Computer Systems at CMU

As told by Onur Mutlu
(on behalf of many)
onur@cmu.edu
Carnegie Mellon University

ECE PhD Orientation 2013

An Exciting and Collaborative Place to Fundamentally Change the Way Future Computing Systems are Designed

Hopefully you are already excited!

Tons of Great People Doing Lots of Diverse, Evolving Work

Pick an Advisor

(not a project)

Computer Systems

- Computer Architecture
  - Seth Goldstein, James Hoe, Diana Marculescu, Todd Mowry, Onur Mutlu, David O'Hallaron
- Security
  - Lujo Bauer, David Brumley, Lorrie Cranor, Anupam Datta, Adrian Perrig, Virgil Gligor
- Networks
  - Hyong Kim, Srinivasan Seshan, Peter Steenkiste, Hui Zhang, Rohit Negi
- Data Storage
  - Greg Ganger, Garth Gibson
- Embedded Systems
  - Priya Narasimhan, Phil Koopman, Raj Rajkumar, Dan Siewiorek, Asim Smailagic
- Fast Computing
  - Franz Franchetti, Markus Püschel
- Silicon Valley
  - Martin Griss, Jason Lohn, Joy Zhang, Pei Zhang, Collin Jackson, Patrick Tague

And Others

PARALLEL DATA LABORATORY
CARNEGIE MELLON UNIVERSITY

CALCM

Carnegie Mellon
CyLab
CONFIDENCE FOR A NETWORKED WORLD

CyLab Usable Privacy & Security Laboratory
HTTP://CUPS.CS.CMU.EDU

Carnegie Mellon
BIOMETRICSLAB

CMU Silicon Valley

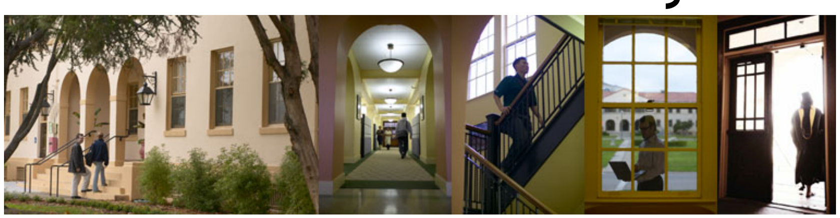

Collaboration is Natural at CMU

- Within the Systems Area: Parallel Data Lab, ISTC-CC, ISTC-EC, CSSI, …
- Across Areas: Circuits, Architecture, Systems, Security, Programming Languages, Theory, …
- Across Departments: ECE, CSD, RI, EPP, …
- With Many Partners in Computing Industry
- With Other Researchers outside CMU

A Personal Example

Students Drive Research

Exciting Areas

Just a few I'll mention:

- Cloud Computing Research
- Computer Architecture Research
- Embedded Systems Research
- Computer Security Research

Cloud Computing (+Big Data)

- Cloud computing: use someone else's computers instead of buying/maintaining your own
- Promises huge benefits in…
  - resource efficiency: consolidation
  - agility and productivity of application development
  - system and service robustness
  - data sharing and end user access
- Huge need for fundamental research and understanding at many levels

The Beginnings: PDL’s Data Center Observatory

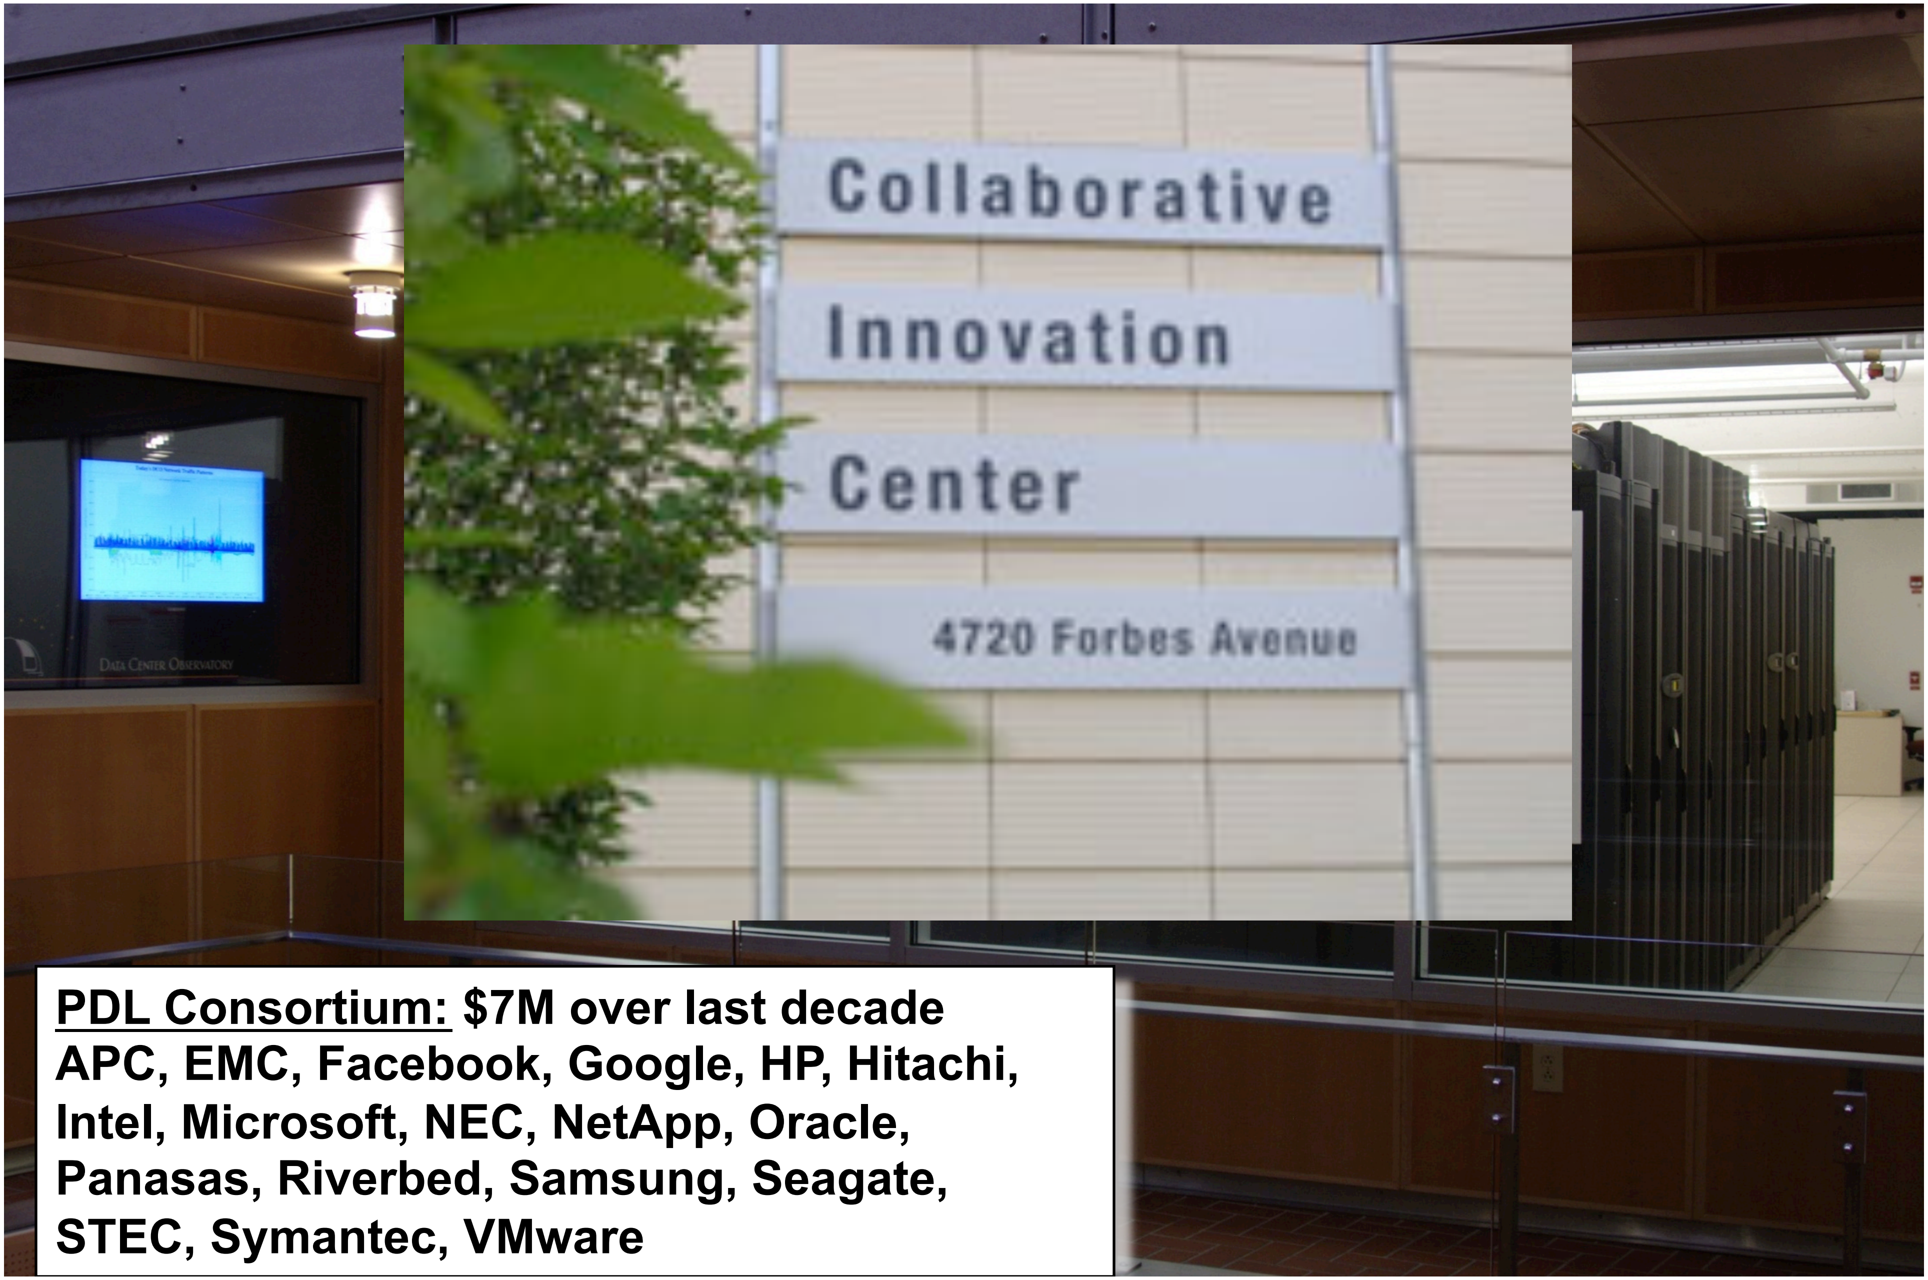

<u>**PDL Consortium:**</u> **$7M over last decade**
APC, EMC, Facebook, Google, HP, Hitachi, Intel, Microsoft, NEC, NetApp, Oracle, Panasas, Riverbed, Samsung, Seagate, STEC, Symantec, VMware

The Parallel Data Lab

Academia's Premier Storage Systems Lab!

