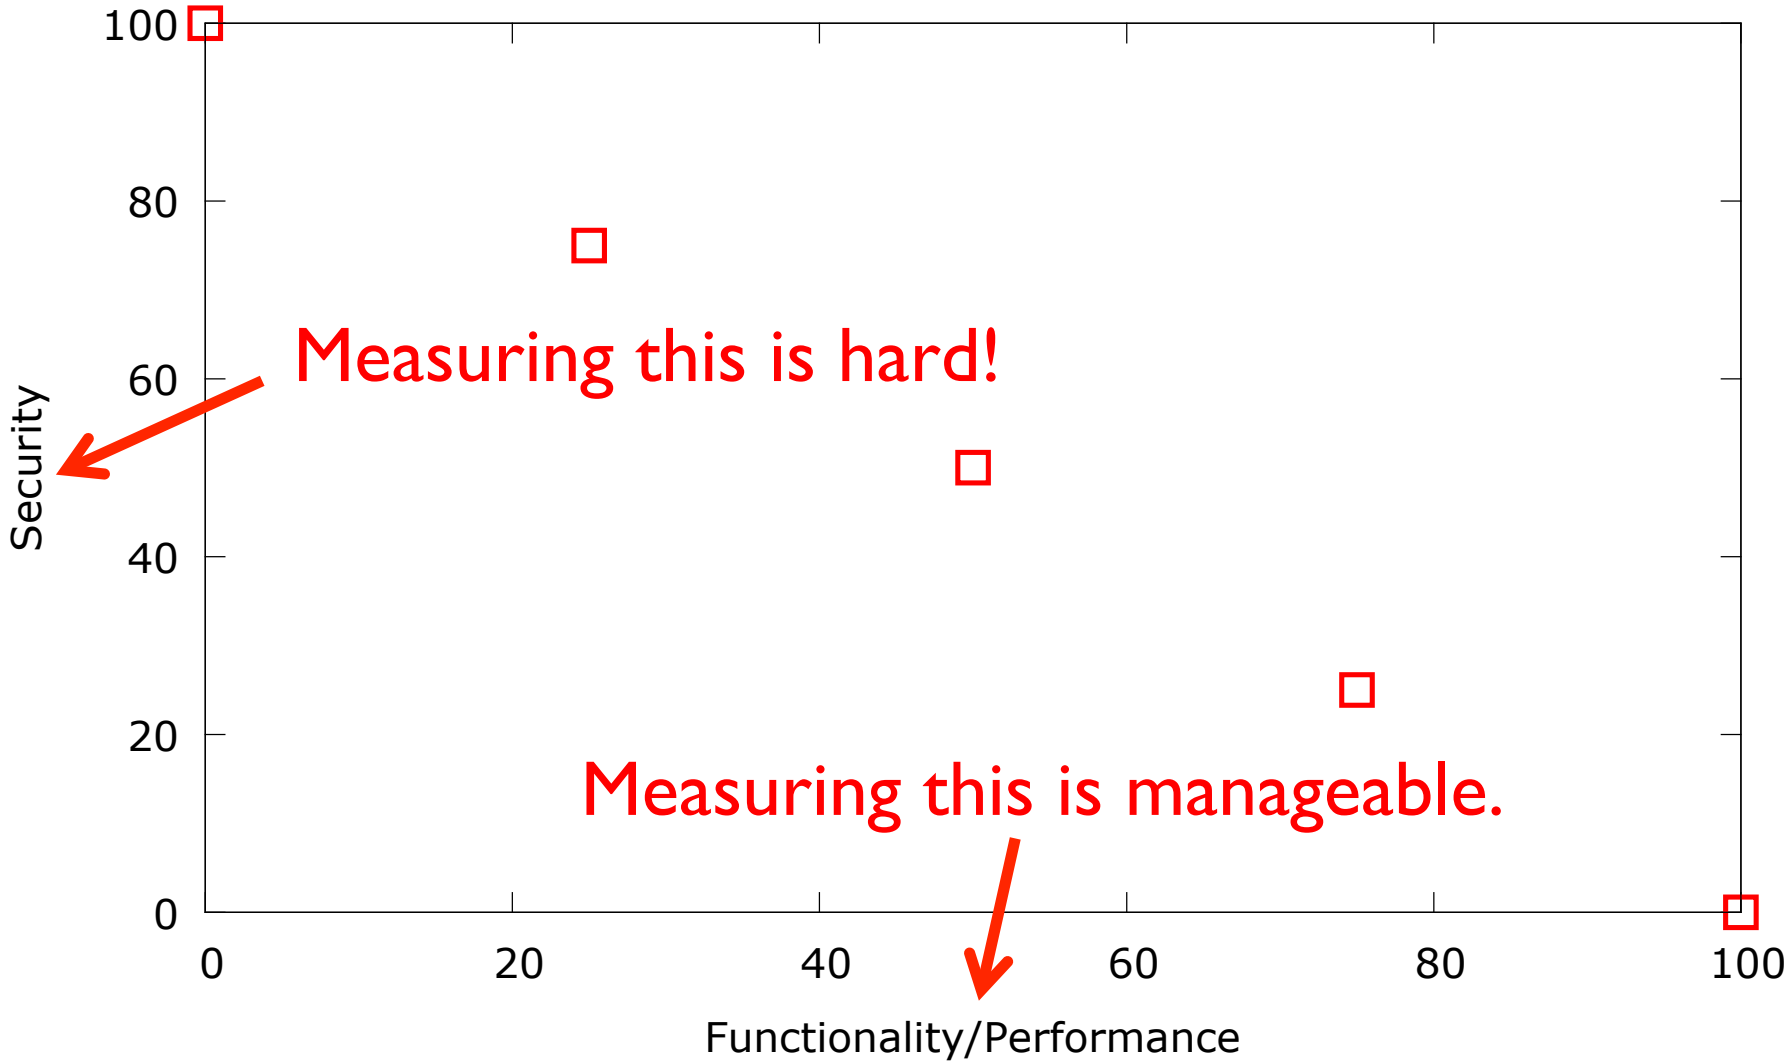

Measurement Example

- A moving target defense (MTD) is intended to make the target environment less predictable
  - Consequence: MTDs should able to quantitatively or qualitative be able to predict how much the defense will impact the predictability of the environment.
  - Q: What to measure?
  - A: Look to the MTD design criteria:
    - Unpredictability in its outcome/use [randomness]
    - Variance of outcomes [scope of possible outcomes]
