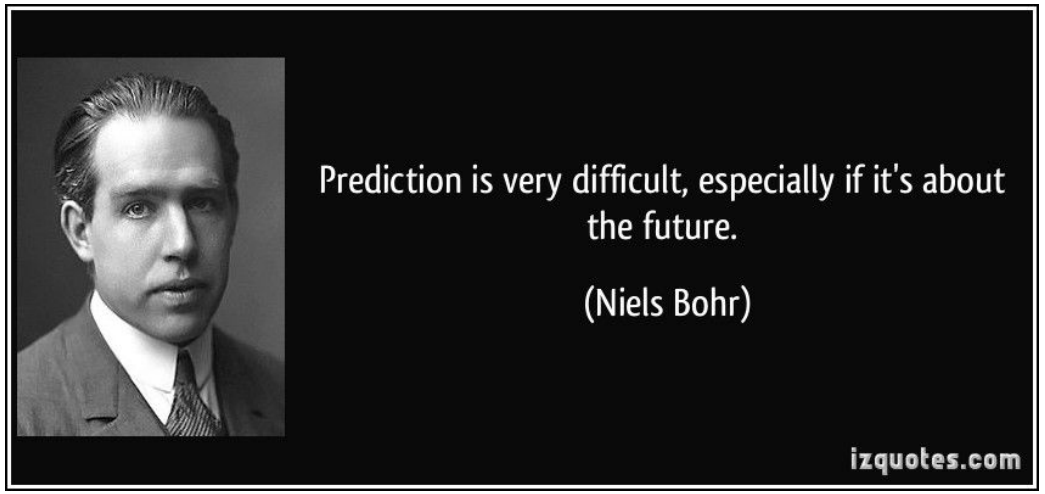

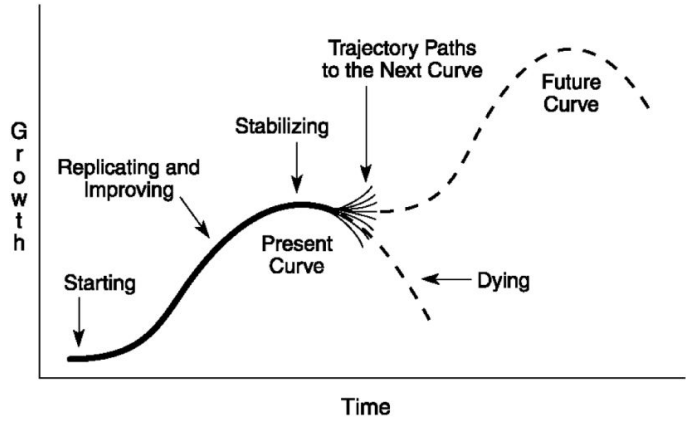

Architecture 2030 @ ISCA'16

Luis Ceze, Tom Wenisch

Mark Hill
(CCC liaison, mentor)

Neha Agarwal, Amrita Mazumdar, Aasheesh Kolli
(Student volunteers)

LIVE!

Context

- Many fantastic community formation/visioning workshops:
  - NSF ACAR, DARPA ISAT Future of Computer Systems without Technology Progress, IEEE Rebooting Computing, …
- These efforts have significant impact on community and funding

Workshop goals

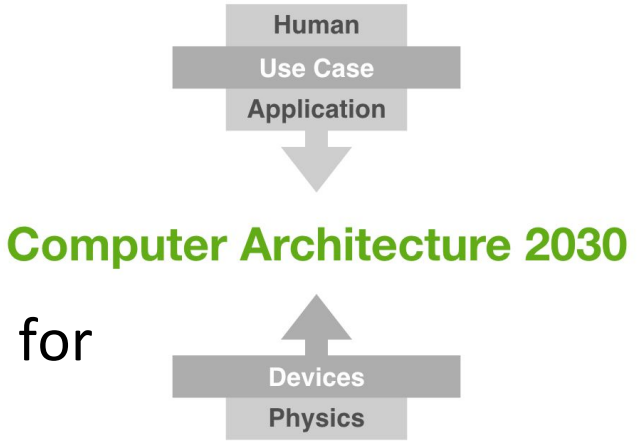

- Kick-off a visioning exercise for Computer Architecture research for the next 15 years
- Hear from Applications and Devices experts
- Increase visibility of architecture to broader CS and funding agencies
- Why now? A lot has changed in the last 5-10 years
  - Hardware design suddenly much more relevant but still (very) hard
  - Deep neural networks “caught us by surprise”, machine learning now a key workload
  - Major platforms emerged (cloud, IoT, etc)
  - Vertical integration (systems companies)
  - Explosion of sensor data (e.g., 1 trillion photos uploaded in 2015, genomics growing fast)
  - Open-source hardware emerging
- Seed a community white-paper similar to the 21st Century Architecture paper

Community input

- Talked to several members of the community
- Survey (~40 replies)

What are the most important challenges to be addressed by architecture research in the next 15 years? *

Long answer text

What are some of these challenges/trends that you are *not* working on? *

Long answer text

What are the main technology opportunities that the community should take advantage of? *

Long answer text

How well do you think the architecture community anticipated today's challenges/opportunities? *

