

ISTC for Cloud Computing: Center Overview

Phil Gibbons (Intel-PI), Greg Ganger (CMU-PI)

August 27, 2015

Executive Sponsor: Rich Uhlig

Managing Sponsor: Scott Hahn

Program Director: Jeff Parkhurst

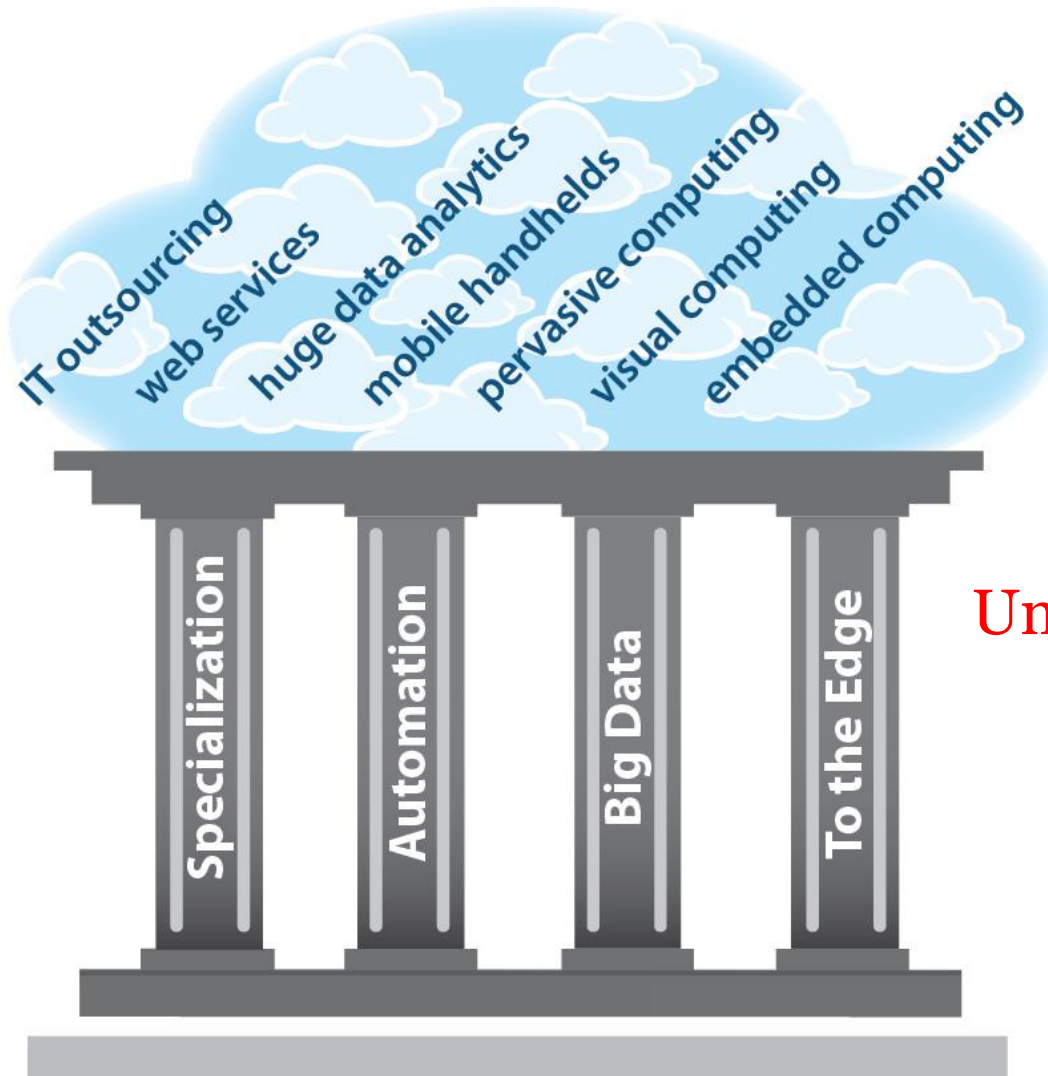
Embedded Researchers:

Michael Kozuch, Babu Pillai, Michael Kaminsky

<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 Highlights at Retreat 2015

Datacenters must embrace specialization (i.e., heterogeneity across nodes and within a node)