Selected Research Topics
(Too many to list here!)

- New storage architectures
- Distributed storage systems
- File & database systems
- Storage management
- Secure storage systems
- Home media systems
- Data organization & search
- Virtual machines
- Applying new technologies

Very Successful History

- Founded in 1992
- Numerous best paper awards
- Excellent partnership with industry
- Successful public code deployments

Significant Industry Involvement

- 17 member companies!
- Two annual on-campus events
- Access to the industry leaders

Check out www.pdl.cmu.edu *for the whole story!*

PDL Facts

- Established: 1992
- Faculty involved: 25
  - Career pubs >1,800, citations ~30,000
- PhD+MS graduated: 57+53
  - Faculty at Berkeley, Stanford, MIT, Toronto, ...
  - Key technical leaders at EMC, Seagate, Intel, ...
- Papers published: >200
  - FAST 2011, 2010, 2009, 2008, 2007, 2005, 2004, 2003, 2002
  - SOSP 2011, 2009, 2007, 2005; HotOS 2010, 2009, 2003
  - NSDI 2011, 2007, 2006; SIGCOMM 2009
  - ATC 2010, 2007, 2006, 2005, 2003, 2002
  - SIGMETRICS 2008, 2007, 2006, 2004, 2002
  - ICDE 2007, 2004, 2002; VLDB 2010, 2007, 2004, 2003, 2002; SIGMOD 2005, 2003
  - SC 2011, 2009, 2005, 2004; HPDC 2010; SciDAC 2007; PDSW 2009, 2008, 2007
  - ISCA 2011-2013; HPCA 2010-2013; MICRO 2010-2012; ISCA 2009; ASPLOS 2004; Computer 2010, 2006, DSN 2010, 2006, 2004; HotDep 2009, 2006
- Awards: Last 5 years: 54; Prior 5 years: 3
  - 2011 SOSP Hall of Fame
  - 2011 IEEE Young Computer Architect
  - 2011 ACM Graduate Student Research
  - ...

PARALLEL DATA LABORATORY
CARNEGIE MELLON UNIVERSITY




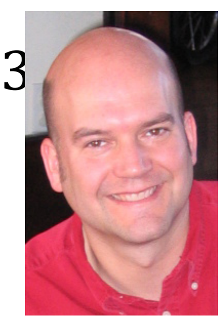

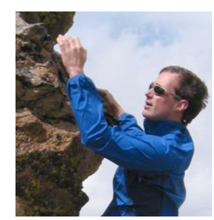

Faculty Associated with PDL

- Greg Ganger, PDL Director
- David Andersen
- Lujo Bauer
- Chuck Cranor
- Lorrie Cranor
- Christos Faloutsos
- Eugene Fink
- Rajeev Gandhi
- Garth Gibson
- Seth Copen Goldstein
- Mor Harchol-Balter
- Bruce Krogh
- Todd Mowry
- Onur Mutlu
- Priya Narasimhan
- David O'Hallaron
- Andy Pavlo
- Adrian Perrig
- Mahadev Satyanarayanan
- Srinivasan Seshan
- Bruno Sinopoli
- Hui Zhang

Cloud and Big Data Research Themes

- Scalability
  - Performance demand drives ever more devices
- Manageability
  - Human costs cannot grow with scale
- Models of Computation: Data-Intensive Computing
  - BigData needs simple tools for processing massive data
- Anticipating & Exploiting New Technology
  - Non-volatile memory devices & non-overwrite disks
  - New hardware/software stacks, new failure modes, new efficiencies

Intel Science & Technology Center for Cloud Computing (ISTC-CC)

- 5-institution community
  - led by Carnegie Mellon and Intel
  - including Georgia Tech, Princeton, and UC-Berkeley
  - homed at Carnegie Mellon; 4 Intel researchers onsite
- Focus on critical enabling technologies
  - foundations on which future clouds will be built
    - complementing evolution of early cloud building blocks
  - work with researchers exploring new cloud apps
    - e.g., ISTCs on visual, pervasive, and embedded
- Broad, aggressive 5-year research agenda

ISTC-CC Research Pillars Build Foundation

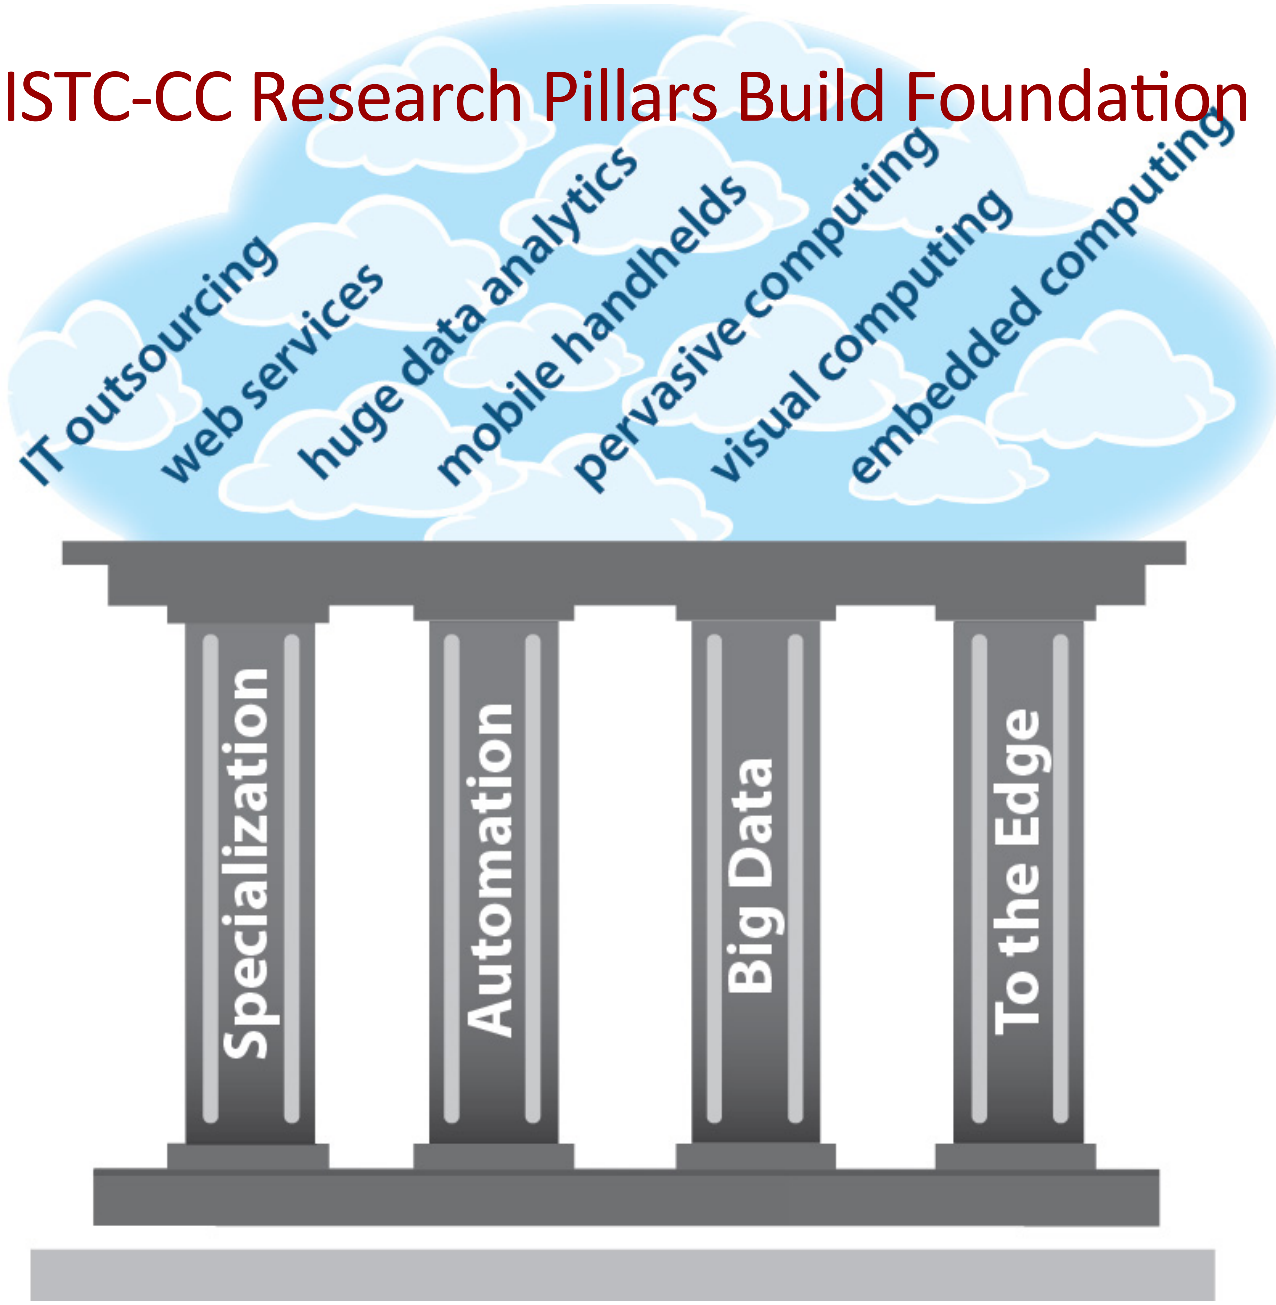

Intel Science & Technology Center for Embedded Computing (ISTC-EC)

