Sponsoring Organization

ifip

In Cooperation

sigmicro

Bronze Sponsor Support

HUAWEI

IBM

NPC'2015
Call for Participation

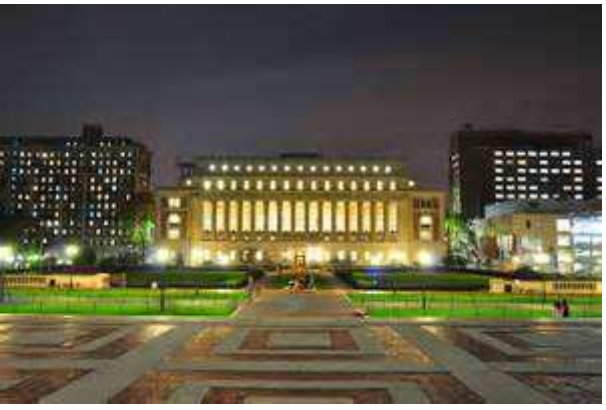

September 17-19, 2015

Columbia University
New York City

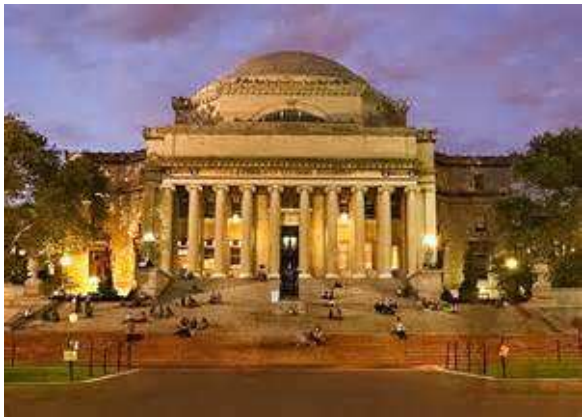

http://npc15.cs.umass.edu

Early registration extended to July 15

12th International Conference on Network and Parallel Computing

Keynote – Abstract

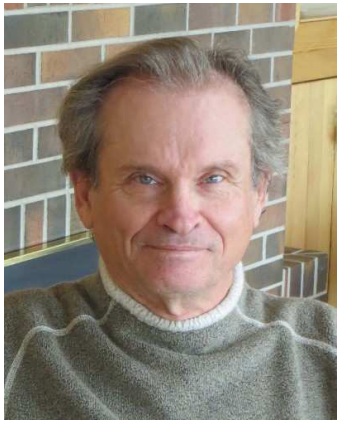

Investigating the Brain's Computational Paradigm

James E. Smith, Professor Emeritus – University of Wisconsin

Understanding and implementing the brain's computational paradigm is the grand challenge facing computer researchers. Not only does it provide computational capabilities far beyond those of conventional computers, its energy efficiency is truly remarkable. Furthermore, I believe strongly that computer architects and designers in particular have an important set of skills and a perspective that should be applied to meeting this grand challenge.

The brain's neocortex is constructed of massively interconnected neurons that compute and communicate via voltage spikes, and a strong argument can be made that precise spike timing is an essential element of the paradigm. I will describe biologically plausible computational elements based on precise spike timing. Through examples, I will illustrate some features of spike-based temporal computation and how it differs from other brain-inspired approaches. A machine learning architecture based on this model is under development and will be described in some detail. Perhaps the most important feature of this approach is that training times are orders of magnitude less than with conventional machine learning models. Although this is work in progress, it clearly illustrates the application of a computer designer's perspective to solving the ultimate computing grand challenge, and I hope that it motivates broader participation from computer architects and designers.

Keynote – Bio

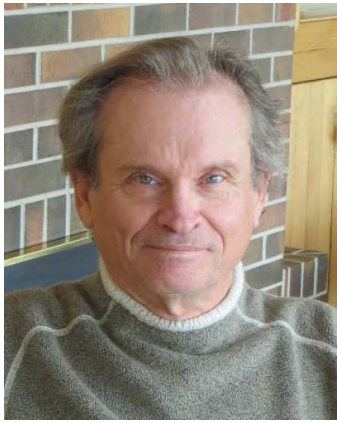

James E. Smith, Professor Emeritus – University of Wisconsin

James E. Smith is Professor Emeritus with the Department of Electrical and Computer Engineering at the University of Wisconsin-Madison. He received his PhD from the University of Illinois in 1976. He then joined the faculty of the University of Wisconsin-Madison, teaching and conducting research – first in fault-tolerant computing, then in computer architecture. He has been involved in a number of computer research and development projects both as a faculty member at Wisconsin and in industry.