1 2 3 4 5

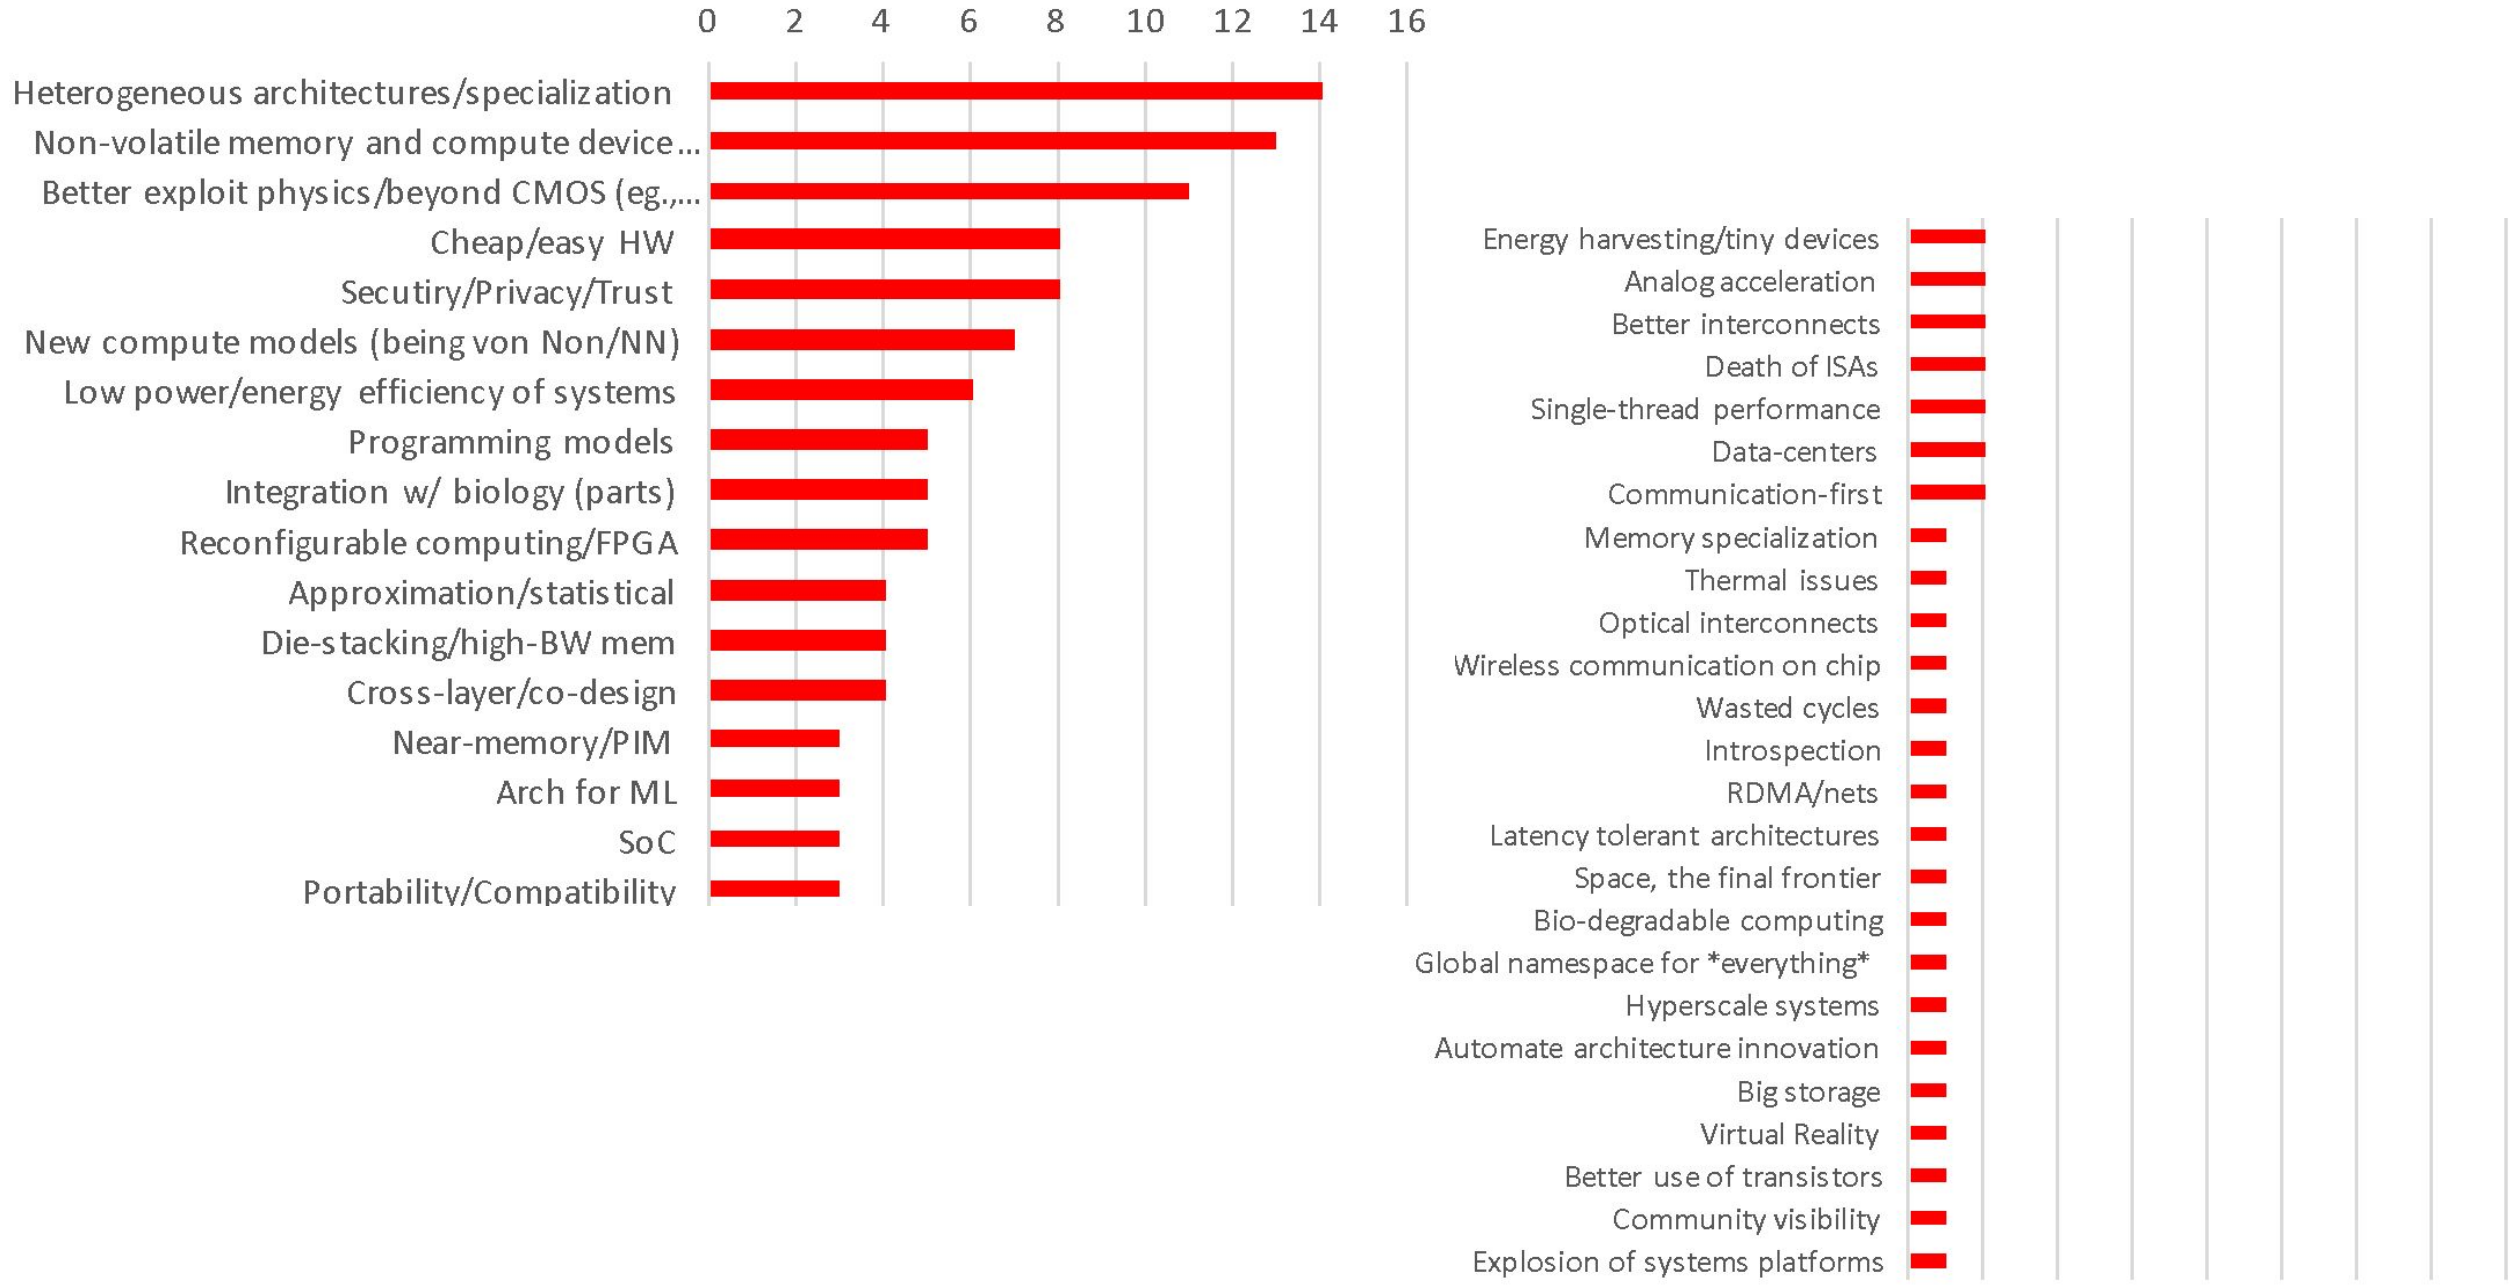
0
2
4
6
8
10
12
14
16
Heterogeneous architectures/specialization
Non-volatile memory and compute device...
Better exploit physics/beyond CMOS (eg.,...
Cheap/easy HW
Secutiry/Privacy/Trust
New compute models (being von Non/NN)
Low power/energy efficiency of systems
Programming models
Integration w/ biology (parts)
Reconfigurable computing/FPGA
Approximation/statistical
Die-stacking/high-BW mem
Cross-layer/co-design
Near-memory/PIM
Arch for ML
SoC
Portability/Compatibility
Energy harvesting/tiny devices
Analog acceleration
Better interconnects
Death of ISAs
Single-thread performance
Data-centers
Communication-first
Memory specialization
Thermal issues
Optical interconnects
Wireless communication on chip
Wasted cycles
Introspection
RDMA/nets
Latency tolerant architectures
Space, the final frontier
Bio-degradable computing
Global namespace for *everything*
Hyperscale systems
Automate architecture innovation
Big storage
Virtual Reality
Better use of transistors
Community visibility
Explosion of systems platforms