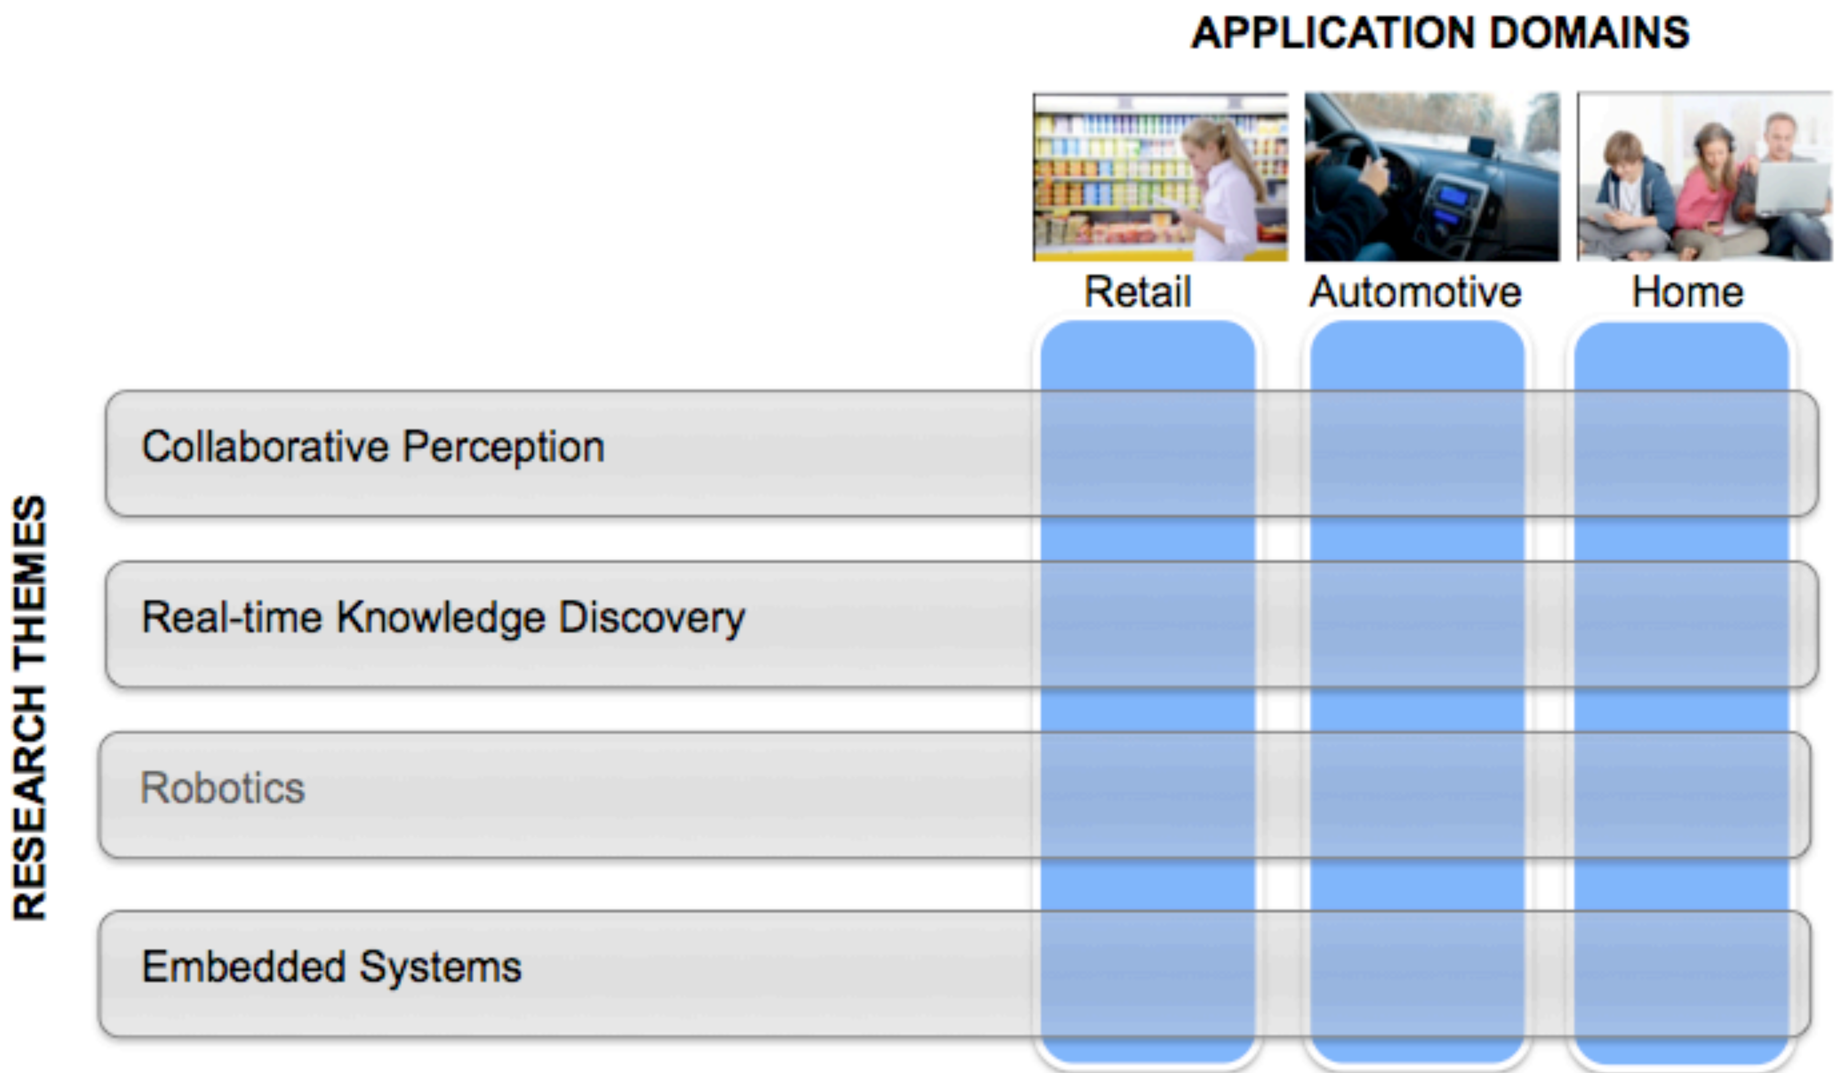

CyLab

World's Premier Academic Research Center for Cyber Security

- One of largest U.S. university-based research & education programs for cyber security, dependability & privacy
- Nationally Recognized
  - National Science Foundation (NSF) CyberTrust Center
  - National Security Agency (NSA) Center of Academic Excellence in Information Assurance Education
- Many awards
  - ACM Best Doctoral Dissertation Award 2011

- Cross-disciplinary, university-wide program
  - Faculty and researchers from six colleges of Carnegie Mellon
  - 50+ faculty/researchers and 130+ graduate students
- Funded by both private & public sectors
  - Gov't like NSA, DARPA, NSF
  - Companies like Lockheed Martin, Booz Allen, Northrop Grumman, and others.