    - Transparent to the adversary [outcomes not readily detectable, work to determine outcome]

MTD Security Measurement

- Example: Address space layout randomization (ASLR), which randomly arranges areas of the process space
  - [randomness] – uses strong randomness
  - [scope of possible outcomes] – size of address space, and fragmentation of address space
  - [outcomes not readily detectable, work to determine outcome] – random guessing, with and without format string attacks

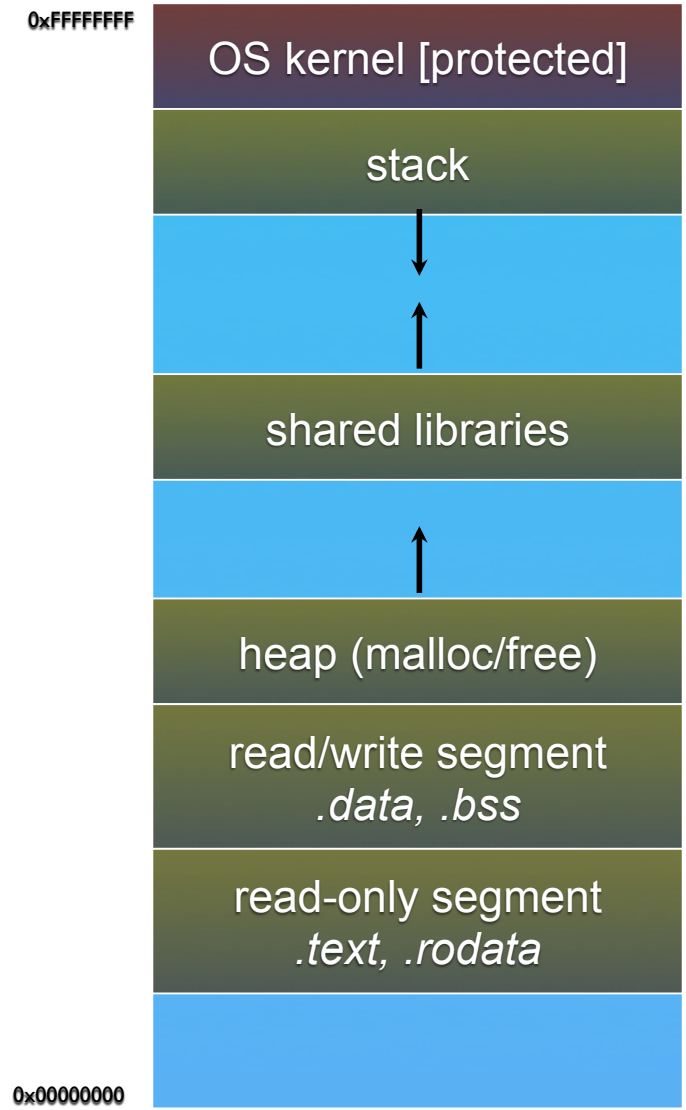

What can we measure?