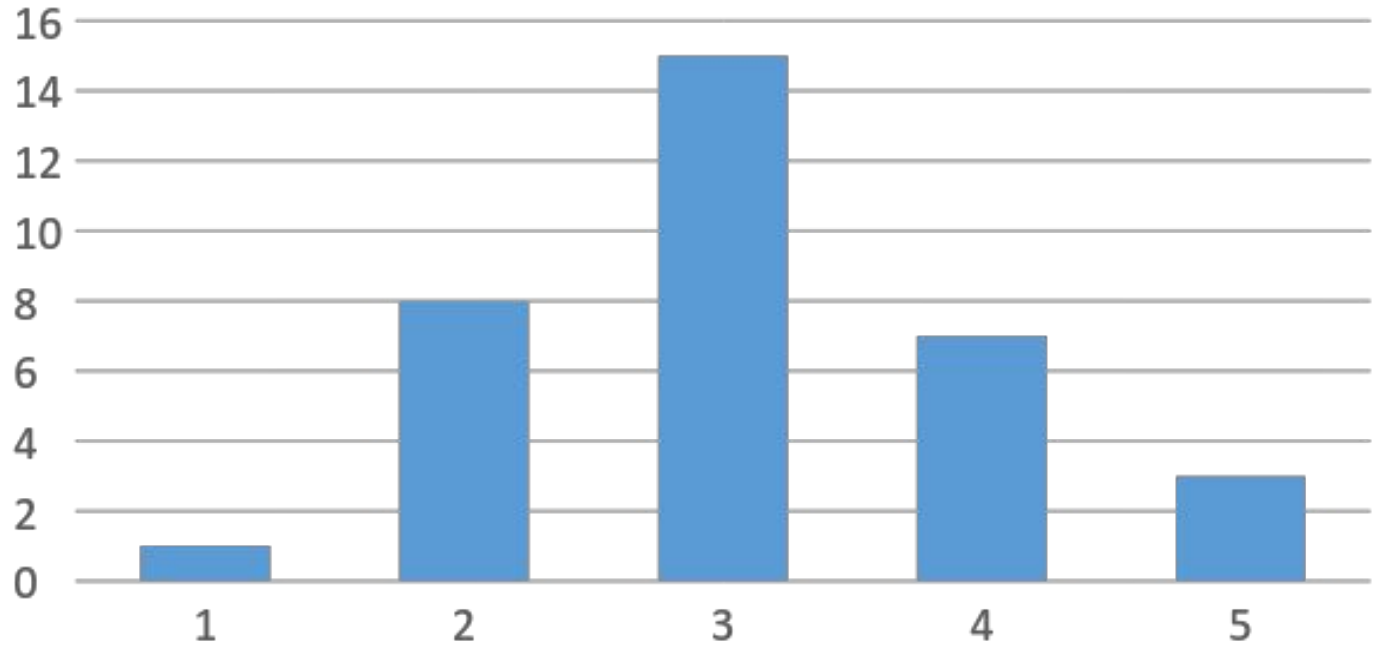

Bad ○ ○ ○ ○ ○ Great!

How well do you think the architecture community anticipated today's challenges/opportunities?