Carnegie Mellon CyLab

CONFIDENCE FOR A NETWORKED WORLD

Cross-Cutting CyLab Research Thrusts

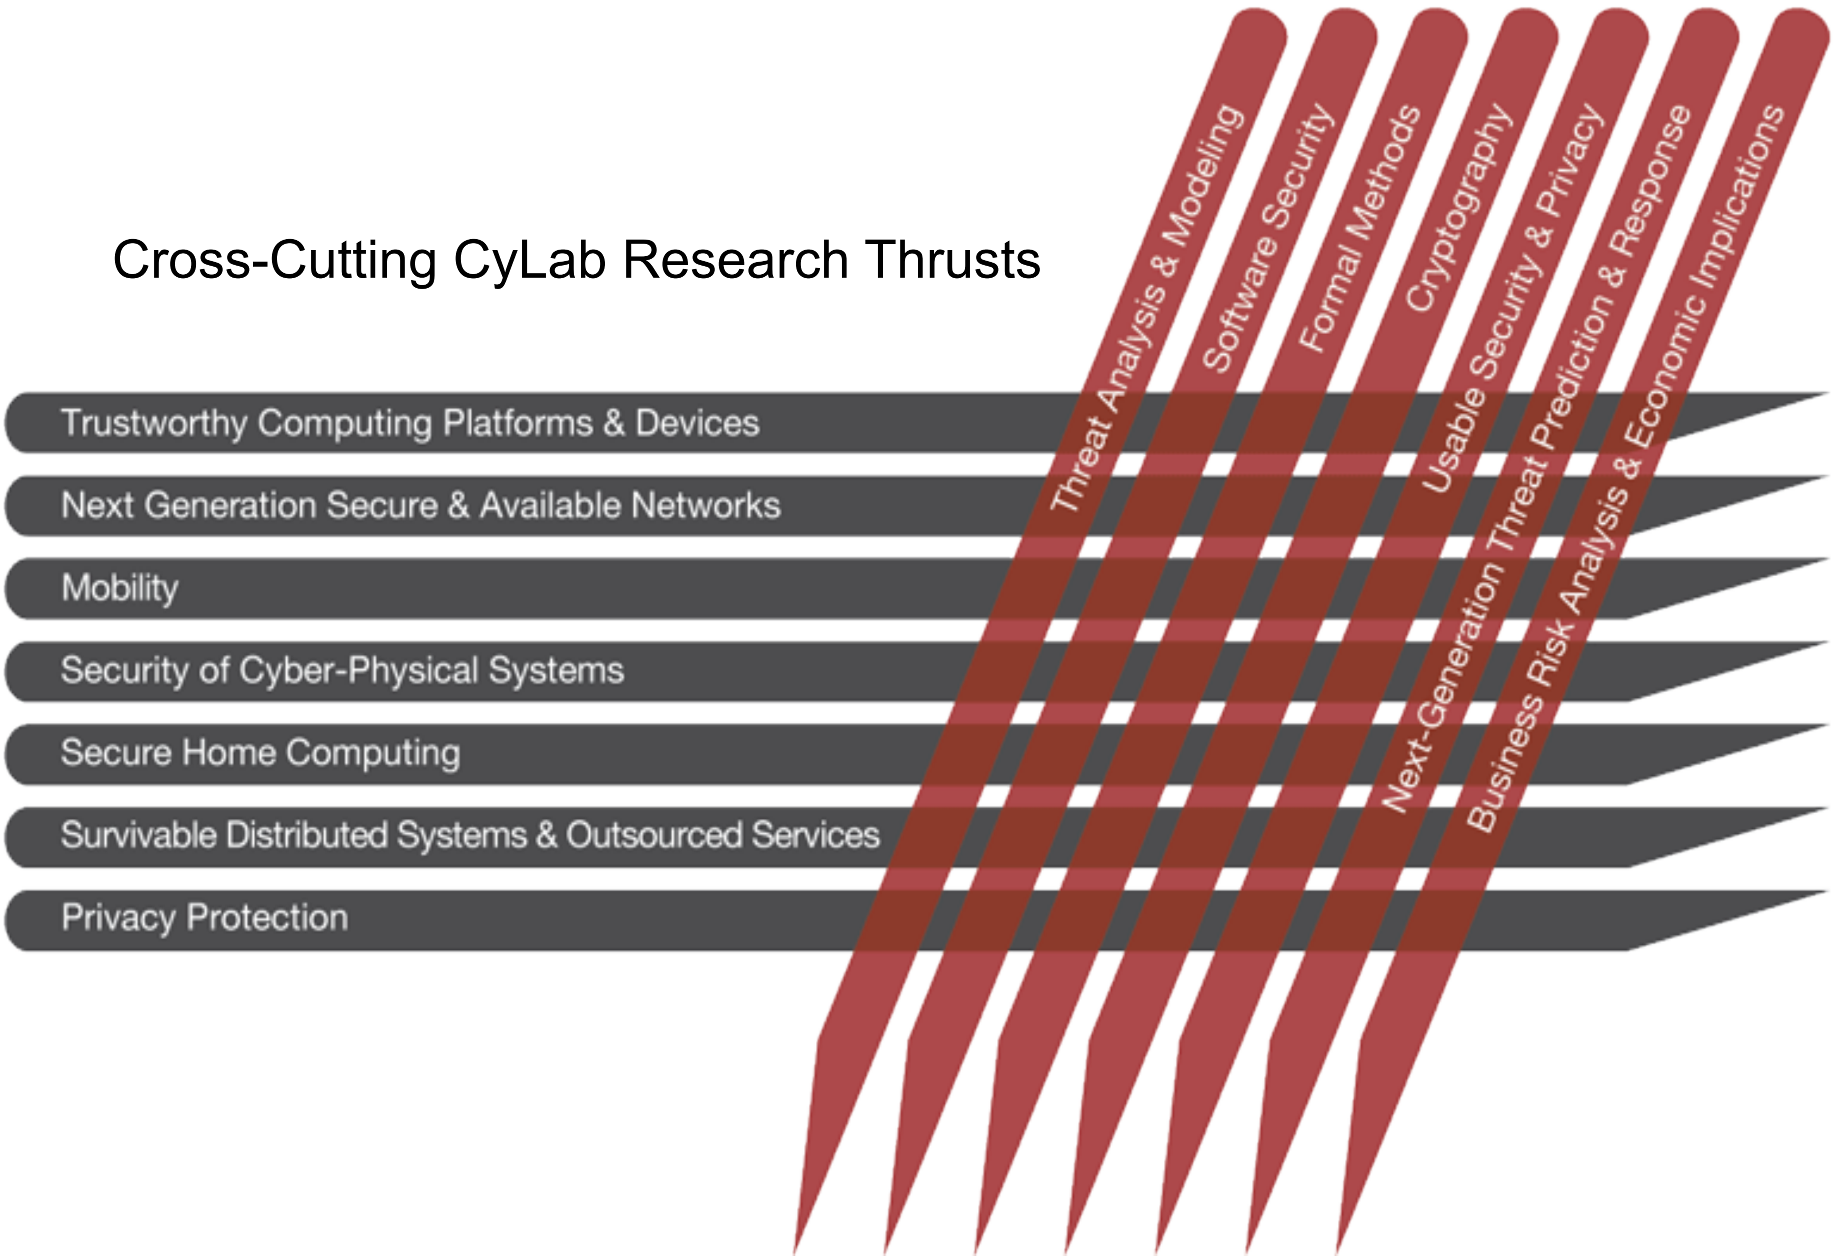

CyLab Researchers & Collaboration

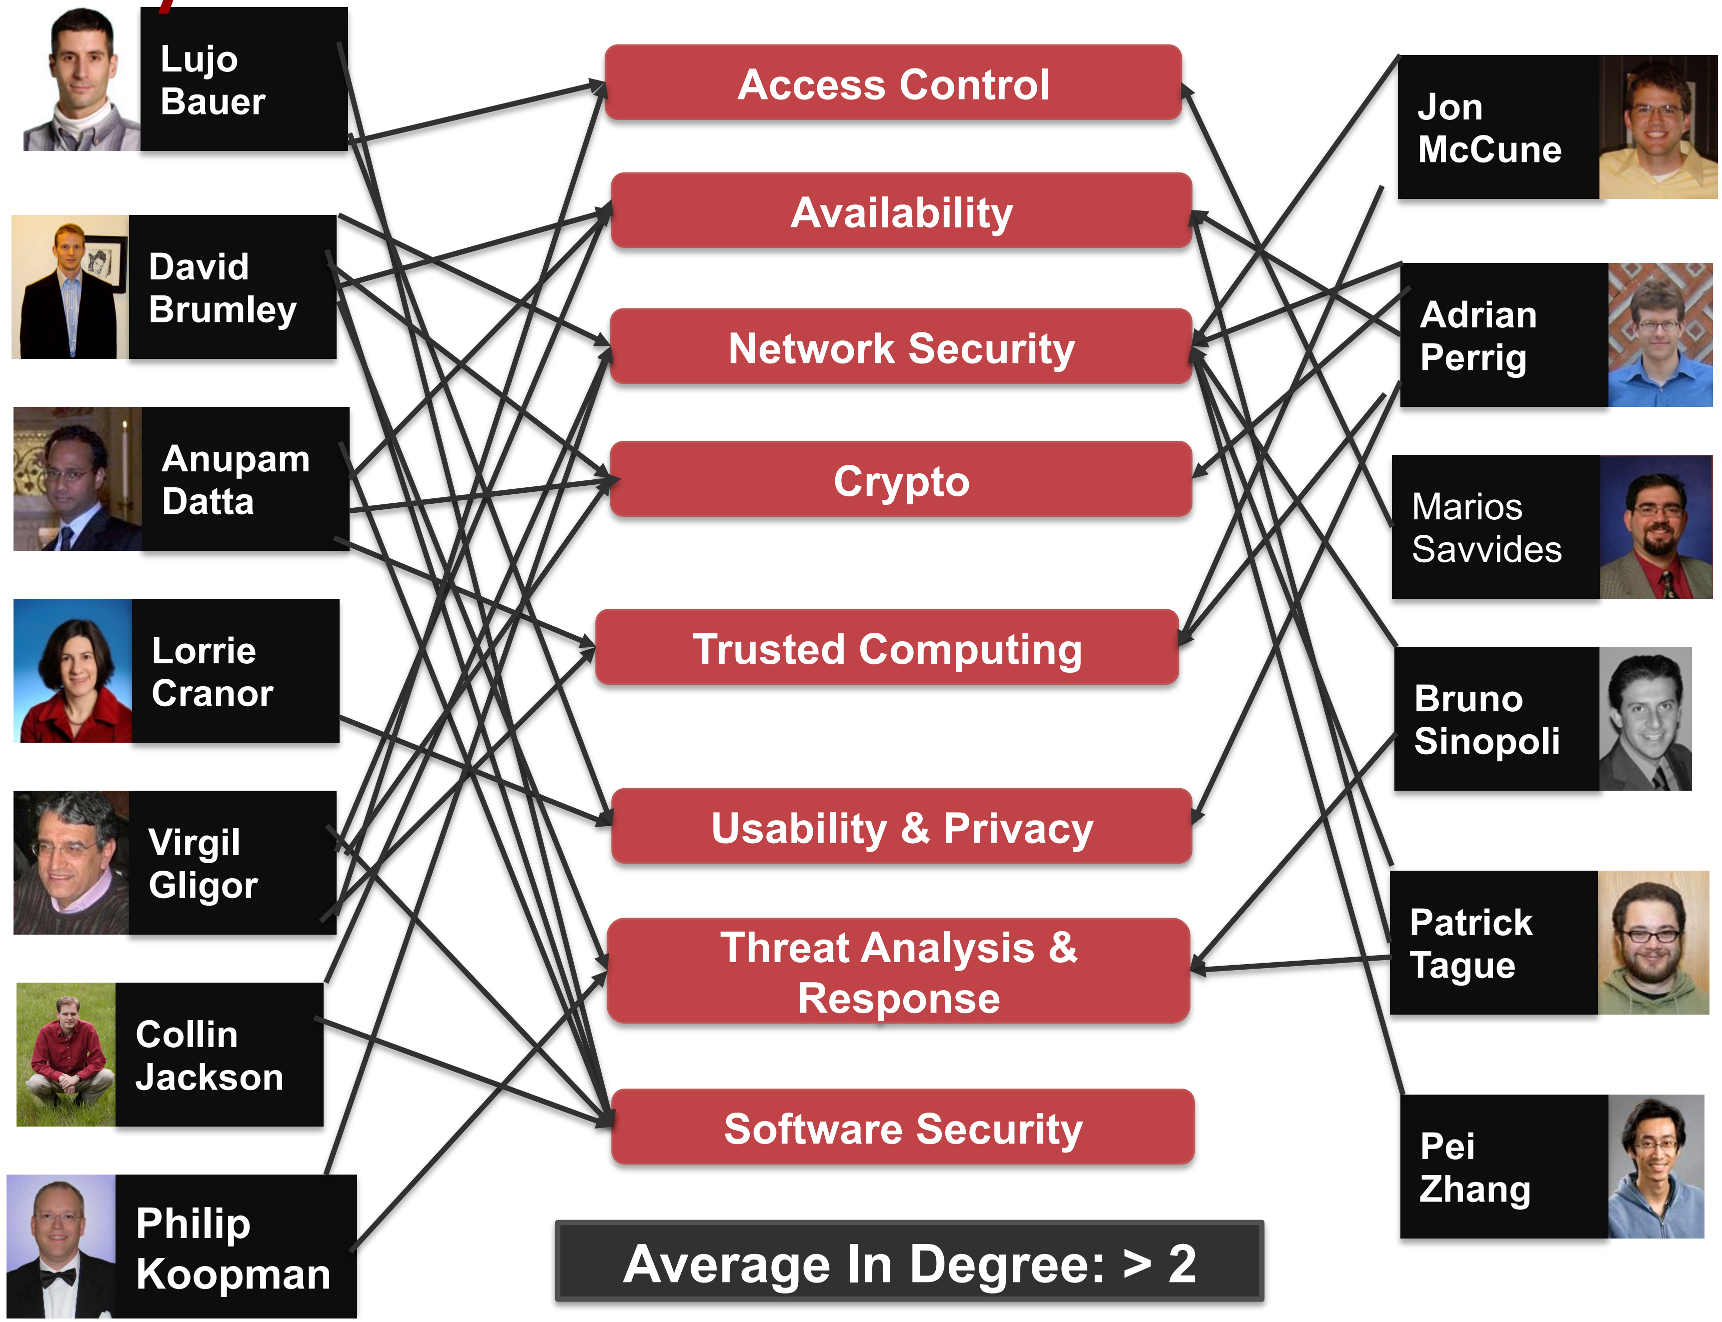

Cross-cutting Centers and Research

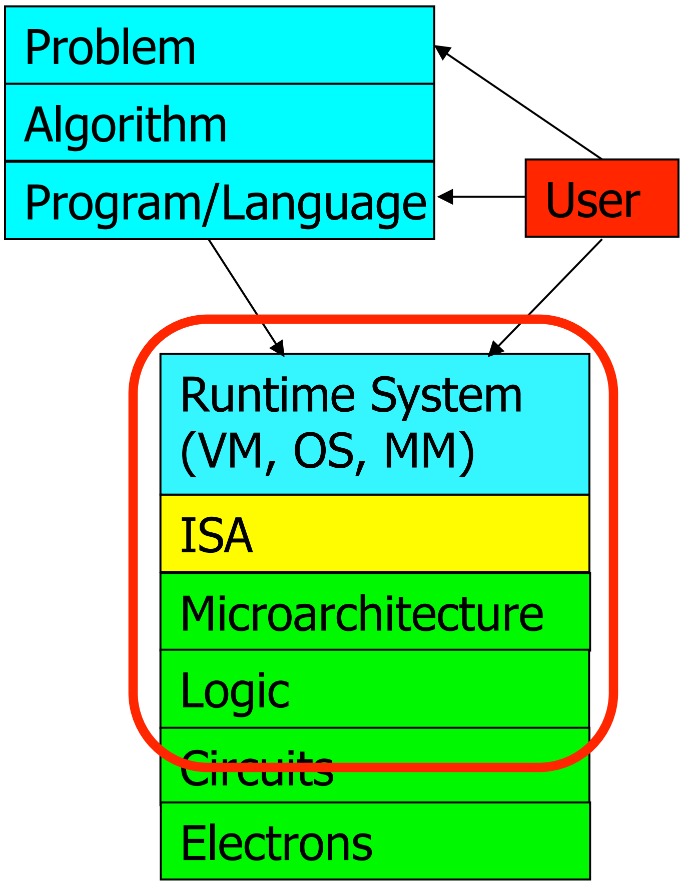

Collaboration at CMU fosters cross-cutting, cross-layer research across the computing stack

Computer Architecture Lab at Carnegie Mellon

- 10-20 years from now
  - What are the computation models?
  - How are computation, communication and memory designed and managed?
  - What are the system software and HW/SW interface primitives?

CALCM

SAFARI

Onur Mutlu
Computer Architecture and Systems

<u>Research Focus:</u> Computer architecture, HW/SW, bioinformatics

