ISTC for Cloud Computing: Center Overview

Greg Ganger & Phil Gibbons

ISTC-CC Retreat 2014

September 4, 2014

http://www.istc-cc.cmu.edu/

ISTC-CC's Research Pillars

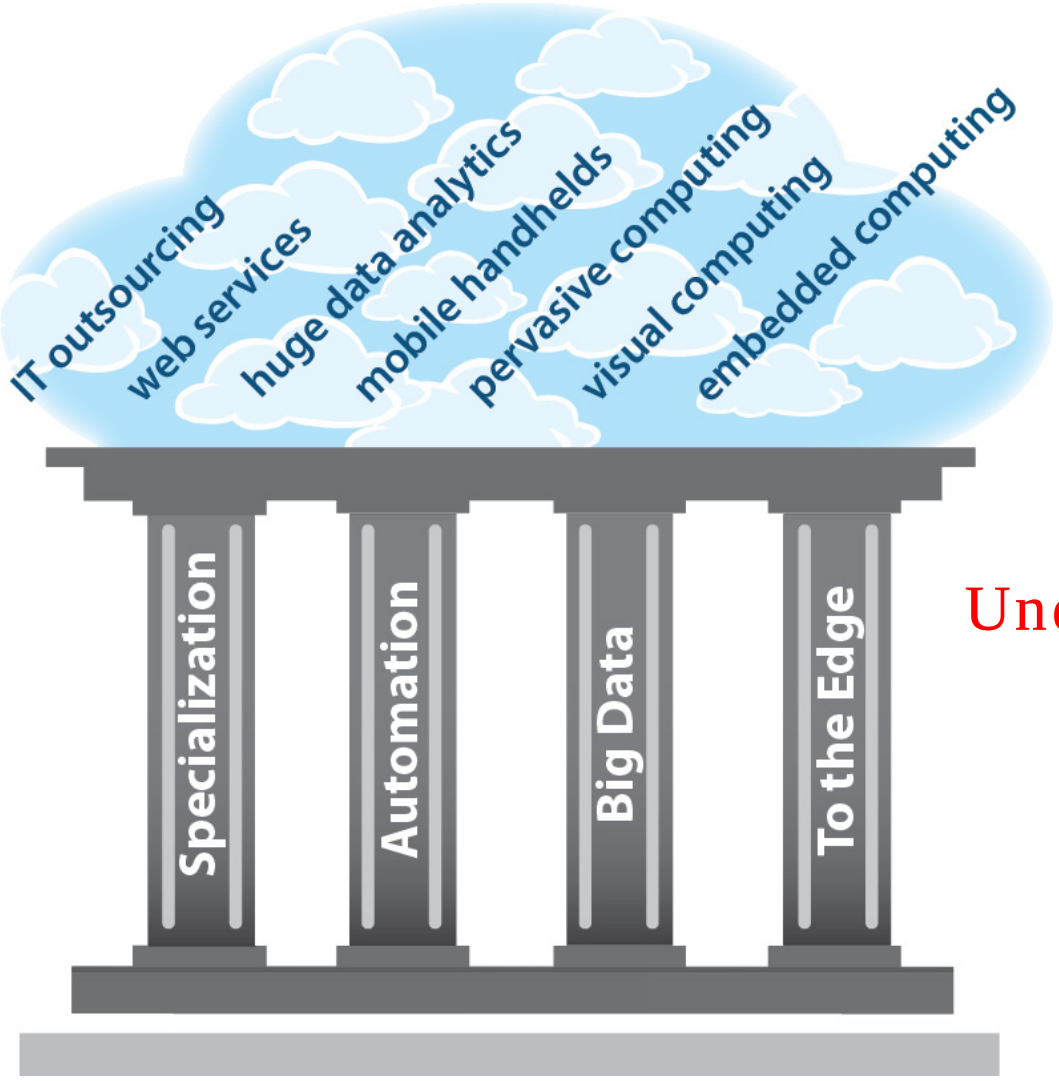

Underlying Infrastructure enabling the future of cloud computing

www.istc-cc.cmu.edu

Specialization Pillar

Low power nodes

Many-core

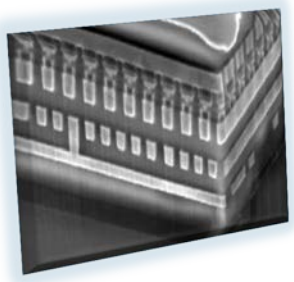

Phase-change memory (PCM)