Big themes

- Making HW as easy to design/write as SW, open sourcing
- New devices/better exploitation of physics/biology
- Post-ISA era
- Post-Dennard/Post-Moore
- Vertical integration (systems companies)
- von Neuman is dead, long live von Neumann

Topic modeling analysis of our community's work

ISCA-1 to ISCA-42 (1974-2015)

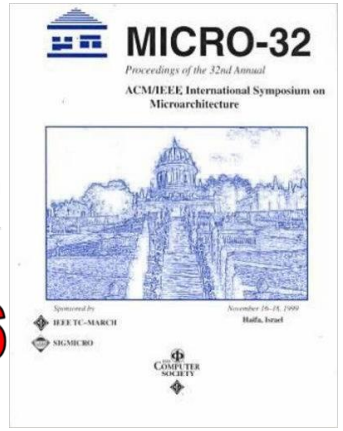

MICRO-6 to MICRO-47 (1972-2015)

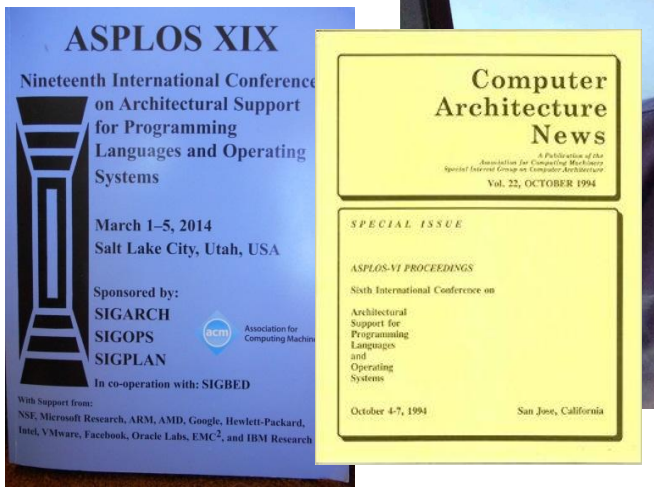

ASPLOS-1 to ASPLOS-20 (1982-2015)

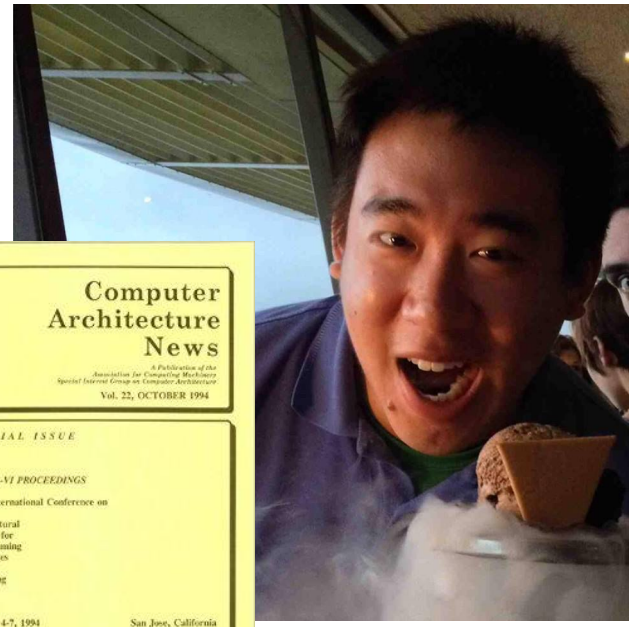

Vincent Lee, UW-CSE

What's in the corpus:

(1) All 3700 papers published on ACM from ISCA, MICRO, and ASPLOS from 1972 to 2015