- *parallel architectures, many-core systems*
- *memory systems, interconnects*
- *hardware/software interaction, novel execution models*
- *energy efficiency, fault tolerance, hardware security*

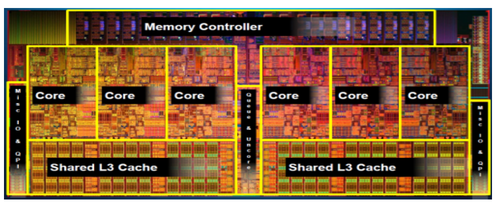

Intel Core i7
8 cores

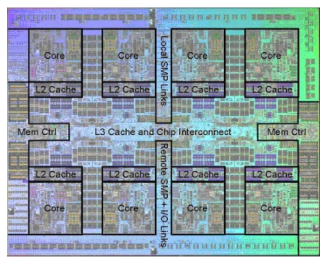

IBM POWER7
8 cores

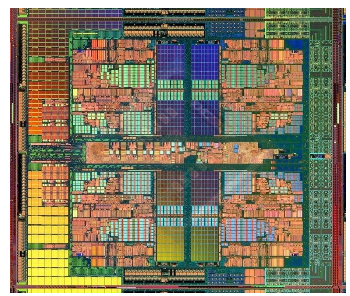

AMD Barcelona
4 cores

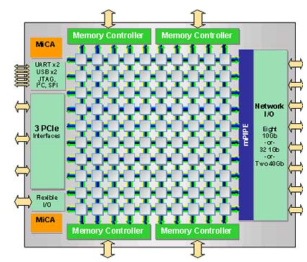

Tilera TILE Gx
100 cores

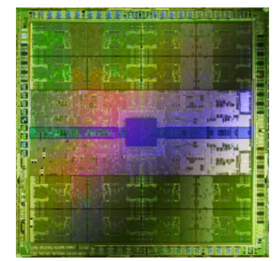

Nvidia Fermi
448 "cores"

<u>Some Questions:</u>

- *How do we design ultra-efficient systems?*
- *Alternatives to multi-core and today's memory?*
- *Can we put the entire data center on a single chip?*
- *What is the HW/SW interface and chip of 2030 AD?*