- **“Scaling Up Clustered Network Appliances with ScaleBricks”** – Hyeontaek Lim (CMU), Ren Wang (Intel), Michael Kaminsky (Intel @ ISTC)
- **“Software Specialization for Heterogeneous Computing Resources”** – Naila Farooqui (GA Tech)
- **“Architecting to Achieve a Billion Requests per Second Throughput on a Single Key-value Store Server Platform”** – Sheng Li (Intel) [short talk]
- **“Can DRAM Do More Than Just Store Data?”**
– Vivek Seshadri (CMU)

+ 8 posters

Automation Highlights at Retreat 2015

- **“Pulse: Telemetry Data Distribution in Datacenters”**
– Daniel Pittman & Joel Cooklin (Intel) [Thursday keynote]
- **“Highlights from Georgia Tech”** – Calton Pu (GA Tech)
- **“Scheduling Heterogeneous Resources in Cloud Datacenters”** – Greg Ganger (CMU)

Posters:

- PriorityMeister: Tail Latency QoS for Shared Networked Storage [short talk]
- Wrangler: Predictable and Faster Jobs in Distributed Processing Systems using Machine Learning [short talk]
- Tetrisched: Space-Time Soft Constraints in Heterogeneous Datacenters
- Affinity-Aware Work-Stealing for Integrated CPU-GPU Processors
- Stellula: An Integrated Information Plane for Federated Sites
- Shared Memory Optimization in Virtualized Cloud

Big Data Highlights at Retreat 2015

- **“Filesystems for Exascale Storage”** – Garth Gibson (CMU)
- **“System Support for Diverse ML Styles”**
– Eric Xing (CMU)
- **“Lightweight Processing on Compressed Graphs”**
– Guy Blelloch (CMU)
- **“Succinct: Enabling Queries on Compressed Data”**
– Rachit Agarwal (UC Berkeley)
- **“GraphLab: What’s New and What’s Next”**
– Carlos Guestrin (U. Washington)
- **“Cloud Computing Needs for Large-scale Visual Computing”** – Kayvon Fatahalian (CMU)
- **“Spark: What’s New and What’s Next”**
– Ion Stoica (UC Berkeley)

**+ 2 short talks
& 13 posters**

To the Edge Highlights at Retreat 2015

Cloud resources at the edge

- **“Near-edge Computing: Enabling Interactive Cloud Computing”** – Valerie Young (Intel) [Friday keynote]
- **“What is Edge Computing Good For?”**
– Satya (CMU)

Posters:

- Towards Real-World Deployment of Cloudlets
- Wearable Cognitive Assistance Applications on Cloudlets
- Adaptive VM Handoff Across Cloudlets

ISTC-CC Talks @ SoCC Aug 27-29

- **“Understanding Issue Correlations: A Case Study of the Hadoop System”**
- **“FastLane: Making Short Flows Shorter with Agile Drop Notification”**
- **“Reducing Replication Bandwidth for Distributed Document Databases”**
- **“ShardFS vs. IndexFS: Replication vs. Caching Strategies for Distributed Metadata Management in Cloud Storage Systems”**
- **“Managed Communication and Consistency for Fast Data-Parallel Iterative Analytics”**
- **“Using Data Transformations for Low-latency Time Series Analysis”**

20% of papers
@ SOCC

**a.k.a. ISTC-CC Retreat’s
“Parallel Session”**



Capstone #1: Big Learning Systems

Goal:

- Deliver world-leading open source Big Data Analytics frameworks,
- tools for auto-selecting among frameworks, and
- published “big picture” whitepaper providing a systems-oriented taxonomy mapping out the space far beyond current efforts

Value to Intel:

- Key enabler for democratizing Cloud:
 - Open source SW ecosystem on par with proprietary SW
- Big Data Analytics is a long-term driver of cloud market growth => a key component to Cloud SW ecosystem
- Also: Value of orchestrating among frameworks & understanding Big Picture

Capstone targets a key enabler for democratizing Cloud

Big Learning Frameworks @ ISTC-CC

Goal: Deliver world-leading open source Big Data Analytics frameworks

- **GraphLab (Dato Core)**

- Carlos Guestrin (U. Washington)



- **Spark**

- Ion Stoica (UC Berkeley)




- **Petuum (Stale Synchronous Parallel)**

- Eric Xing, Greg Ganger, Garth Gibson (CMU)
Phil Gibbons (embedded researcher)



- **Distributed ML in C++ (DMLC)**

- Alex Smola, Dave Andersen (CMU)



- **Ligra**

- Guy Blelloch (CMU)