(2) No workshop papers

Publication Corpus By the Numbers

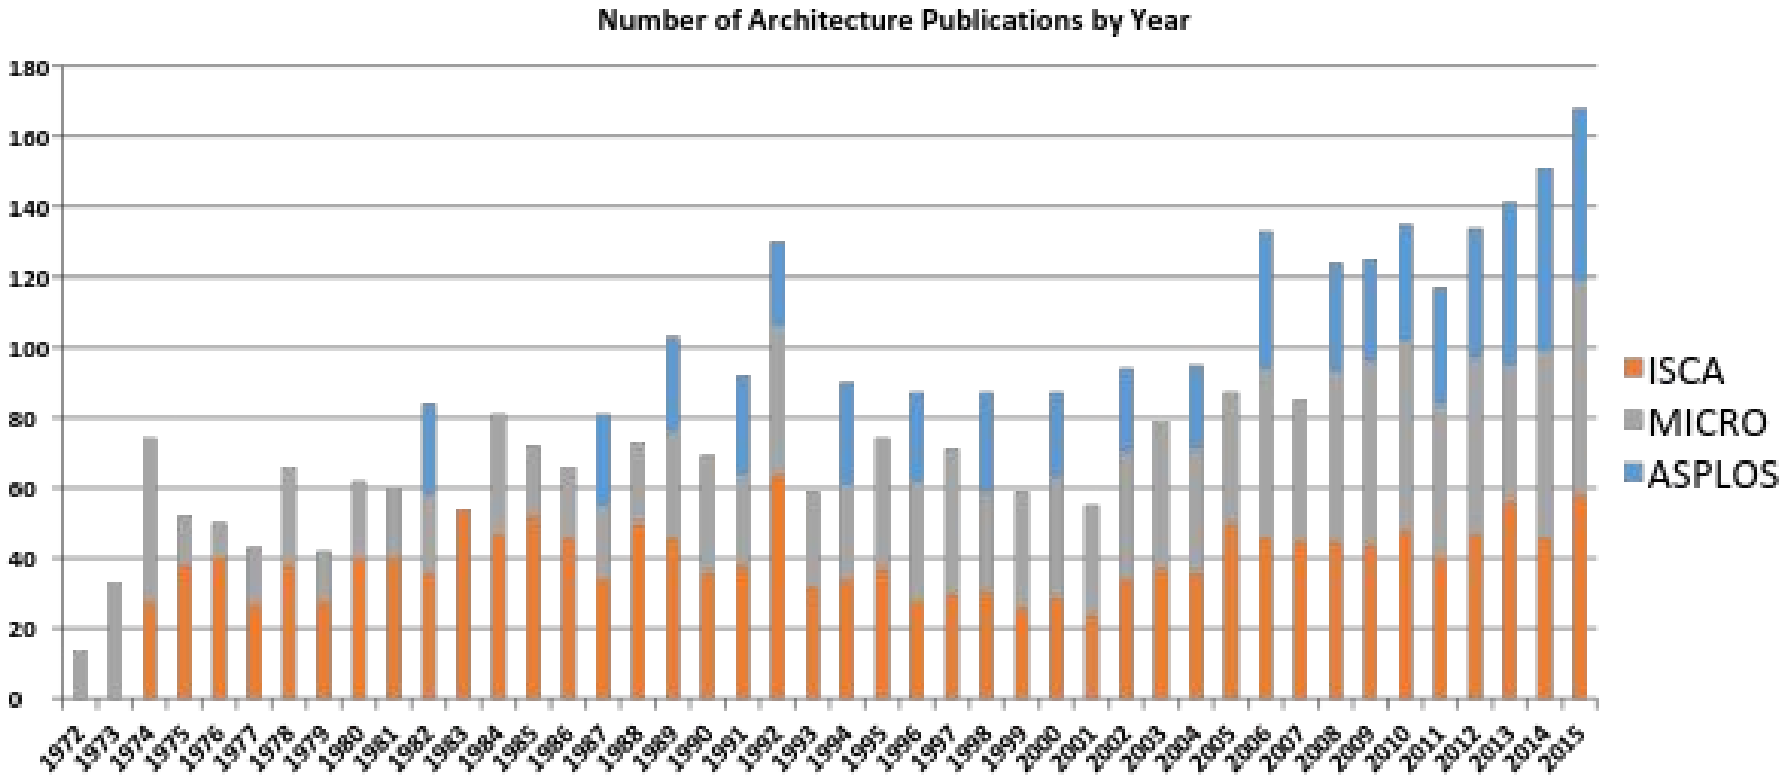

Comments:

(1) Number of publications in 1992 exceptionally high

(2) ASPLOS occurring every other year could potentially increase strength of topics during the year it took place

Clear Topics that Manifest in the Model

Quantum Computing	DRAM	Hardware Accelerators	Voltage Scaling
Fault Tolerance	Graphics	Cache Performance	Network Architectures
Database Architectures	Encryption	Microcoded Machines	Context Switching
Virtualization	Systolic Array Architectures	Compiler Optimizations	Neural Networks
Graph Processing	Datacenter Architectures	Replacement Policies	Network Interface Architectures
Branch Prediction	Prefetching	Die Stacked Memory	Floating Point
VLIW	Log Based Debugging	Microarchitecture	Memory Management
Cache Coherence	Memory Consistency	Scheduling	Concurrency Bugs
Energy Efficiency	Synthesis and Verification	NVM and Persistent Memory	Vector Processing

Starting with Our Roots…

Fundamental Microarchitecture Research

address, register, control, registers, operand, alu, instruction, memory, field, bit,

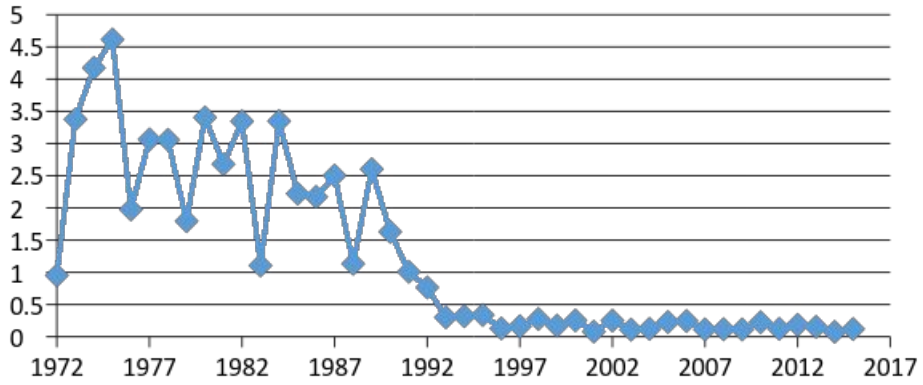

Microcoded Machines and Programs

micro, control, microcode, microprogram, microinstruction, compaction, mi, set, mo, fig,

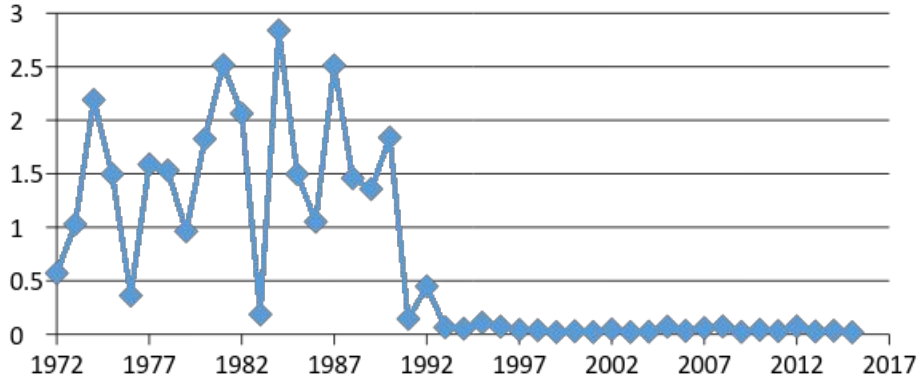

We Cared About Much Simpler Things…

computer, control, time, data, processing, speed, structure, significant, level, number,

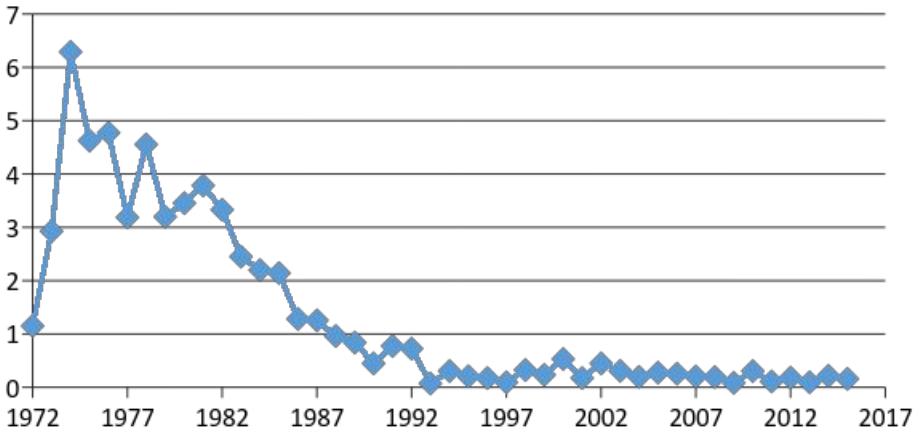

Vincent Lee, UW-CSE

Things That Trended then Died Off…

VLIW and Wide Issue Processors?

unit, units, functional, architecture, execution, cycle, multiple, issue, pp, vliw