- Specialization is fundamental to efficiency
  - No single platform best for all application types
  - Called **division of labor** in sociology
- Cloud computing must embrace specialization
  - As well as consequent heterogeneity and change-over-time
  - Stark contrast to common cloud thinking
- New approaches needed to enable…
  - Effective mixes of targeted and general platform types
  - Nimble incorporation of new technologies and accelerators

Specialization Projects

- **S1: Specialized Platforms of Wimpy Nodes**
  - exploring + extending range of apps that run (most) efficiently on such platforms by overcoming OS limits, memory limits, and scalability issues
- **S2: Specialized Platforms of Heterogeneous Multi-Cores**
  - exploring best ways to devise and use heterogeneity on multi-core nodes, considering core types, accelerators, and DRAM/NVM memory, with a focus on cloud's virtualized, multi-tenancy workloads

Specialization Highlights at Retreat 2014

- **"Hardware-Software Interface Issues in Heterogeneous Systems: Design, Verification, and Programming"** – Margaret Martonosi (Princeton)
- **"Cuckoo Hashing: From ISTC-CC to Intel's DPDK"** – Dong Zhou (CMU) & Ren Wang (Intel Labs)
- **"Some New Ideas in Memory System Design for Data-Intensive Computing"** – Onur Mutlu (CMU)
- **"Using RDMA Efficiently for Key-Value Services"** – Dave Andersen (CMU)
- **"Optimizing Performance and Productivity on Heterogeneous Processors"** – Sudha Yalamanchili (Georgia Tech)

+ 17 posters

Automation Pillar

- Automation is crucial to cloud reaching potential
  - We suspect that no one here needs to be convinced of this…
- Management is very hard, but cloud makes it worse
  - Much larger scale
  - Much more varied mix of applications/activities
  - Much less pre-knowledge of applications
  - And, we're adding in platform specialization ☺
- Leaps forward needed on many fronts…
  - Diagnosis, scheduling, instrumentation, isolation, tuning, …

Automation Projects

- **A1: Resource Scheduling for Heterogeneous Cloud Infrastructures**
  - maximizing the effectiveness of a cloud composed of diverse specialized platforms servicing diverse app types
  - enabling software framework specialization via hierarchical scheduling
- **A2: Problem Diagnosis and Mitigation**
  - new tools and techniques for rapid, robust diagnosis of failures and performance problems
  - automated mitigation based on “quick and dirty” online diagnoses

Automation Highlights at Retreat 2014

- **"Scheduling Heterogeneous Resources in Cloud Datacenters"** – Alexey Tumanov (CMU)

Posters:

- **PriorityMeister: Tail Latency QoS for Shared Networked Storage**
- **SpringFS: Bridging Agility and Performance in Elastic Distributed Storage**
- **Tetrisched: Space-Time Scheduling for Heterogeneous Datacenters**
- **The Power of Choice in Data-aware Cluster Scheduling**

Big Data Pillar

Customer Database

~600 TB

HD Internet Video

12 EB/yr

Particle Physics

300 EB/yr

- Extracting insights from large datasets
  - "Analytics" or "Data-intensive computing"
  - Becoming critical in nearly every domain
    - likely to dominate future cloud data centers
- Need right programming/execution models
  - For productivity, efficiency, and agility
  - Resource efficient operation on shared, specialized infrastructures

10^6 MEGA 10^9 GIGA 10^{12} TERA 10^{15} PETA 10^{18} EXA

Estimating the Exaflood, Discovery Institute, January 2008
Amassing Digital Fortunes, a Digital Storage Study, Consumer Electronic Association, March 2008

Big Data Projects

- **B1: Big Learning Systems**
  - new programming abstractions and execution frameworks enabling efficiency and productivity for large-scale Machine Learning on Big Data

- **B2: Big Data Storage**
  - exploring trade-offs and new approaches in Big Data storage, including support for high ingress and multi-framework sharing of data

Big Data Highlights at Retreat 2014

- **"Tachyon: A Reliable Memory Centric Storage for Big Data Analytics"** – Haoyuan Li (UC Berkeley)
- **"What's New with GraphLab"** – Yucheng Low (GraphLab)
- **"Spark: What's New and What's Next"** – Ion Stoica (UC Berkeley)
- **"Scaling Distributed Machine Learning with the | Parameter Server"** – Alex Smola (CMU)
- **"GraphX: Unifying Data-Parallel and Graph-Parallel Analytics"** – Joey Gonzalez (UC Berkeley)
- **"Exploiting Bounded Staleness to Speed Up Big Data Analytics"** – Garth Gibson (CMU)

+ 15 posters

To the Edge Pillar

- Edge devices will participate in cloud activities
  - Serving as bridge to physical world (sense/actuate)
  - Enhancing interactivity despite location / connectivity
- Need new programming/ execution models
  - For adaptive cloud
    + edge cooperation