ISTC-CC Analytics Publications in 2014

1. **On Model Parallelization and Scheduling Strategies for Distributed Machine Learning.** S. Lee, J. K. Kim, X. Zheng, Q. Ho, G. A. Gibson, E. P. Xing. Proceedings of 2014 Neural Information Processing Systems (**NIPS'14**), December 2014.
2. **Dependent Nonparametric Trees for Dynamic Hierarchical Clustering.** A. Dubey, Q. Ho, S. Williamson, and E. P. Xing. Proceedings of 2014 Neural Information Processing Systems (**NIPS'14**), December 2014.
3. **Exploiting Iterative-ness for Parallel ML Computations.** Henggang Cui, Alexey Tumanov, Jinliang Wei, Lianhong Xu, Wei Dai, Jesse Haber-Kucharsky, Qirong Ho, Greg R. Ganger, Phil B. Gibbons, Garth A. Gibson, Eric P. Xing. ACM Symposium on Cloud Computing 2014 (**SoCC'14**), November 2014.
4. **Fast Iterative Graph Computation: A Path Centric Approach.** Pingpeng Yuan, Wenya Zhang, Changfeng Xie, Hai Jin, Ling Liu and Kisung Lee. Proceedings of IEEE International Conference for High Performance Computing, Networking, Storage and Analysis (**SC'14**), November 2014.
5. **Scaling Distributed Machine Learning with the Parameter Server.** Mu Li, David G. Andersen, Jun Woo Park, Alexander J. Smola, Amr Ahmed, Vanja Josifovski, James Long, Eugene J. Shekita, and Bor-Yiing Su. Proceedings of 11th USENIX OSDI (**OSDI'14**), October 2014.
6. **GraphX: Graph Processing in a Distributed Dataflow Framework.** Joseph E. Gonzalez, Reynold S. Xin, Ankur Dave, Daniel Crankshaw, Michael J. Franklin, and Ion Stoica. Proceedings of 11th USENIX OSDI (**OSDI'14**), October 2014.
7. **Scaling Queries over Big RDF Graphs with Semantic Hash Partitioning.** Kisung Lee and Ling Liu, Proceedings of the 40th IEEE International Conference on Very Large Databases (**VLDB'14**), September 2014.
8. **Jointly Modeling Aspects, Ratings and Sentiments for Movie Recommendations.** Qiming Diao, Minghui Qiu, Chao-Yuan Wu, Alexander J. Smola, Jing Jiang, and Chong Wang. Proceedings of ACM SIGKDD Conference on Knowledge Discovery and Data Mining (**KDD'14**), August 2014.
9. **Efficient Mini-Batch Training for Stochastic Optimization.** Mu Li, Tong Zhang, Yuqiang Chen, and Alexander J. Smola. Proceedings of ACM SIGKDD Conference on Knowledge Discovery and Data Mining (**KDD'14**), August 2014.
10. **Reducing Data Loading Bottleneck with Coarse Feature Vectors for Large Scale Learning.** Shingo Takamatsu and Carlos Guestrin. **BigMine'14**, August 2014.
11. **GraphLens: Mining Enterprise Storage Workloads Using Graph Analytics,** Yang Zhou, Sangeetha Seshadri, Larry Chiu and Ling Liu, IEEE 2nd International Congress on Big Data (**Big Data'14**), June-July 2014.
12. **Exploiting Bounded Staleness to Speed up Big Data Analytics.** Henggang Cui, James Cipar, Qirong Ho, Jin Kyu Kim, Seunghak Lee, Abhimanu Kumar Jinliang Wei, Wei Dai, Gregory R. Ganger, Phillip B. Gibbons, Garth A. Gibson, Eric P. Xing. USENIX Annual Technical Conference (**ATC'14**), June 2014.
13. **Stochastic Gradient Hamiltonian Monte Carlo.** Tianqi Chen, Emily B. Fox, and Carlos Guestrin. Proceedings of 31st International Conference on Machine Learning (**ICML'14**), June 2014.
14. **Learning Everything about Anything: Webly-Supervised Visual Concept Learning.** Santosh Kumar Divvala, Ali Farhadi, and Carlos Guestrin. Proceedings of **CVPR'14**, June 2014.
15. **Sharp Threshold for Multivariate Multi-Response Linear Regression via Block Regularized Lasso.** W. Wang, Y. Liang and E. P. Xing, IEEE Transactions on Information Theory (**TOIT**), 2014.
16. **Toward Combining Online & Offline Management of Big Data Applications,** Brian Laub, Chengwei Wang, Karsten Schwan, and Chad Huneycutt, MBDS Track, ACM International Conference on Autonomic Computing (**ICAC'14**), June 2014.
17. **Exploring Graph Analytics for Cloud Troubleshooting.** Chengwei Wang, Karsten Schwan, Brian Laub, Mukil Kesavan, Ada Gavrilovska. International Conference on Autonomic Computing (**ICAC'14**), June 2014.
18. **Aggregation and Degradation in JetStream: Streaming Analytics in the Wide Area.** Ariel Rabkin, Matvey Arye, Siddhartha Sen, Vivek S. Pai, Michael J. Freedman. 11th USENIX Symposium on Networked Systems Design and Implementation (**NSDI'14**), April 2014.
19. **GRASS: Trimming Stragglers in Approximation Analytics.** Ganesh Ananthanarayanan, Michael Chien-Chun Hung, Xiaoqi Ren, Ion Stoica, Adam Wierman, Minlan Yu. 11th USENIX Symposium on Networked Systems Design and Implementation (**NSDI'14**), April 2014.