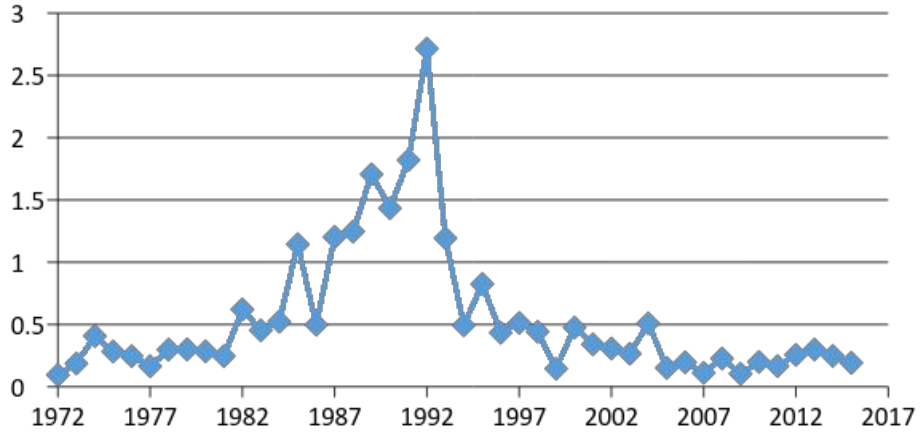

Support For Ancient Languages

prolog, goal, variable, list, lisp, uni, logic, clause, machine, cation

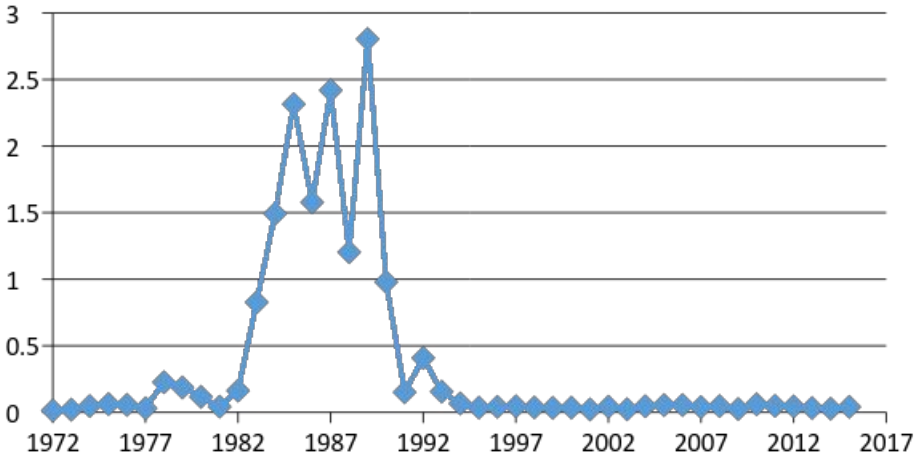

RISC vs. CISC Instruction Set Wars

instruction, vax, instructions, risc, mips, architecture, memory, set, byte, speci,

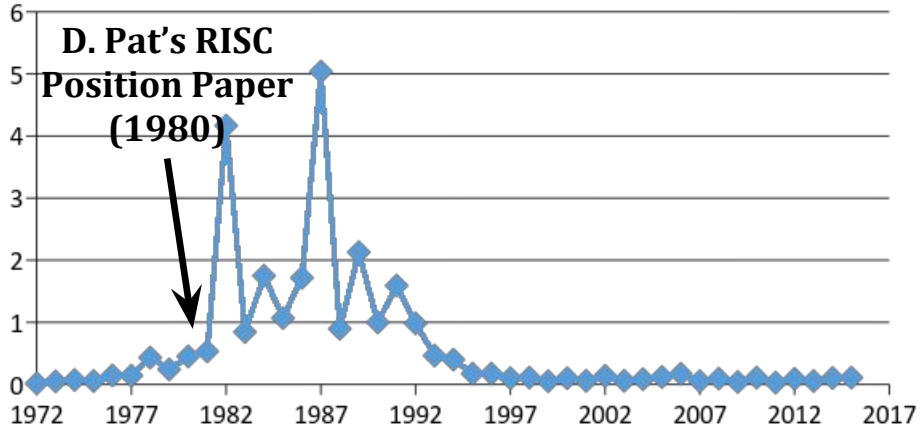

Branch Prediction

branch, branches, conditional, prediction, instructions, performance, slots, static, target, delay,

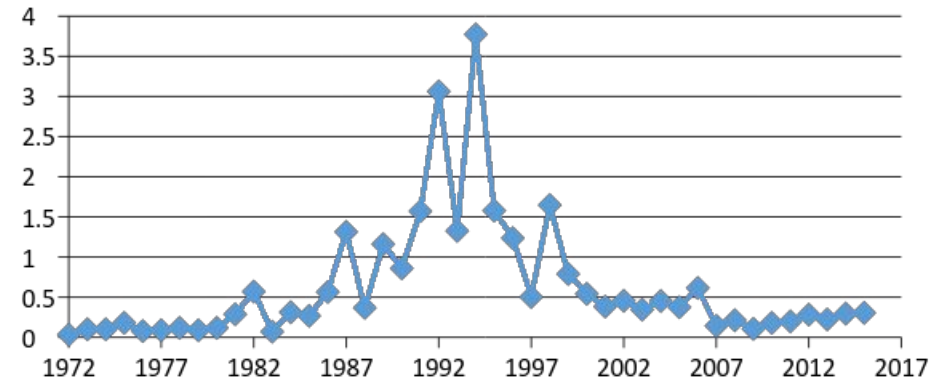

Trending Now Research Areas

Approximate Computing

approximate, error, quality, output, approximation, application, precision, applications, precise, input,

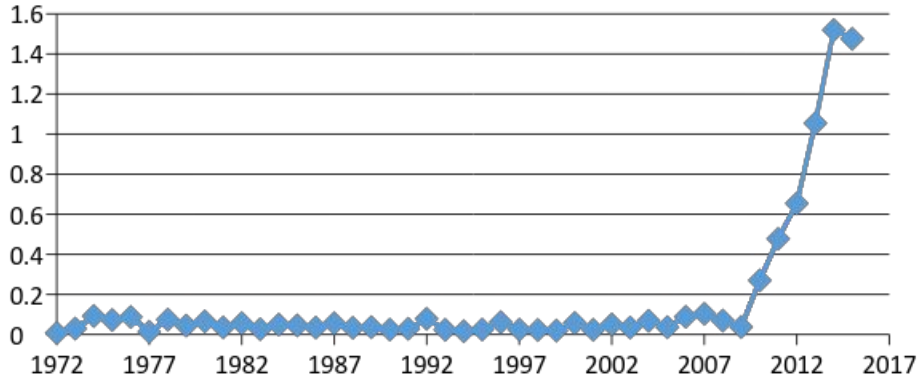

Hardware Security

side, channel, attacks, attacker, information, timing, signal, random, security, channels,