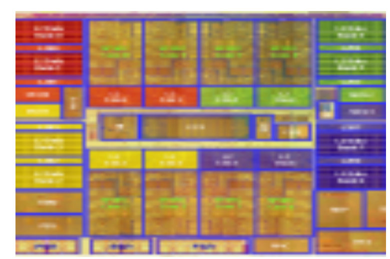

Sun Niagara II
8 cores

Heterogeneous parallel architectures

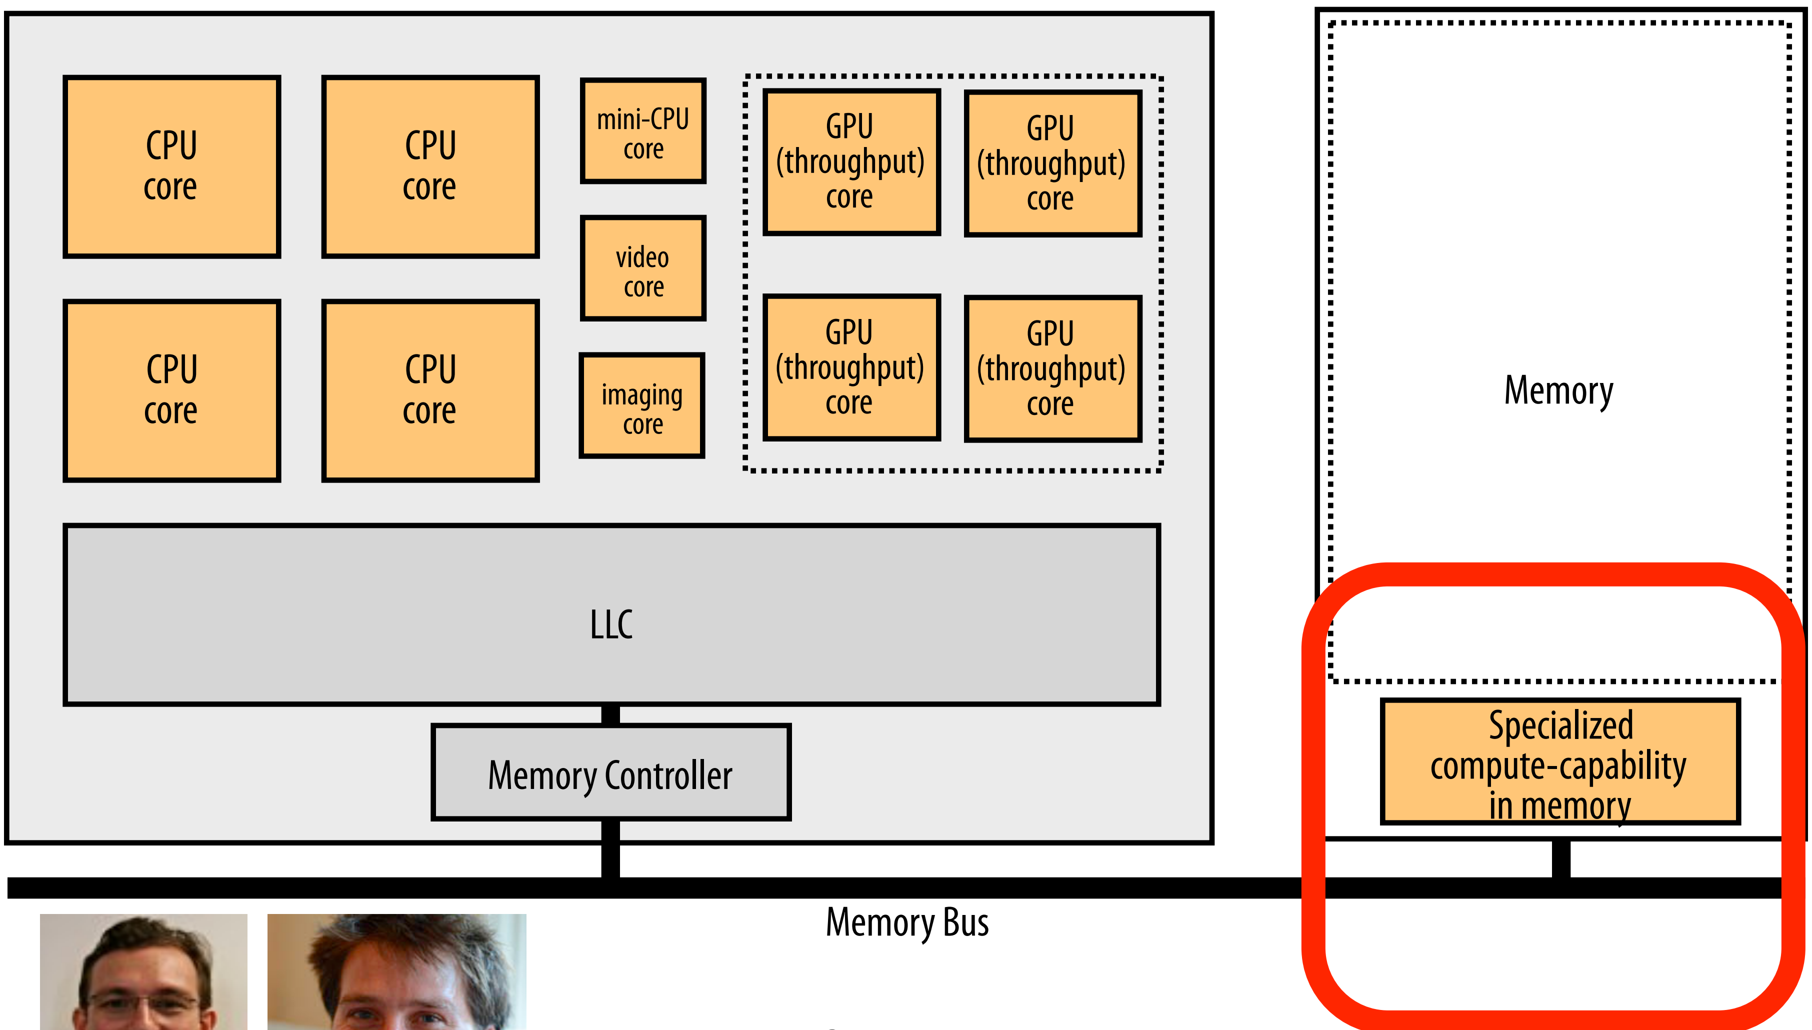


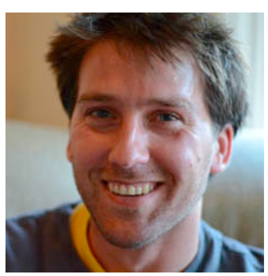

Onur Mutlu and Kayvon Fatahalian

Enabling Ultra-Efficient (Visual) Search

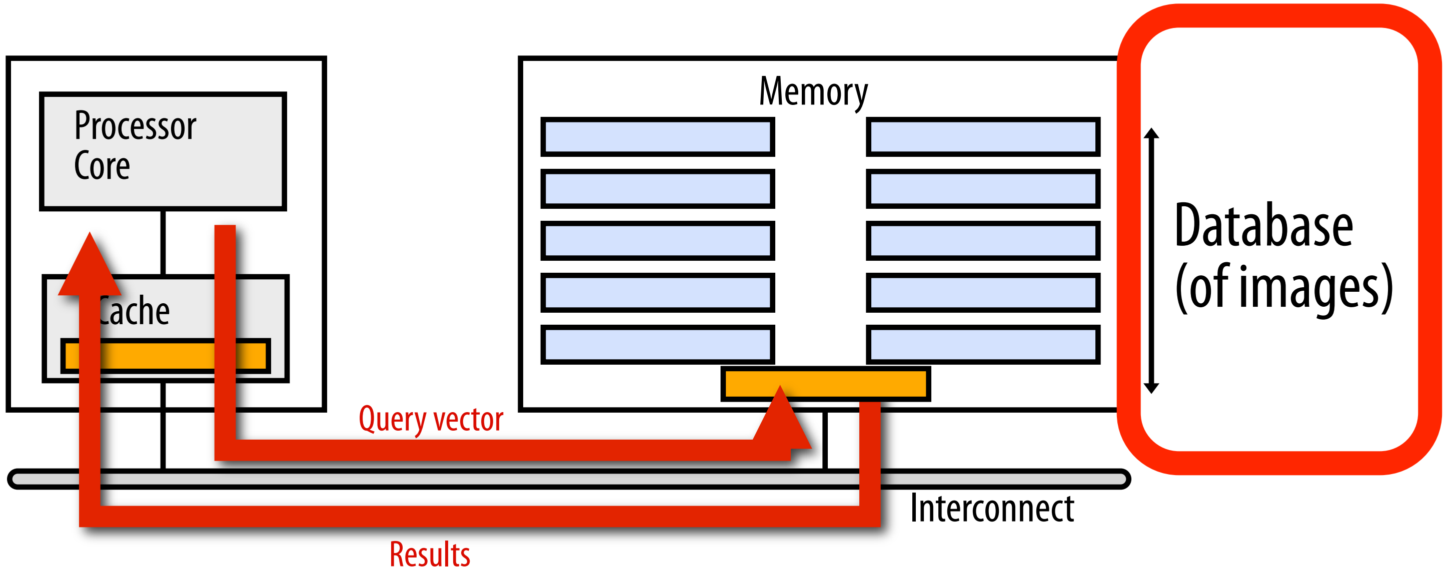

- What is the right partitioning of computation capability?
- What is the right low-cost memory substrate?
- What memory technologies are the best enablers?
- How do we rethink/ease (visual) search algorithms/applications?

CoRAM Memory Architecture

Can we make reconfigurable fabric a first-class computing device?

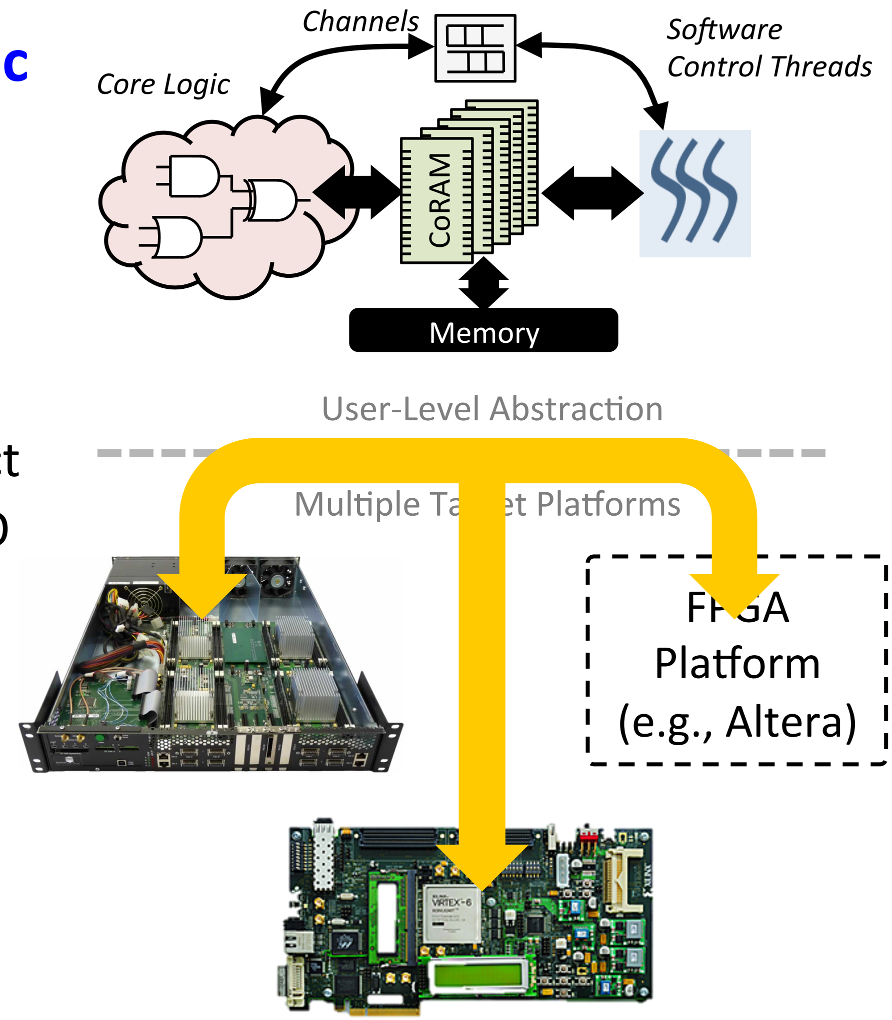

Memory and I/O is Hard on FPGA

Proposed Solution *[Chung, FPGA'11]*

- extend capability of FPGA SRAMs to act as universal portals to memory and I/O
- built-in NoC provides data transport
- replace rigid platform interfaces with **portable software control threads**

James Hoe

Spiral: Computers Programming Computers

Result

Same performance, 1/10th human effort, non-expert user

Key ideas

restrict domain, use mathematics, performance portability

Franz Franchetti

D. McFarlin, F. Franchetti, M. Püschel, and J. M. F. Moura: High Performance Synthetic Aperture Radar Image Formation On Commodity Multicore Architectures. in Proceedings SPIE, 2009.