Capstone #2: Resource Management for Specialization

Goal:

- **Deliver leading open source resource management system for heterogeneous clouds**, which runs workload mixes over mix of specialized resources, maximizing aggregate effectiveness

Value to Intel:


- Recall: Key enabler for democratizing Cloud:
 - Open source SW ecosystem on par with proprietary SW
- **Orchestration is a key component to Cloud SW ecosystem**
- **Builds on & exposes Intel accelerators**

Capstone targets another key enabler for democratizing Cloud,
exposing IA value in Cloud

Capstone #3: Edge-Cloud Resources

Goal:

- Deliver world-leading open source SW for adaptive cloudlet (local cloud) exploitation,
- with wearable task-specific assistance use cases, and
- provide **technical leadership** for Intel-internal local cloud efforts



Value to Intel:

- **New market segment (local cloud) for servers**
 - “This may be the largest new vertical market segment which will need high-compute over the new few years.” – Ariel Moshkovitz (IL/STG)
- **New value proposition** for Intel’s Mobile Edge Computing (MEC) efforts in Telco space

Capstone targets new market opportunities

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



Highlights of Year 4 Honors

- **New IEEE Fellow** – Liu
 - ISTC-CC: 8 ACM Fellows, 9 IEEE Fellows, 2 NAE members
- **Marie R. Pistilli Women in EDA Achievement Award** – Martonosi
- **ACM Doctoral Dissertation Award** – Zaharia (Stoica student)
- **ACM Student Research Competition win** – Mutlu/Mowry advisees
- **Program Chairs** – Freedman (SoCC), Kozuch (Middleware industry track), Pu (CLOUD)
- **Keynotes** – Gibbons (x3), Liu (x2), Mutlu (x2), Yalamanchili, Gibson, Satya, Siewiorek, Harchol-Balter, Guestrin
- **Best Paper Awards** – Andersen/Kaminsky in SOCC'14 & ISCA'15, Gibson in SC'14, Xing in EMNLP'14, Yalamanchili in IISWC'14, Katz in LISA'14, Gavrilovska in HiPC'14, Liu in Mobiquitous'14, Mutlu in HPCA'15, Liu in CCGrid'15, Martonosi in Top Picks'14, Smola in KDD'14
- **Fellowships** – Mesa (Google), Park (Samsung)
- **Graduate Student Teaching Award** – Tumanov (CMU) for cloud class
- Many grants awarded: Amplifying funding
 - => **Intel support for 20 students, yet working with 62**

Year 4 Publication Highlights

- OSDI (Oct'14) – 3 papers
- ICCD (Oct'14) – 3
- IISWC (Oct'14) – 2
- SoCC (Nov'14) – 4
- SC (Nov'14) – 2
- MICRO (Dec'14) – 3
- NIPS (Dec'14) – 3
- HPCA (Feb'15) – 3
- VEE (Mar'15) – 2
- SIGMOD (May'15) – 3
- NSDI (May'15) – 3 papers
- ICWS (Jun'15) – 2
- CLOUD (Jun'15) – 3
- SIGMETRICS (Jun'15) – 2
- SPAA (Jun'15) – 2
- **ISCA (Jun'15) – 7**
- DSN (Jun'15) – 3
- SIGCOMM (Aug'15) – 3
- VLDB (Aug'15) – 2
- **SoCC (Aug'15) – 6**

95 published papers highlighted in
ISTC-CC Newsletter for Year 4

Open Source Code Releases in Year 4

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

- **GraphBuilder 1.0** released open source in Jun'13
- **Dato Core** released open source in early 2015
- **Spark 1.4.0** release mid-2015 – Apache incubator
- **Tachyon 0.7.1** released open source in Aug'15
- **Mesos 0.23** released July'15 – Apache
- **Cuckoo hashing** in Intel DPDK (internal)
 - Intel plans to release in DPDK v2.1 soon
- Other open source releases include:
CBT, MICA, Concurrent Cuckoo, two Parameter Server systems, IndexFS, Eiger, EPaxos, Parrot, Cloudlet OpenStack++, DMLC, CuckooFilter, RankSelect, MemC3, NVMalloc, etc.