- What are the security cubits?
  - Design artifacts (e.g., ASLR examples)
  - Compliance (with policies)
  - Adversarial action
  - Performance
  - Function
  - People
  - Risk
  - History

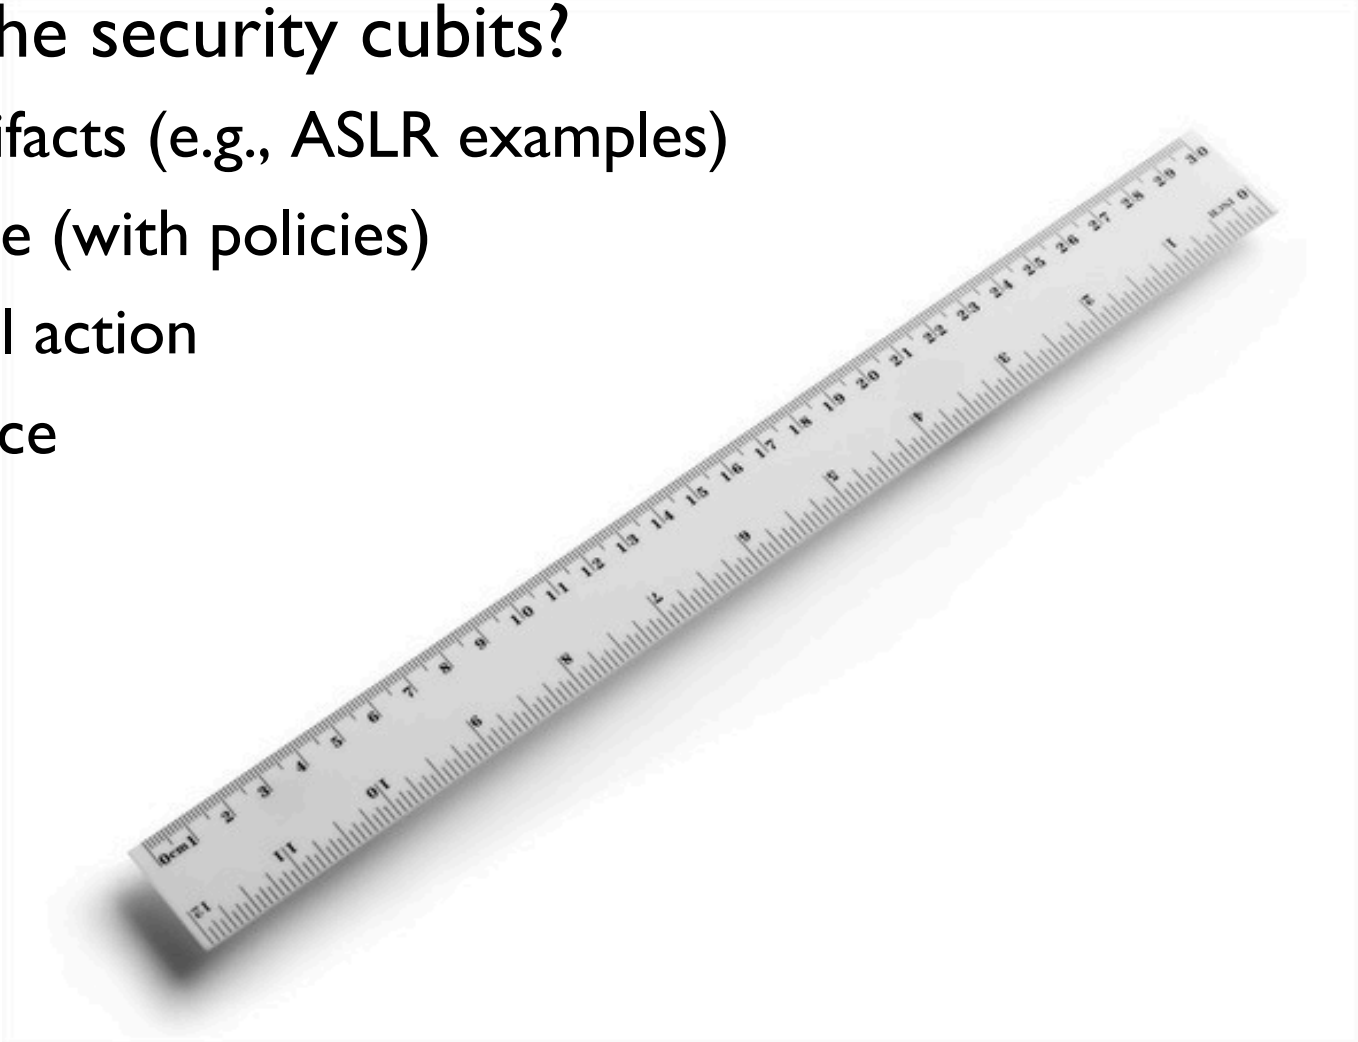

The metric challenge …

- The whole security measurement game is an attempt to move from digital signal to human intent, and ultimately predict the likelihood of future outcomes.

How do you map what you have (ASLR randomness) onto what you want to know (vulnerability)?

- Like the study of auto-accidents, *this will never be an exact science*.

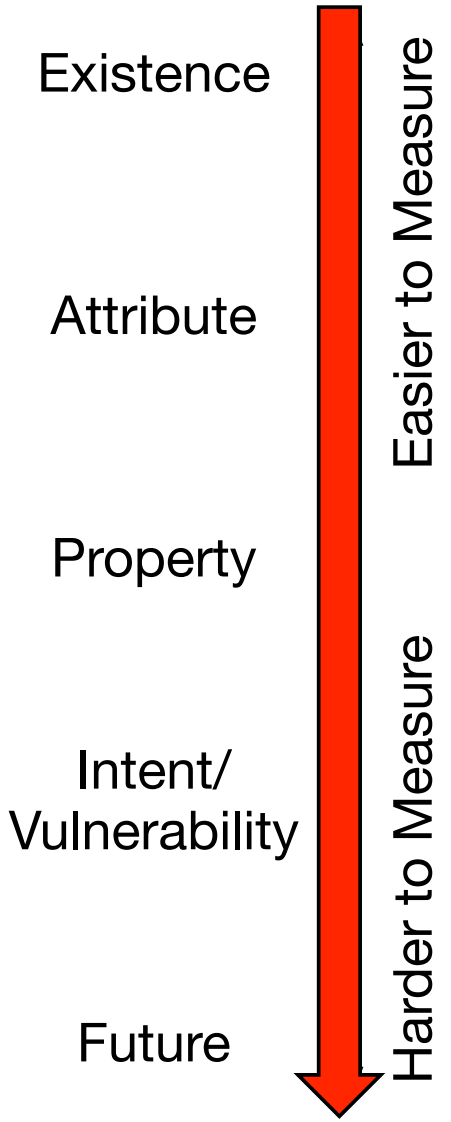

Metrics in practice …

- Need structured thinking of how to move from basic measurements to real actionable knowledge:
  - Attack trees
  - Security ontologies
  - Common vulnerability scoring system (CVSS)
  - Adversarial goal measure (ADVISE)

- This is a whole different talk …

12:30 p.m. -1:30 p.m.

Tutorial 4: Security-Metrics-Driven Evaluation, Design, Development and Deployment
William H. Sanders, University of Illinois at Urbana-Champaign
Location: Room 1040 NCSA

Ok Mike, is it science?

PENNSTATE 1855

- Absolutely: defining ways of projecting low level measurements onto high level concepts:
  - generalizations of metrics to characteristics
  - ontologies to represent and relate different aspects of systems
  - Empirically determining through simulation and in situ where these mappings work and where they don't
  - refining all of the above …

- Desired end state: theory for mapping environmental measurements to security characterizations

Conclusions

- A science of security relevant for the near future must be better grounded in realistic outcomes:
  - Probabilistic metrics of (in)security
  - Based on structured reasoning using
    - low level metrics,
    - historical information, and
    - models of humans and their needs/goals

- Party like it is 1999: going down the security metric rat-hole is necessary for us to make progress.
  - Focus on what we can reliably measure.
  - Work from the measurable to actionable characteristics …

Contact

- Lead PI: P. McDaniel, mcdaniel@cse.psu.edu
- URL: http://cra.psu.edu/