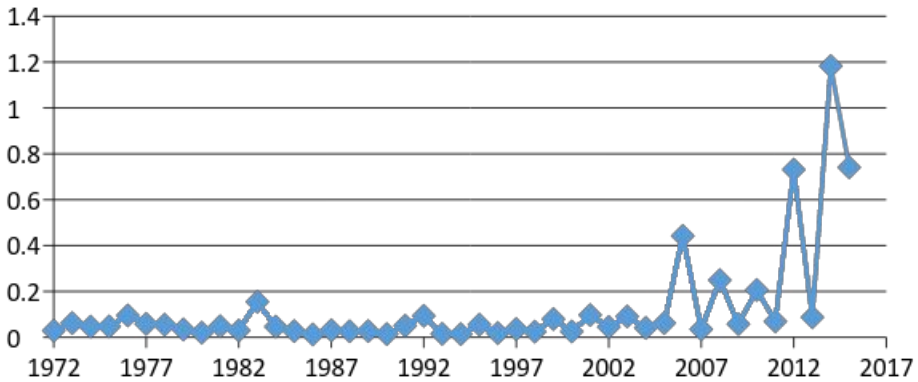

Accelerators

accelerator, fpga, hardware, gurable, application, design, accelerators, recon, ow, figure,

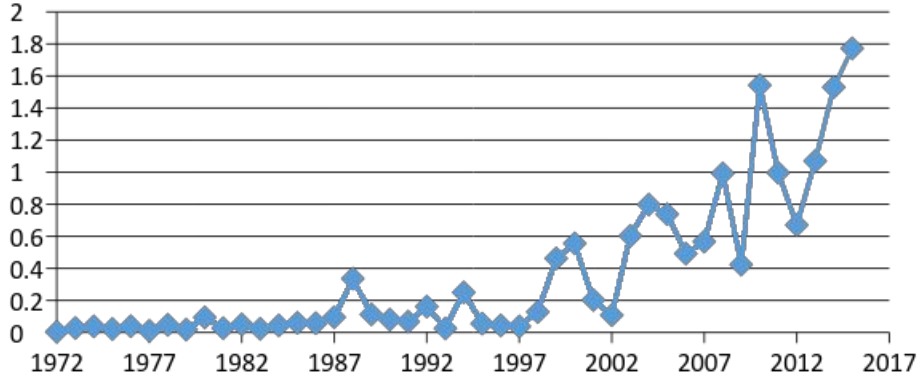

Datacenter Architectures

network, server, throughput, gb, nic, host, http, processing, interface, memcached,

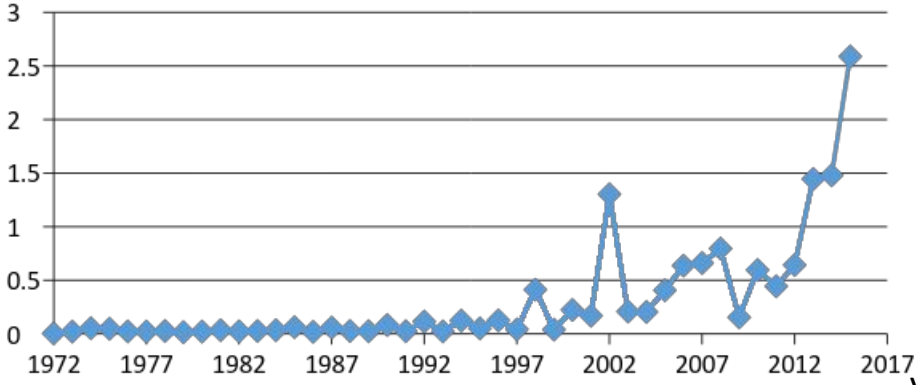

“Eternally Relevant” Research

Making Sequential Things Faster

parallel, parallelism, sequential, speedup, level, parallelization, execution, dependencies, serial, paral,

