A very personal list of caveats and wishes for systems research

Saurabh Bagchi
Assistant Professor
School of Electrical and Computer Engineering & CERIAS
Purdue University
West Lafayette, IN
sbagchi@purdue.edu

PURDUE UNIVERSITY

CERIAS

What are the characteristics systems should strive for?

- Configurability

- Landscape today:
  - Total cost of ownership of computer systems today is dominated by personnel cost
  - Large number of system crashes and web site unavailability is caused by system management errors
- Systems we do not want to see:
  - 156 parameters to tune, three of these are sensitive enough that a 5% deviation from ideal will cause a catastrophic failure
  - But if all the parameters are perfectly tuned, we will get every ounce of performance out of the system
- Systems we want to see instead:
  - Few high level parameter settings – lower level parameters are derived from them
  - Stability envelope for operation which is never violated

What are the characteristics systems should strive for?

- Malleability

- System should be easy to increment
  - As requirements, underlying hardware, or middleware change
- Systems we (mostly) do not want to see: A one-off system which is designed for one and only one problem and for a specific environment
- Systems we want to see:
  - The system has encapsulated components which are possible to replace in response to changes
  - The system has interfaces that can support differing models of applications
- Both configurability and malleability need to be validated through real user trials
  - Carefully classify your users by role, expertise level
  - Have statistically meaningful sample sizes

Intrusion Response System

- The need for IRS
  - A survivable system needs to provide functionality/confidentiality/integrity at best
  - Human intervention after IDS reporting of an intrusion can be costly and slow
  - IRS takes reports from IDS (usually bundled together), thinks for a while, and carries out actions to counter the intrusion
- Existing examples of IRS
  - Anti-virus software which disables access to worm executables or files infected with virus.
  - Routers/firewalls which actively block worm traffic.
  - Mail server which removes virus infected mail attachments.

Silo-way of building IRS for Distributed Systems

- Most of them are stand-alone and are tied with one single and specific detector
- Need IRS for distributed systems
  - An environment of multiple interconnected boxes with mixed and cooperating services
- Each IRS/IDS pair doesn't leverage the detection reports from the other IDSs
  - Existing research on correlation IDS have shown clear advantages on doing so
- Each IRS/IDS pair doesn't consider the effects from the response actions carried out by the other pairs
  - This can lead to redundant response actions at least and unnecessary denial-of-service of the system at worst
- At best, we have only each IRS/IDS pair trying to react optimally in a local manner. There's no guarantee on system wise global optimality from these individual IRS/IDS pairs' actions.

Experiment – Incorrect Initial Conditions

- ADEPTS where initial settings (effectiveness of responses, IV values, etc.) are incorrect, say due to inexperienced sysadmin

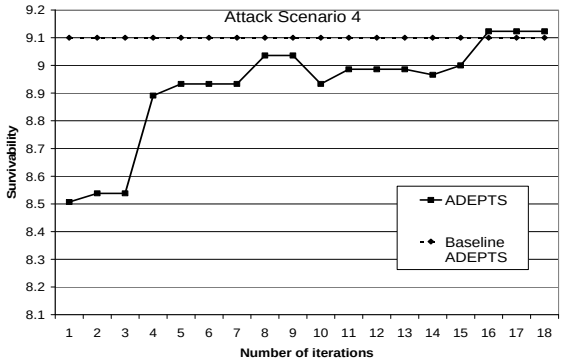

After 16 iterations of the attack, the effect of incorrect initial parameters disappears

Role of rigorous analysis in systems research?

- Analysis can and often does lay out the upper bound of the relevant metric – capacity, speed, dependability
- This serves as useful benchmark for comparison of the actual achieved metric of the system
- Analysis does not mean taking the easy way out by making patently untrue simplifying assumptions
  - Look around at the use of sophisticated modeling tools and solvers
- Examples abound
  - Shannon's channel capacity sets upper bound on rate at which information can be reliably communicated over a channel
  - Byzantine fault tolerance limit spurred systems work on more constrained fault models
  - Our group's work on capacity of a covert timing channel to leak secure information

Systems research in academe and industry: Synergy or opposition?

- Industry can provide an energizing context for the work
  - Personal Example: Motorola providing a mesh network solution to firefighters as they enter a burning building
  - Graph theoretic formulation for ensuring connectivity and coverage
- The problem context can lead to fundamental technical contributions
  - Abstracting out the problem into a realistic but tractable model
  - Solution by developing fundamentally new or modified algorithm
- Solution of multiple grades
  - Highly constrained and somewhat unrealistic problem space – very efficient solutions possible
  - As more assumptions are loosened, the algorithms may resort to heuristics or greedy decisions
- Concerted effort needed among industry engineers and end-consumers to have greater appetite for risk to drive systems innovations

New challenges in wireless systems research

- Wireless devices becoming an ubiquitous part of our landscape – leads to healthy systems research
- Even more experimental variability than in wireline systems research
  - Embedded in changing, possibly hazardous, environment
  - Miniaturized low cost devices
  - Presence of devices interfering on the em spectrum
- Some desirable characteristics of wireless systems research
  - Define what parameters you controlled for
  - Lay out the claim carefully accounting for the specific environment and system configuration
  - Solutions reasonably agnostic to the specifics of the device

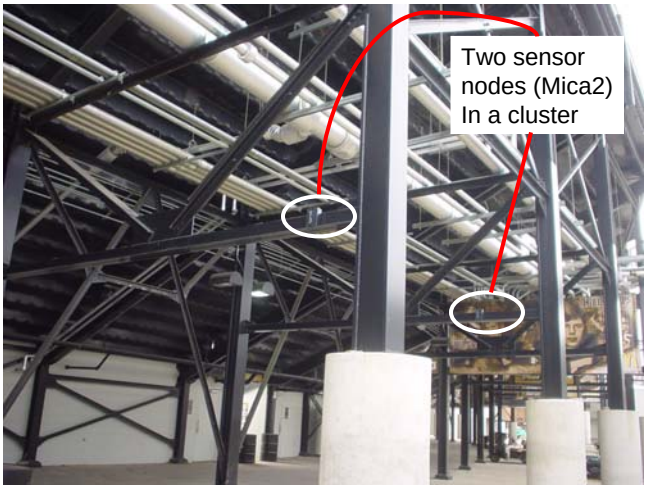

Wrap-up

1. Want to see systems research (also) stress on the metrics of configurability and malleability
2. Analysis should be brought to bear on system design and evaluation
3. Industry can spur challenging and meaningful systems research
4. Wireless systems research has unique promise and challenges