Franz Franchetti

Project 1: Extending Spiral to new algorithms and platforms

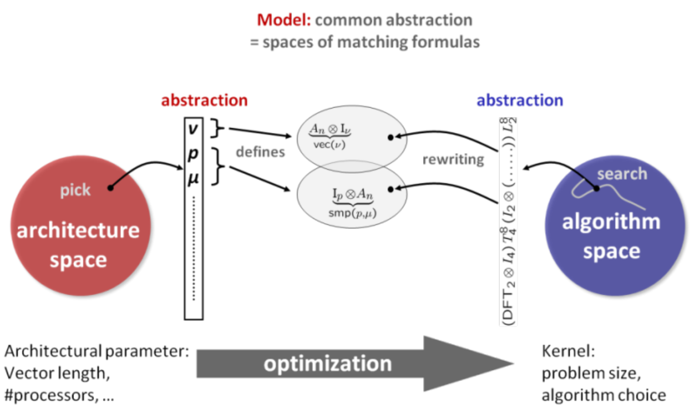

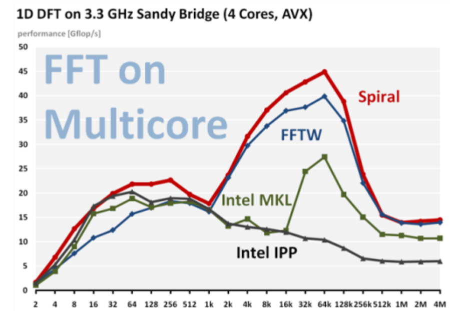

Project 2: domain-specific power efficient computing

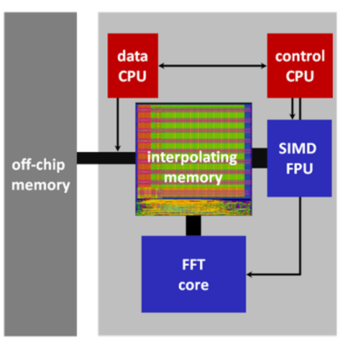

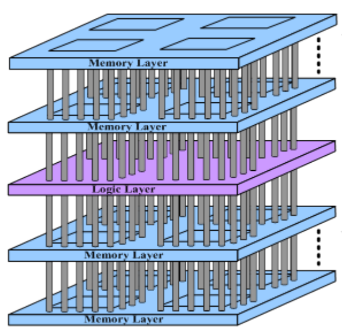

Project 3: Code/proof co-synthesis for robotic vehicles

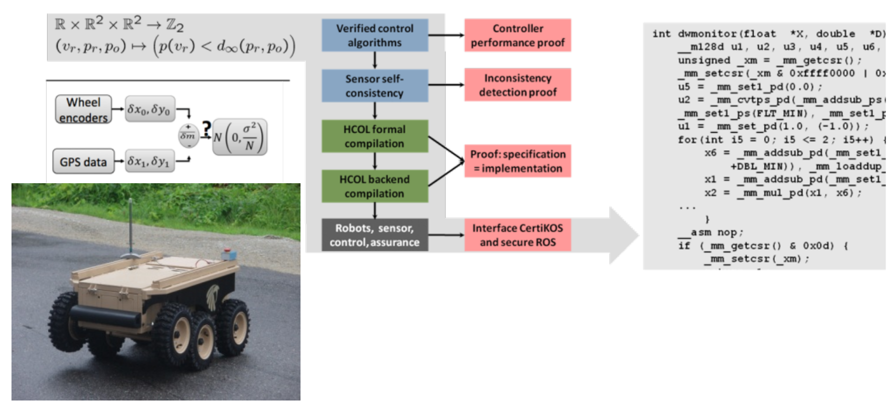

Diana Marculescu: Resource Mgmt. for Many-Core

- ***Holistic* power & perf. mgmt. for many-core systems**

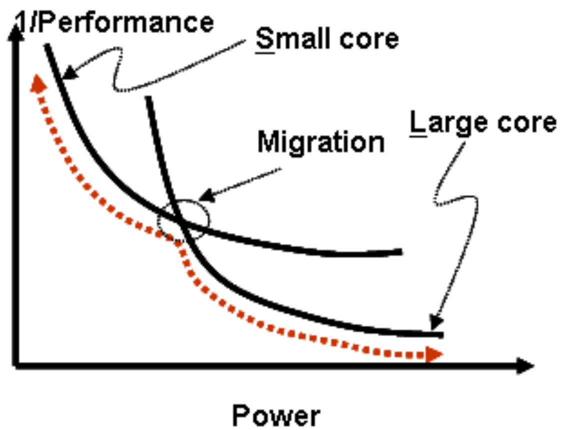

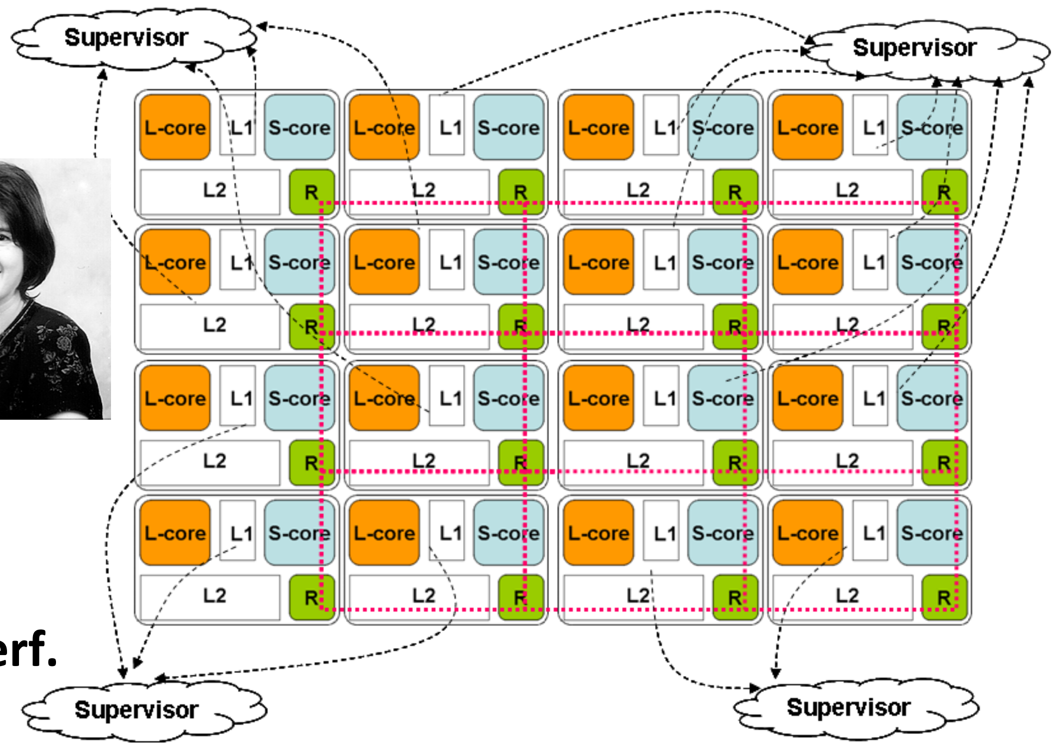

- **10-20X lower power/delivered perf. by exploiting**
  - Workload variation
  - Near-threshold computing
  - Voltage/frequency overscaling
  - Implicit and explicit resource heterogeneity

Challenge: ***determine best online policy for maximizing performance under TDP constraints***

P. Koopman

Associate Professor
Department of ECE
& National Robotics Engineering Center

http://www.ece.cmu.edu/~koopman

Courses:

18-348: Embedded System Engineering

18-649: Distributed Embedded Systems

18-849: Dependable Embedded Systems

Safety Critical Systems

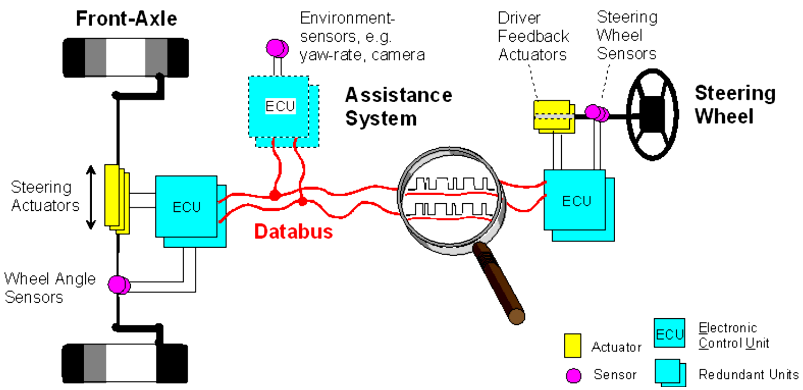

- Cars will be controlled via networks
  - Steer-by-wire; Brake-by-wire; …
  - Can we make them safe+affordable?
- Safety even in the presence of bugs
- General Motors Sponsorship

Robot Robustness

- Is this 5-ton robot acting in a safe way?
  - (How can you be sure?)
- Software stress testing
  - Do bad inputs make it go crazy?
- Run-time safety monitoring
  - Needs to be safe even with bugs

Electrical & Computer ENGINEERING

koopman@cmu.edu

Radu M.: Design of Highly Integrated Multicore Systems

Research Focus*: Embedded Systems, Cyber Physical Systems*

Vision: Demonstrate low-power and reliable approaches for multiprocessor communication at nanoscale (i.e., developing sort of internet-on-a-chip).
(collaboration w/ D. Marculescu, S. Blanton)

Heterogeneous Network-on-Chip platform

Multi-frequency/voltage NoC

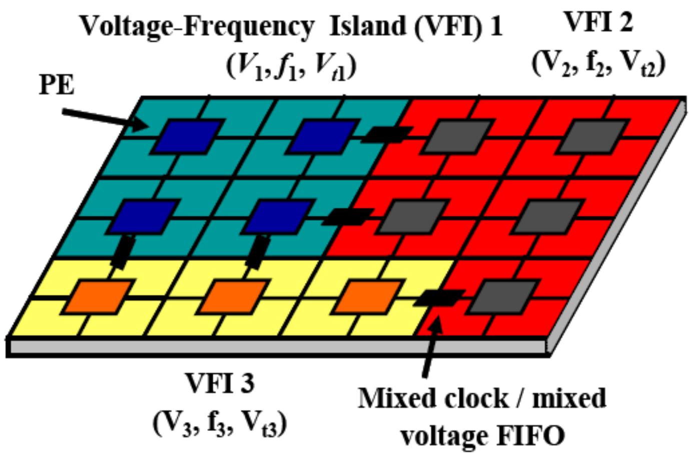

Tasks: Develop distributed energy/performance optimization techniques for multicore systems; Develop on-chip fault-tolerant communication techniques

Radu M.: Design of Cyber Physical Systems (CPS)

Research Focus*: Embedded Systems, Cyber Physical Systems*

CPS operation from physical processes (e.g., volcanic activity monitoring, traffic conditions, etc.) to workloads. This consists of data monitoring and compression/communication to data centers for further analysis.

Distributed controllers dynamically estimate the workload and decide on prioritizing data transmission or allocating more resources.

Examples include implantable devices like pacemakers, body-area-networks, nano/bio devices for health care applications

Radu M.: Social Networks Modeling and Analysis

Research Focus*: Social Systems*

Social networks involve both technology and human interaction and show remarkable dynamics in response to various real-world events. Our research targets new approaches for characterizing complex interactions and collective behaviors in bursty events in social media.