Cloudlet demo

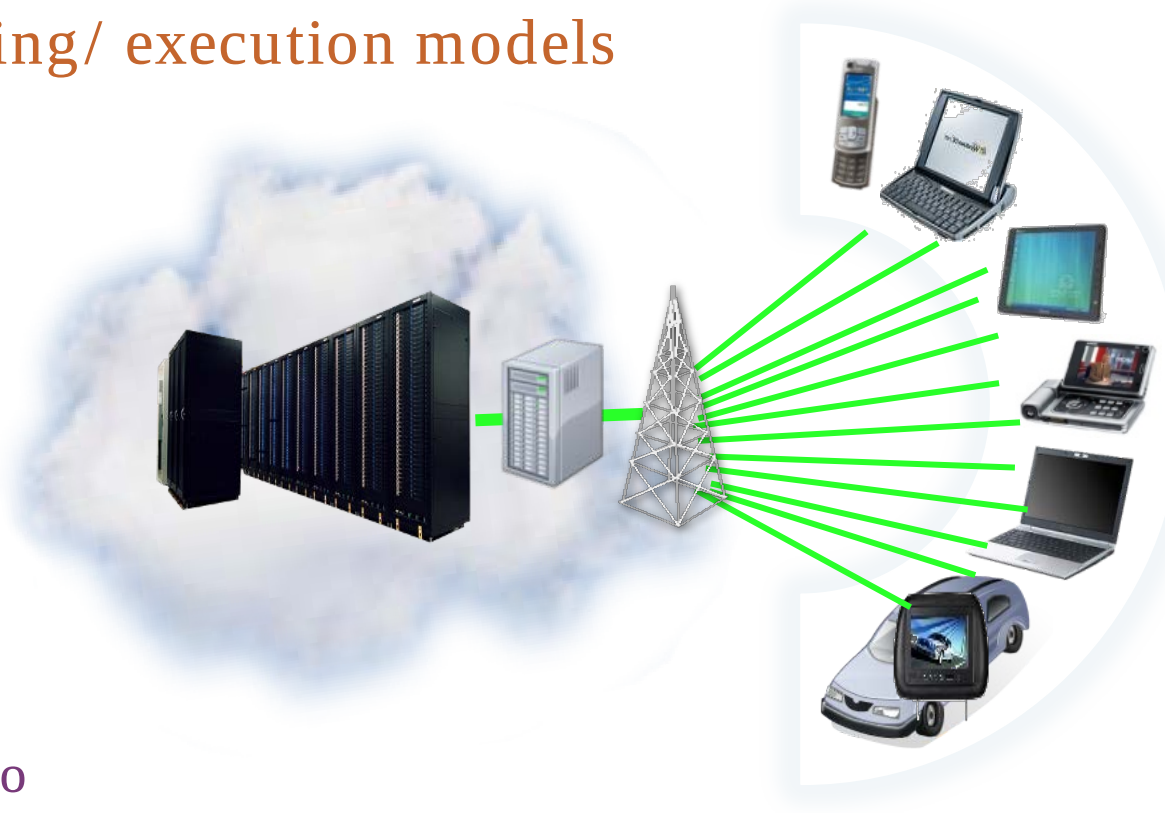

To the Edge Projects

- **E1: Cloud-Assisted Mobile Client Computations**
  - new abstractions and system architectures for dynamic exploitation of edge-local cloud resources to enable rich edge device experiences

- **E2: Geographically Distributed Data Storage**
  - new techniques for geographically distributed data storage/caching that reduce both access latency & reliance on expensive WAN-uplink bandwidth, while providing the desired scalability, fault tolerance, consistency & findability

To the Edge Highlights at Retreat 2014

- **"Personal Clouds: Sharing and Integrating Networked Resources to Enhance End User Experiences"**
  – Ada Gavrilovska (Georgia Tech)

- **"Towards Wearable Cognitive Assistance"**
  – Satya (CMU)

- Plus Friday Keynote:

  "Visual Cloud – Video delivery, Cloud Graphics and Visual Understanding" – Jim Blakley (Intel)

Toward ISTC-CC Capstones

Capstone #1: Big Learning Systems

- We're creating the leading frameworks for big learning
  - e.g., GraphLab, Spark, and Parameter Server systems
- Ideal: given an ML app + system, use best mechanisms
  - requires mapping the space and understanding tradeoffs

Capstone #2: Resource mgmt for specialization

- We've asserted need for mixes of specialized platforms
  - Significant progress on both exploiting & scheduling fronts
- Putting it together: Flexible apps on heterogeneous clouds

