

Welcome to the ISTC-CC Retreat!

Phillip B. Gibbons
Intel Labs
Co-PI for ISTC-CC


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



Future Cloud Computing will:

- incorporate heterogeneous mixes of specialized platforms exploiting a variety of emerging technologies
- be dominated by big data analytics over stored and live data feeds
- encompass billions of edge devices through new paradigms for meshing clients and cloud
- rely on significant advances in automation to realize the desired efficiency and productivity

ISTC-CC Pillars Build Foundation



1: Specialization

Three initial ISTC-CC research foci:

- **S1: Specialized platforms of wimpy nodes**
- **S2: Specializing heterogeneous many-core platforms for cloud computing**
- **S3: Exploitation of NVM in cloud computing**

Specialization	Automation	Big Data	To the Edge
----------------	------------	----------	-------------

2: Automation

Four initial ISTC-CC research foci:

- **A1: Adaptive sizing/scheduling of elastic services**
- **A2: Run-time validation of cloud software and patches**
- **A3: Problem diagnosis tools & techniques**
- **A4: Dynamic scheduling of heterogeneous mixes of frameworks and services**

Specialization

Automation

Big Data

To the Edge

3: Big Data

Three initial ISTC-CC research foci:

- **B1: APIs and frameworks for advanced machine learning on Big Data**
- **B2: Characterization and better programming of Big Data applications**
- **B3: Online high ingest for Big Data cloud data stores**



4: To the Edge

Two initial ISTC-CC research foci:

- **E1: Adaptive work division among clients, edge servers, and core cloud**
- **E2: Efficient, client-effective use of limited wide-area bandwidth**

Specialization	Automation	Big Data	To the Edge
----------------	------------	----------	-------------

Key ISTC-CC leaders

- **PIs:** Greg Ganger (CMU), Phil Gibbons (Intel)
- **Intel oversight**
 - Wen-Hann Wang (Executive Sponsor)
 - Rich Uhlig (Managing Director)
 - Limor Fix (Director of APR)
 - Matt Haycock (Director of ISTC Program)
 - Jeff Parkhurst (Program Director)
- **Advisory Board (inc. Wen-Hann and Rich)**
 - Jason Waxman (Gen. Mgr. Cloud Infra. Group., Intel)
 - Balint Fleischer (Dir. DCG Planning and Arch., Intel)
 - Frans Kaashoek (Professor of CS&Eng, MIT)
 - Randy Bryant (Dean of School of CS, CMU)
 - Pradeep Khosla (Dean of Engineering, CMU)
- **On site Intel Labs Researchers**
 - Michael Kozuch, Babu Pillai, Michael Kaminsky



World-renown Cloud researchers

- 2 NAE members
- 7 ACM Fellows
- 5 IEEE Fellows
- Winners of ACM Dissertation Award, PECASE, IEEE Young Computer Architect Award, etc
- Leaders of Open Cirrus, etc.
- In first 3 months, papers in SOSP (3), SOCC (4), MICRO (5), ASPLOS (2), PPOPP, LISA, among others
 - Research agenda formulated in March/April
- Already > \$3M in amplifying funding

Today's Agenda

[8:50-9:30] Keynote #1: Jason Waxman (Intel, GM, Cloud Infrastructure Group)

[9:30-10:00] Poster previews (1 minute each) - Babu Pillai (Intel Labs) ★

[10:00-Noon] Poster session #1

[Noon-1:15] Lunch

[1:15-2:45] Research Talks session #1

- Hyeontaek Lim (CMU), “SILT: A memory-efficient, high performance key-value store”
- Matei Zaharia (UC Berkeley), “Scalable machine learning with Spark”
- M. Satyanarayanan (CMU), “Cloudlets as enablers of the post-PC world”

[2:45-3:15] Break

[3:15-4:45] Research Talks session #2

- Karsten Schwan (GA Tech), “Overview of ISTC cloud research at Georgia Tech”
- Mor Harchol-Balter (CMU), “Dynamic, robust auto-scaling of data center capacity”
- Wyatt Lloyd (Princeton), “Why settle for eventual: Scalable causal consistency...”

[4:45-6:00] Breakout session #1, 8 sessions ★

[6:00-6:30] Group Photo★ **then Transportation to dinner**

[6:30-8:00] Dinner at the Pittsburgh Athletic Club (+ Poster previews)★

[8:00-10:00] Poster session #2, Pittsburgh Athletic Club

Tomorrow's Agenda

[8:00-8:45] Continental breakfast

[8:45-9:30] Welcome - Greg Ganger (CMU) & Keynote #2:

- Rich Uhlig (Intel Labs, Manager, Systems Architecture Lab), "Optimizing for the Cloud: Tech Trends, Testbeds and Working Together"

[9:30-10:00] Break

[10:00-11:30] Research Talks session #3

- Andy Konwinski (UC Berkeley), "Mixed-framework data centers with Mesos"
- Swapnil Patil (CMU), "Performance debugging scalable table stores"
- Margaret Martonosi (Princeton), "Exploiting and improving the predictability of data center workloads"

[11:30-noon] Madness session (3 mins each) – Michael Kaminsky (Intel Labs) 

[noon-1:30] Lunch

[1:30-3:00] Research Talks session #4

- Justin Meza (CMU), "Locality-aware data placement in DRAM-PCM hybrid memory systems"
- Chengwei Wang (GA Tech), "A flexible system integrating monitoring and analytics for managing large-scale data centers"
- Raja Sambasivan (CMU), "Diagnosing performance changes by comparing request flows"

[3:00-4:00] Breakout session #2 

[4:00-4:30] Wrap-Up

Concluding Thoughts

Retreat Goal: **Benefit the research projects**

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

Who to see about...

- Poster previews: see Babu Pillai
- Minutes of Madness: see Michael Kaminsky
- Logistics questions/issues: see Jennifer Gabig

Jason Waxman bio

- Jason Waxman is the General Manager in Intel's Cloud Infrastructure Group. His role includes a focus on blade servers, Internet datacenters and technology for future dense data center architecture. He is also responsible for Intel's initiatives in Cloud Computing.

Waxman holds executive positions in industry design efforts including the board of Blade.org and the Server System Infrastructure Forum. He has previously served as the director for Intel's Xeon processor and chipset product lines and the related platform enabling and customer relationships, and has spent the last 11 years of his Intel career in Enterprise computing focused on server products and technologies involved in the introduction of over 12 new platforms.