This can be used for real-time management of social networks in applications like micro-blogging, viral marketing, emergency response, and disaster management.

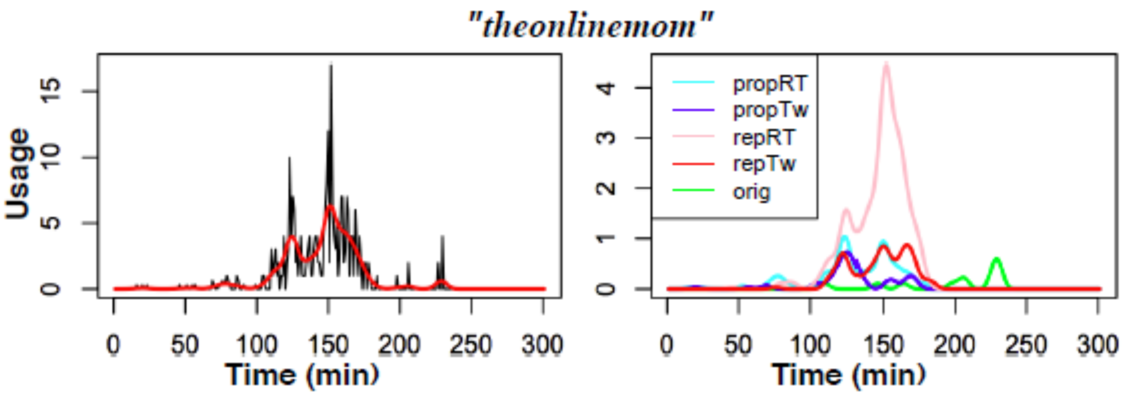

Peter Steenkiste

The eXpressive Internet Architecture

The Internet has been wildly successful, but the architecture faces several challenges

- Security should be built in, not layered on top of, the architecture
- Better support for content and service-oriented networking
- Mobile, wireless users should be first class citizens

XIA is a future Internet architecture proposal that radically improves trustworthiness and evolvability of the Internet

- Intrinsic security properties that are part of the architecture
- Communication between diverse principals: hosts, content, …
- Support for rich services, e.g. transcoding, mobility, DTN, …

This new interdisciplinary project is looking for students …

- … interested in networking, distributed systems, security, ..
- … with good system building skills and a sense of adventure!

Dan Siewiorek and Asim Smailagic: Virtual Coaches

Vision: Use low cost sensors and machine learning to determine user context and respond proactively anticipating user needs

Dan Siewiorek and Asim Smailagic

Machine Learning on Embedded Platforms

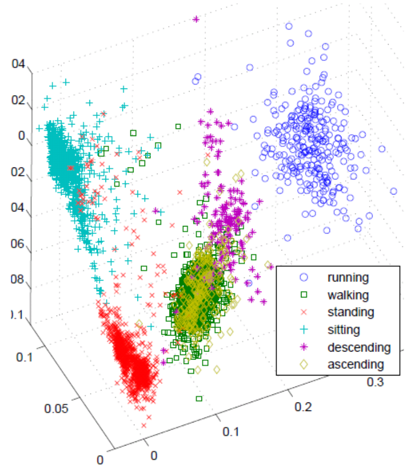

Feature space for accelerometers

Wheel Chair Propulsion Pattern to Avoid Repetitive Use Damage to Shoulders and Wrists

Wrist worn sensors

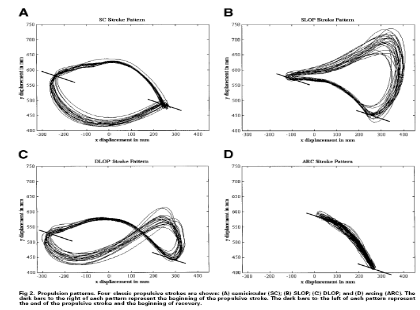

Observed propulsion patterns

Tasks: Develop machine learning models for multi-sensor systems; provide appropriate engaging guidance and encouragement in domains of health care, occupational therapy, and social networking; Cloudlet Architecture to provide mobile computational resources

David Brumley – Computer Security

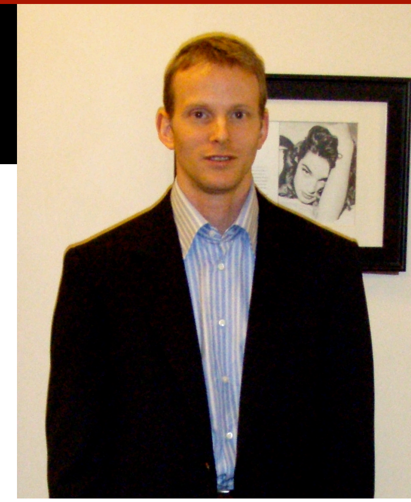

Research Vision:

Develop techniques to *Automatically* Check the World's Software for *Exploitable* Bugs

Example Current Projects

- Vetting the Security of Systems
- Decompilation
- Automatic Exploit Generation
- Malware Analysis

Example Project: Automatic Exploit Generation

Example Research Challenges

$$\forall x . P(x)$$

1. Formal Exploit Specifications
e.g., buffer overflow, information leaks

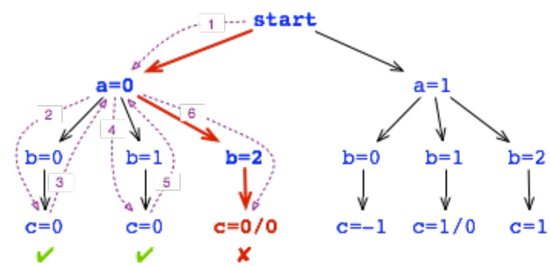

2. Binary Program Verification
e.g., Symbolic Execution

Windows 7 ubuntu

3. Exploits against real code
We get a shell in seconds

Lujo Bauer

Research interests: computer security and privacy, usable security

- Passwords

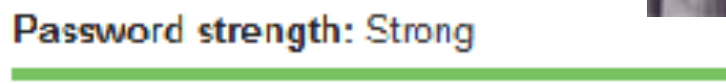

  - What makes a password hard for an attacker to guess but easy for a user to remember?
  - How to we induce users to create such passwords?
- Access control over personal data (a.k.a. privacy)
  - What kind of access-control / privacy policies do users want? What kind of policies are easy for users to understand?
  - What kind of interfaces are best for creating and editing policies?
- Application and system security
  - How do we prevent apps and plugins from stealing our information?
  - How do we build applications and systems that we can *prove* enforce users' policies correctly?

Nicolas Christin

Research Focus: Information security and policy, including security economics, network security, and usable security.

- Modeling behaviors and incentives in information security
  - Attackers
    - Understanding financial incentives, underground economies
  - Targets/Victims
    - Understanding biases, incentives
  - Combines network measurements, economics, psychology

- Integrating human factors in secure system design
  - Designing right set of incentives
  - Designing usable and secure systems
    - e.g., effect of password composition policies on security/usability

Nicolas Christin

Example focus area: Modeling Online Crime Supply Chains

- Understanding the motivations of the attackers
  - Increasing monetization of Internet security attacks
    - Financial incentive dominating attackers' motivations
  - Attacks are nowadays the product of multiple interactions between various parties
    - "Supply chains" of attacks
- *Overarching goal: Model these supply chains and understand them to better disrupt them*

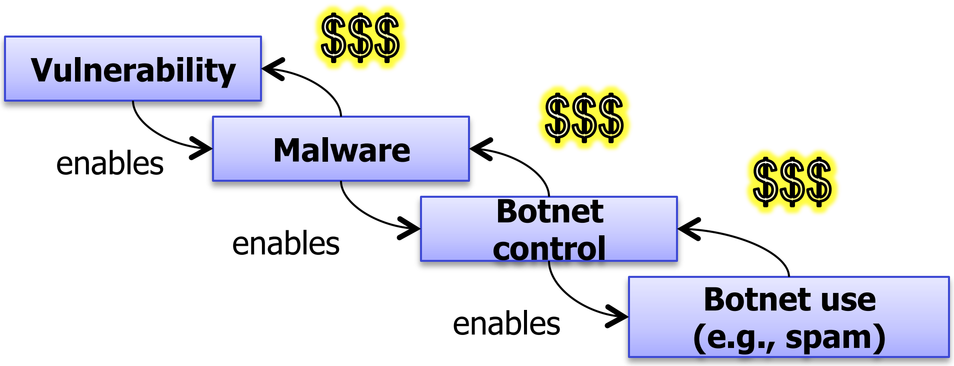

- Example: Japanese One Click Frauds
  - Gathered and analyzed 4 years worth of data
  - Found that top eight groups responsible for >50% of all frauds

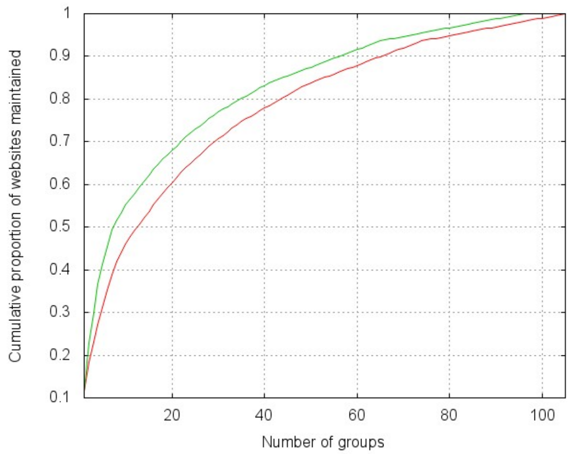

  - Found they could make **a lot** ($100K/yr+) of money with little risk or investment needed
    - Strong incentive to misbehave!

CMU is a Great Research Environment:

Collaborative, Open, Broad

Collaborative (and Social)

PDL Retreat 2011

Strong Industry Following, Collaboration, Feedback

Collaboration → High Impact → More Collaboration

An Exciting and Collaborative Place to Fundamentally Change the Way Future Computer Systems are Designed

You can change the world!

Computer Systems at CMU

As told by Onur Mutlu
(on behalf of many)
onur@cmu.edu
Carnegie Mellon University

ECE PhD Orientation 2013