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

ISTC-CC impact at Intel (Examples, 1/2)

ISTC-CC's research “really opens my horizon,” helping me and my group to “think about problems differently.” – Balint Fleischer, General Manager DCG (after 1st BoA mtg)

“We used [the FAWN team's] quantitative research in multiple forums as we made a case for a multi-atom low power server. It is clear that FAWN data was instrumental in securing a positive disposition for pursuing Avoton and its communications variant Rangeley.” – Matt Adiletta, Sr. Fellow DCG

ISTC-CC's GraphLab => Sponsor Lab => Graph Analytics Operation (GAO) in 2013 => IL Venture Tribeca => DCG/CPG Intel Analytics Toolkit in 2014 (GraphBuilder open source from Intel also an outcome)

CuckooSwitch [Michael Kaminsky (embedded researcher)] improves SDN Ethernet forwarding lookup by 15x memory efficiency & 2x throughput => Sponsor Lab in 2014 w/intern => DCG DPDK 2.1; Follow-on: scalable Switch Route Forward (S2RF) improves cluster thruput by 30% w/ low latency

ISTC-CC impact at Intel (Examples, 2/2)

- **Cloudelets:** near-edge cloud resources
 - Intel MEC team working with others (e.g., Vodaphone) on standardization efforts and prototypes
 - Intel Labs project being explored
- **TetriSched:** scheduling for heterogeneous clouds
 - Key enabler to democratizing clouds with diverse server capabilities, a core focus in DCG's CPG
 - Exploring collaboration with CPG's SDI team
 - Already strongly engaged with earlier Mesos effort
- **Spark:** fast, large-scale data processing
 - Exploded in community size, including Intel
 - SSG partnered with DataBricks to optimize Spark for IA
 - Used in many internal analytics efforts, e.g., improving SDN

ISTC-CC Students @ Intel

Blue Badge Employees:

- Ameya Ambardekar – CMU
- Parag Dixit – CMU
- Naila Farooqui – GA Tech
- Soila Kavulya – CMU
- Min Lee – GA Tech
- Lavanya Subramanian-CMU

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
- Kavita Chandrasekar (2015) – GA Tech
- Nanley Cherry (2014) – GA Tech
- Chris Fallin (2012) – CMU
- Naila Farooqui (2014,2015) – GA Tech
- Prasun Gera (2014) – GA Tech
- Tae Jun Ham (2015) – Princeton
- 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 (2013,2015) – GA Tech
- Vivek Seshadri (2012) – CMU
- Priyanka Tembey (2012) – GA Tech
- Blaise Tine (2015) – GA Tech
- Fish Tung (2015) – CMU
- Nandita Vijaykumar (2014) – CMU
- Manzil Zaheer (2014) – CMU
- Dong Zhou (2014) – CMU

Agenda for Rest of Today- update This

[10:00-11:00] Research Talks session #1

- Garth Gibson (CMU),
Hyeontaek Lim (CMU) / Ren Wang(Intel Labs) / Michael Kaminsky (Intel Labs)

[11:00-11:15] Break

[11:15-12:05] Research Talks session #2

- Eric Xing (CMU), Naila Farooqui (Georgia Tech)

[12:05-12:35] Keynote #1 – Daniel Pittman (Intel)



[12:35-1:05] Lunch (grab food and eat during talks)

[1:05-2:20] Research Talks session #3

- Guy Blelloch (CMU), Rachit Agarwal (Berkeley), Calton Pu (Georgia Tech)

[2:20-2:45] Break / Group Photo



[2:45-3:35] Research Talks session #4

- Carlos Guestrin (U.Washington), Kayvon Fatahalian (CMU)

[3:35-4:15] Short Talks and Poster session Previews – Babu Pillai



[4:15-5:30] Poster session [Pod]



[5:30] Reception & Dinner [Pod]