Prof. Smith has made a number of significant contributions to the development of superscalar processors. These contributions include basic mechanisms for dynamic branch prediction and implementing precise traps. He has also studied vector processor architectures and worked on the development of innovative microarchitecture paradigms. He received the 1999 ACM/IEEE Eckert-Mauchly Award for these contributions. Today, almost every microprocessor makes use of these techniques from Prof Smith.

Currently, he is studying computational neuroscience at home along the Clark Fork near Missoula, Montana.

Keynote – Abstract

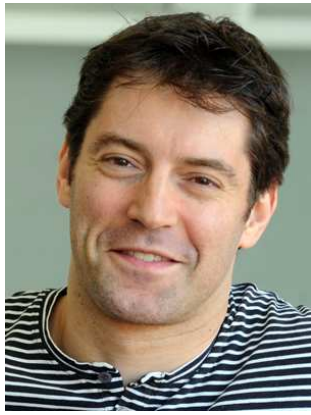

Energy-Harvesting Active Networked Tags (EnHANTs)

Dan Rubenstein,

Associate Professor of Computer Science, Columbia University

We are developing Energy-Harvesting Active Networked Tags (EnHANTs). EnHANTs are small, flexible, and energetically self-reliant devices that can be attached to objects that are traditionally not networked (e.g., books, furniture, walls, doors, toys, keys, clothing, and produce), thereby providing the infrastructure for various novel tracking applications. Examples of these applications include locating misplaced items, continuous monitoring of objects (items in a store, boxes in transit), and determining locations of disaster survivors.

In order for EnHANTs to rely on harvested energy, they have to spend significantly less energy than Bluetooth, Zigbee, and IEEE 802.15.4a devices. Moreover, the harvesting components and the ultra-low-power physical layer have special characteristics whose implications on the higher layers have yet to be studied (e.g., when using ultra-low-power circuits, the energy required to receive a bit is significantly higher than the energy required to transmit a bit). I will talk about our work in developing ultra-low power neighbor discovery algorithms, and some early ideas on higher-layer tracking protocols being developed for EnHANTs.

Keynote – Bio

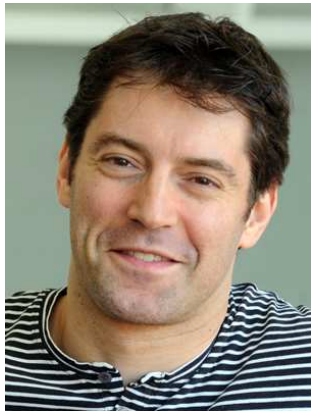

Dan Rubenstein,

Associate Professor of Computer Science, Columbia University

Dan Rubenstein is an associate professor in the Department of Computer Science at Columbia. His research interests are in network technologies, applications, and performance analysis. He is a former editor of IEEE/ACM Transactions on Networking, program chair of Association for Computing Machinery (ACM) Sigmetrics 2011, and has received an NSF CAREER Award, IBM Faculty Award, and several best paper awards. Rubenstein is also Co-Founder and former Chief Scientist at Infinio Systems, and has been a visiting scientist at Google.

Rubenstein received his Ph.D. in computer science from the University of Massachusetts, Amherst; an MA in mathematics from UCLA; and a BS in mathematics from MIT.

Keynote – Abstract

On the Future Evolution of Supercomputers in the Cloud

Kemal Ebcioğlu, President Global Supercomputing Corporation

Since the early days of mainframes, the sharing of hardware resources has provided benefits at least in terms of cost efficiency. Today, hardware resource sharing is best accomplished by service provisioning through Cloud Computing, which brings forth economies of scale, and which has become one of the most successful IT trends of our age. Emerging future applications such as the Internet of Things will furthermore have repercussions in the design of *supercomputers in the cloud*, tasked to control a massive amount of connected devices, and will continue to augment the realm of traditional HPC in the cloud beyond technical applications. Cloud Computing, whose market size has been predicted to exceed 200 billion dollars by 2020, has indeed become a *disruptive technology*, using Clayton Christensen's term, virtually replacing traditional in-house computing centers.