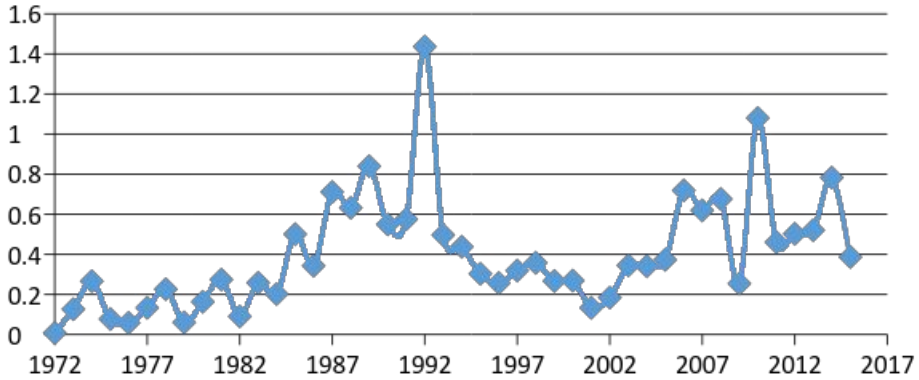

Graphs

node, graph, nodes, edges, edge, graphs, de, set, figure, number,

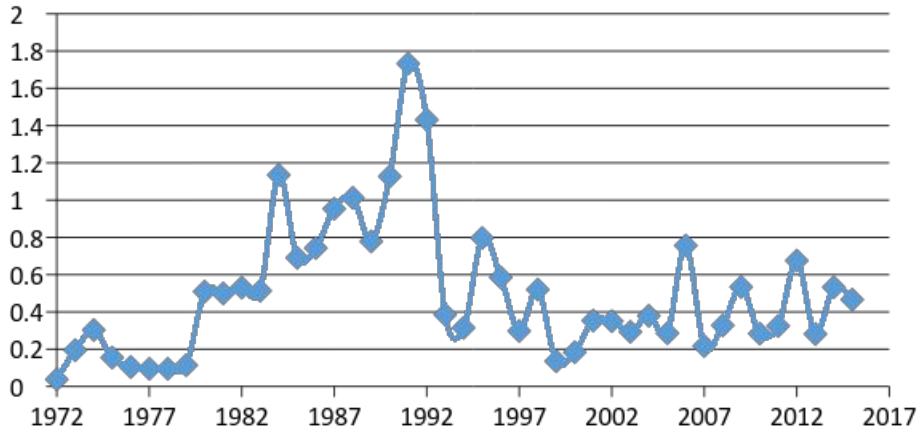

Caches

miss, misses, cache, direct, spec, size, data, mapped, cpi, benchmarks

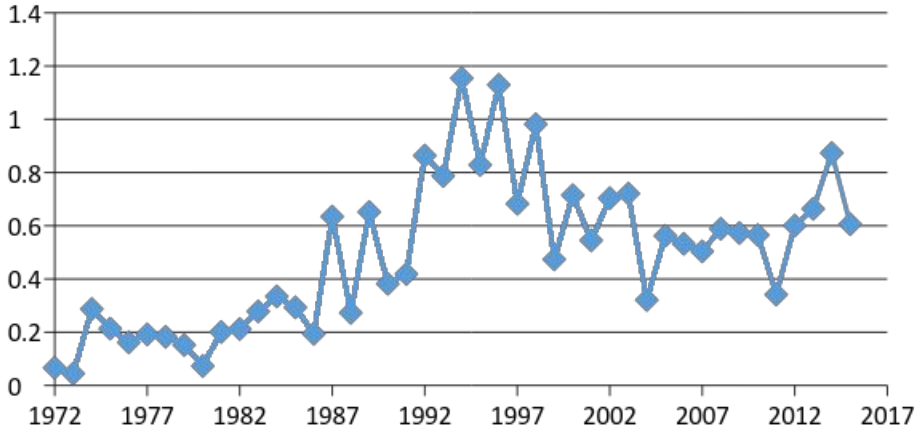

Vincent Lee, UW-CSE

Agenda

8:30	Intro remarks by Luis Ceze and Tom Wenisch
8:50	Mark Hill (Wisconsin) on “21st Century Computer Architecture”
9:10	Tom Conte (GeorgiaTech) on “IEEE Rebooting Computing Initiative & International Roadmap of Devices and Systems”
9:30	Devices Keynote: Philip Wong (Stanford) on “Device Technologies for the N3XT 1,000X Improvement in Computing Performance”
10:30	***Break***
11:00	Steve Keckler (nVidia/UT Austin) on “The Influence of Academic Research on Industry R&D”
11:25	Michael Taylor (UCSD) on “Open Source HW: Architecture’s Only Hope for Survival”
11:45	Alvy Lebeck (Duke) on “Computing and Biomolecules”
12:05	Yuan Xie (UCSB) on “Technology-driven Architecture Innovation: Challenges and Opportunities”
12:30	***Lunch***
14:00	Applications Keynote: Kayvon Fatahalian (CMU) on “100 Quadrillion Live Pixels: The Challenge of Continuously Interpreting, Organizing, and Generating the World’s Visual Information”
15:00	Breakout session kick off
15:30	***Coffee Break***
16:00	Break-out session
17:00	Report-out/discussion
17:30	Wrap-up

Arch 2030 Break-out Session

Break-out session discussion topic

- Articulate a grand-challenge/big-idea for the architecture community
  - Short description accessible to broader CS
  - What is the expected benefit if successful?
  - How will it push the field forward?
  - Which related disciplines will it draw from (PL, OS, ML, etc)? And how?

Example: DARPA grand challenges

- Smart collaborative spectrum allocation (2016)

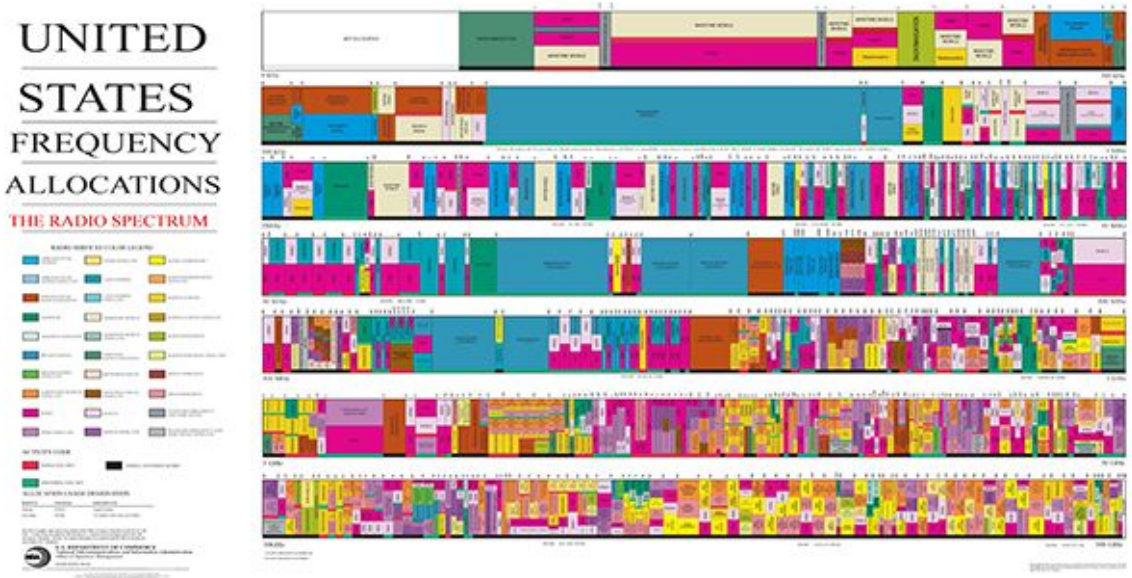

- Self-driving car in urban settings (2007)

Logistics

- Prep a 5-minute report-out
  - Please use the Google Slides template
  - https://goo.gl/ltWjWU
- 5 groups (you should have a paper with a number)
  - Groups 1 in this room, 2, 3, 4 and 5 along the hallway rooms
- Leaders:
  - 1. Ras Bodik
  - 2. Joel Emer
  - 3. Sarita Adve
  - 4. Babak Falsafi
  - 5. David Wood

Report-out: 5pm

15:00	Breakout session kick off
15:30	***Coffee Break***
16:00	Break-out session
17:00	Report-out/discussion
17:30	Wrap-up

Break-out session questions

- Articulate a grand-challenge for the architecture community
  - Short description accessible to broader CS
  - What is the expected benefit if successful?
  - How will it push the field forward?
  - Which related disciplines will it draw from (PL, OS, etc)? And how?

Heilmeier’s Catechism

1. What are you trying to do? Articulate your objectives using absolutely no jargon.
2. How is it done today, and what are the limits of current practice?
3. What's new in your approach and why do you think it will be successful?
4. Who cares? If you're successful, what difference will it make?
5. What are the risks and the payoffs?
6. How much will it cost? How long will it take?
7. What are the midterm and final assessments to check for success?