Capstone #3: Edge-Cloud Resources

- We've asserted need for cloudlets near edge
- Working with IL & BUs to create demo and pilots

ISTC-CC: Intel Oversight & BoA

- Intel oversight
  - Rich Uhlig (Executive Sponsor)
  - Scott Hahn (Managing Director)
  - Chris Ramming (Director of UCO)
  - Jeff Parkhurst (Program Director)

- Board of Advisors (including Rich and Scott)
  - Randy Bryant (Dean of School of CS, CMU)
  - Jeff Chase (Professor of CS, Duke)
  - Balint Fleischer (Gen. Mgr. Data Center Grp Arch., Intel)
  - Frans Kaashoek (Professor of CS&Eng, MIT)
  - Pradeep Khosla (Chancellor, UC San Diego)
  - Jason Waxman (Gen. Mgr. Cloud Platform Group, Intel)

ISTC-CC: Institutions & Faculty

- Carnegie Mellon University
  - Greg Ganger (PI), Dave Andersen, Guy Blelloch, Garth Gibson, Mor Harchol-Balter, Todd Mowry, Onur Mutlu, Priya Narasimhan, M. Satyanarayanan, Dan Siewiorek, Alex Smola, Eric Xing
- Georgia Tech
  - Greg Eisenhower, Ada Gavrilovska, Ling Liu, Calton Pu, Karsten Schwan, Matthew Wolf, Sudha Yalamanchili
- Princeton University
  - Mike Freedman, Margaret Martonosi
- University of California at Berkeley
  - Anthony Joseph, Randy Katz, Ion Stoica
- University of Washington
  - Carlos Guestrin
- Intel Labs
  - Phil Gibbons (PI), Michael Kaminsky, Mike Kozuch, Babu Pillai

Carnegie Mellon University

Georgia Tech

PRINCETON UNIVERSITY

UC Berkeley

UNIVERSITY of WASHINGTON

intel

ISTC-CC: Cross-Institution Collaboration

	Project	Personnel
S1	Specialized Platforms of Wimpy Nodes	**Andersen[C], Schwan[G], Freedman[P], Kaminsky[I], Kozuch[I], Pillai[I]**
S2	Specialized Platforms of Heterogeneous Many-Cores	**Mowry[C], Mutlu[C], Gavrilovska[G], Schwan[G], Yalamanchili[G], Martonosi[P], Gibbons[I], Kozuch[I]**
A1	Resource Scheduling for Heterogeneous Cloud Infrastructures	**Joseph[B], Katz[B], Stoica[B], Ganger[C], Harchol-Balter[C], Kozuch[I]**
A2	Problem Diagnosis and Mitigation	**Ganger[C], Narasimhan[C], Eisenhauer[G], Liu[G], Schwan[G], Wolf[G]**
B1	Big Learning Systems	**Stoica[B], Andersen[C], Blelloch[C], Ganger[C], Gibson[C], Smola[C], Xing[C], Guestrin[W], Gibbons[I]**
B2	Big Data Storage	**Stoica[B], Andersen[C], Ganger[C], Gibson[C], Xing[C], Pu[G], Schwan[G]**
E1	Cloud-Assisted Mobile Client Computations	**Satya[C], Siewiorek[C], Gavrilovska[G], Liu[G], Schwan[G], Martonosi[P], Pillai[I]**
E2	Geographically Distributed Data Storage	**Andersen[C], Satya[C], Siewiorek[C], Freedman[P], Kaminsky[I], Pillai[I]**

Year 3 Publication Highlights

- ICCD (Sep'13) – 2 papers
- SOCC (Oct'13) – 3
- SOSP (Nov'13) – 4
- MICRO (Dec'13) – 2
- HPCA (Feb'14) – 4
- FAST (Feb'14) – 2
- NSDI (Apr'14) – 3
- EuroSys (Apr'14) – 2
- CCGrid (May'14) – 2
- SIGMETRICS (Jun'14) – 2
- ISCA (Jun'14) – 2
- ATC (Jun'14) – 3 papers
- ICAC (Jun'14) – 2
- SPAA (Jun'14) – 3
- ICWS (Jul'14) – 2
- ICDCS (Jul'14) – 2
- CLOUD (Jul'14) – 4
- KDD (Aug'14) – 2
- SIGCOMM (Aug'14) – 2
- PACT (Aug'14) – 3

To appear:

- OSDI (Oct'14) – 3
- SC (Nov'14) – 2

98 published papers highlighted in
ISTC-CC Newsletter for Year 3

Open Source Code Releases in Year 3

Open Source page: www.istc-cc.cmu.edu/research/ossr/

- GraphBuilder 1.0 released open source in Jun'13
- GraphLab 2.2 released open source in Jul'13

- Spark 0.9 release Feb'14 – Apache incubator
- Tachyon 0.5 released open source in Jul'14

- Mesos 0.14 released Oct'13 – Apache

- Cuckoo hashing in Intel DPDK (internal)
  - Intel plans to release in DPDK v1.8 end of this year

- Other open source releases include:
  CBT, MICA, Concurrent Cuckoo, two Parameter Server systems, Eiger, EPaxos, Parrot, Cloudlet OpenStack++, CuckooFilter, RankSelect, MemC3, NVMalloc, etc.

Also Benchmarks page: www.istc-cc.cmu.edu/research/benchmarks/

Highlights of Year 3 Honors

- New IEEE Fellows – Gibbons, Gibson, Yalamanchili
  - ISTC-CC: 8 ACM Fellows, 8 IEEE Fellows, 2 NAE members
- 2 IEEE Design & Automation Conf. Awards – Siewiorek
- Steven J. Fenves Award – Ganger
- Program Chairs – Gibson (ATC), Kozuch (HotCloud), Schwan (Middleware & MBDS), Stoica (NSDI)
- Conference Chair – Kaminsky (SOSP)
- Best Paper Awards – Multu in HPCA'14 & RTAS'14, Satya in IC2E'14, Smola in KDD'14
- Fellowships – Chang (Intel), Meza & Subramanian (Bertucci)
- 2013 SPEC Distinguished Dissertation Award – Gandhi
- CMU SCS Outstanding Dissertation Award – Gonzalez
- Many grants awarded: Amplifying funding

  => Intel support for 25 students, yet working with 63

ISTC-CC Students @ Intel

Blue Badge Employees:

- Ameya Ambardekar – CMU
- Parag Dixit – CMU
- Soila Kavulya – CMU
- Min Lee – GA Tech

Intel Fellowship winners:

- Kevin Chang – CMU
- Dan Lustig – Princeton
- Yoongu Kim – CMU
- Priyanka Tembey – GA Tech
- Michelle Goodstein - CMU

Interns @ Intel:

- Hrishikesh Amur (2012) – GA Tech
- Nanley Cherry (2014) – GA Tech
- Chris Fallin (2012) – CMU
- Naila Farooqi (2014) – GA Tech
- Prasun Gera (2014) – GA Tech
- Liting Hu (2013) – GA Tech
- Sudarsun Kannan (2013) – GA Tech
- Yoongu Kim (2012) – CMU
- Xiaozhou Li (2013) – Princeton
- Jamie Liu (2012) – CMU
- Dan Lustig (2013) – Princeton
- Alex Merritt (2012,2013) – GA Tech
- Dipanjan Sengupta (2013) – GA Tech
- Vivek Seshadri (2012) – CMU
- Priyanka Tembey (2012) – GA Tech
- Nandita Vijaykumar (2014) – CMU
- Dong Zhou (2014) – CMU

Agenda for Rest of Today

[10:oo-11:15] Research Talks session #1

- Margaret Martonosi (Princeton), Ada Gavrilovska (GA Tech), Dong Zhou (CMU) / Ren Wang(Intel Labs)

[11:15-11:40] Break

[11:40-12:30] Research Talks session #2

- M. Satyanarayanan (CMU), Haoyuan Li (Berkeley)

[12:30-12:45] Poster session #1 Previews - Babu Pillai (Intel Labs)

[12:45-2:00] Lunch /Poster session #1

[2:00-2:15] Break

[2:15-3:30] Research Talks session #3

- Yucheng Low (GraphLab), Ion Stoica (Berkeley), Alex Smola (CMU)

[3:30-4:00] Break / Group Photo

[4:00-4:50] Research Talks session #4

- Onur Mutlu (CMU), Alexey Tumanov (CMU)

[4:50-5:05] Poster session #2 Previews – Michael Kaminsky (Intel Labs)

[5:15-8:30] Reception w/Poster session #2 & Dinner [JF5 Café/Patio]

Concluding Thoughts

Retreat Goal:

Benefit the research projects (direction & impact)

- Community building
- Brainstorming/feedback (but not ARs ☺) on:
  - ideas & approaches for tackling the research challenges
  - jump-start collaborations
  - synergies, connections & tech transfer opportunities
- Learn, share & have fun

Who to see about...

- Poster Previews/Sessions: Babu Pillai or Michael Kaminsky
- Logistics questions/issues: Tanya McCutchen