However, on one end of the spectrum of hardware resource sharing, there is the Application Specific Integrated Circuit, specifically designed for accelerating a single instance of a single application, which achieves low power and high performance by eliminating all interpretation overhead; while on the other end of the same spectrum, there is the general-purpose multi-core processor, burdened with the design complexity, power wall, memory wall, and frequency wall difficulties inherent within the highly parallel, pipelined interpretation of processor instructions. Today's cloud computing industry is focused on the latter, interpretive end of the hardware resource sharing spectrum, which is far from providing energy efficiency; in fact, large cloud data centers have already been reported to require hundreds of megawatts of power. Furthermore, today's software-based hypervisors/operating systems running on such multi-core processors are not scalable enough to support massive parallelism. But between the two endpoints of the hardware resource sharing spectrum, there is an unexplored, vast research area, which we believe to be the basis for power-efficient hardware resource sharing in future massively parallel cloud computing data centers. It is this unexplored area that our research vision is focused on. In this talk, we will describe a vision comprising next generation, easy-to-create, power efficient, custom hardware-accelerated cloud computing data centers with exascale computation capabilities.

Keynote – Bio

Kemal Ebcioğlu, President Global Supercomputing Corporation

Kemal Ebcioğlu conducted research on architectures, compilers, and languages for fine-grain parallelism at the IBM T.J. Watson Research Center, Yorktown Heights, NY, from 1986 to 2005. Dr. Ebcioğlu proposed, launched, and led pioneering IBM Research projects on fine-grain parallel architectures, including VLIW (Very Long Instruction Word) and DAISY (Dynamically Architected Instruction Set from Yorktown), a binary translation project. His last position at IBM was co-leader of Programming Model and Tools, a 40-person group that was part of a US Defense Advanced Research Projects Agency-funded IBM supercomputer research project, emphasizing high programmer productivity for HPC.

Dr. Ebcioğlu received two IBM Outstanding Technical Achievement awards, and an IBM Divisional award. In 2006, he retired from IBM and founded Global Supercomputing Corporation, where he currently is president. Ebcioğlu received a Ph.D. degree in computer science from the State University of New York at Buffalo in 1986.

Dr. Ebcioğlu has over 70 technical publications and 12 US patents. He has served as the International Federation for Information Processing Working Group 10.3 (Concurrent Systems) Chair in the period 2001-2006, and as the ACM Special Interest Group on Microarchitecture (SIGMICRO) Chair in the period 1999-2005. He has served as general chair, program chair, and steering committee chair for various conferences related to fine grain parallelism.

Dr. Ebcioğlu received the IEEE Computer Society B. Ramakrishna Rau Award in 2013, which is presented in recognition of substantial contributions in the field of computer microarchitecture and compiler code generation.

Ebcioğlu's present research interests include parallel scalable cloud computing and virtualization, high-productivity exascale systems, overcoming the memory wall barrier, and dynamic binary translation and optimization.

Panel

Resolved: "Everyone doing network and parallelism research should be doing it on mobile."

Panelists:

- David F. Bacon, Google
- Alan Bivens, IBM
- Jason Mars, University of Michigan
- Vijay Janapa Reddi, University of Texas, Austin
- Craig Stunkel, IBM
- Vijayalakshmi Srinivasan **Moderator**

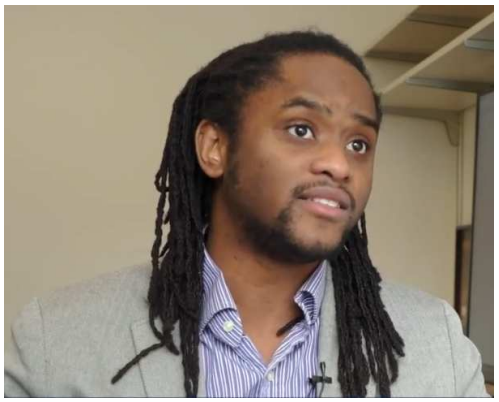

Technical Papers and Posters - 1

Program Committee: http://npc15.cs.umass.edu/organization

SESSION 1: NETWORKS

HyperFatTree : A Large-scale Tree-based Network with Low-radix Switches,
Yong Su, Zheng Cao, Zhiguo Fan, Zhan Wan, Xiaoli Liu, En Shao, Xuejun An and Ninghui Sun.

An Opportunistic Network Coding Routing for Opportunistic Networks,
Jiansheng Yao, Chunguang Ma, Gang Du and Qi Yuan.

Parallel Algorithms for Generating Random Networks with Given Degree Sequences,
Maksudul Alam and Maleq Khan.

SESSION 2: APPLICATIONS: SCALABILITY, VIRTUALIZATION AND SECURITY

CovertInspector: Identification of Shared Memory Covert Timing Channel in Multi-tenanted Cloud,
Sheng Wang, Weizhong Qiang, Hai Jin and Jinfeng Yuan.

Vshadow: Promoting Physical Servers into Virtualization World,
Song Wu, Yongchang Li, Xinhou Wang, Hai Jin and Hanhua Chen.

DMR: A Deterministic MapReduce for Multicore Systems,
Yu Zhang and Huifang Cao.

SHORT PAPERS + POSTERS

Optimal Deployment Model of Key Pre-distribution Protocol for Heterogeneous Wireless Sensor Networks,
Qi Yuan, Chunguang Ma, Xiaorui Zhong, Peng Wu

Hummer: Mitigating Stragglers with Partial Clones,
Li Jia, Wang Changjian, Li Dongsheng

Testing of SDN Applications for Design Flaws and Implementation Bugs,
Jiangyuan Yao, Zhiliang Wang, Xia Yin, Jun Bi, Xingang Shi, Jianping Wu, Yahui Li

Technical Papers and Posters - 2

Program Committee: http://npc15.cs.umass.edu/organization

SESSION 3: MULTICORES AND ACCELERATORS

RECU: Rochester Elastic Cache Utility -- Unequal Cache Sharing is Good Economics,
Chencheng Ye, Jacob Brock, Chen Ding and Hai Jin.

An Optimal Page-Level Power Management Strategy in PCM-DRAM Hybrid Memory,
Jinbao Zhang, Xiaofei Liao, Hai Jin, Dong Liu, Li Lin and Kao Zhao.

Instruction Fusion for Multiscalar and Many-Core Processors,
Yaojie Lu and Sotirios Ziavras.

Two-level Task Scheduling for Irregular Applications in GPU Platform,
Jing Li, Lei Liu, Yuan Wu, Xiaobing Feng and Chengyong Wu.

SESSION 4: PARALLELISM

Hierarchical Read-write Optimizations for Scientific Applications with Multi-variable Structured Datasets,
Preeti Malakar and Venkatram Vishwanath.

Performance Evaluation and Enhancement of Process-Based Parallel Loop Executions,
Xingjing Lu, Long Chen and Zhiyuan Li.

Determinism at Standard Library Level in Transactional Memory Based Applications,
Vesna Smiljkovic, Osman Unsal, Adrian Cristal and Mateo Valero.

Thursday, September 17

7:45 –	8:30	Registration
8:30 –	8:45	Opening Remarks *(General Co-Chairs, Program Co-Chairs)*
8:45 –	10:00	**Keynote 1**: James E. Smith
10:00 –	10:20	Break
10:20 –	11:40	**Session 1**: Networks
11:40 –	1:15	Lunch
1:15 –	2:30	**Keynote 2**: Dan Rubenstein
2:30 –	2:45	Break
2:45 –	4:05	**Session 2**: Apps: Scalability, Virtualization & Security
4:05 –	5:15	Poster Session *(10-minute presentations, then authors at posters)*

TECHNICAL PAPER PRESENTATIONS: 25 minutes each, including questions

Friday, September 18

8:00 – 9:00 Registration

9:00 – 10:15 **Keynote 3**: Kemal Ebcioğlu

10:30 – 10:50 Break

10:50 – 12:35 **Session 3**: Multicores and Accelerators

12:35 – 2:00 Lunch

2:00 – 3:30 **Panel** "*Everyone doing network and parallelism research should be doing it on mobile.*"

David F. Bacon, Alan Bivens, Jason Mars,

Vijay Janapa Reddi, Craig Stunkel

Moderator: Vijayalakshmi Srinivasan

3:30 – 3:50 Break

3:50 – 5:10 **Session 4**: Parallelism

Saturday, September 19

OUTING

- New York 9-hour Guided Tour:
  - Times Square
    - *(Departure point)*
  - Statue of Liberty
  - Wall Street
  - Grand Central Terminal
  - Federal Hall
  - St Patrick's Cathedral
  - Empire State Building:
    - Tallest building in world for almost 40 years
  - And more
- Lunch in Little Italy.

Times Square
View Map »


Saint Patrick's Cathedral


911 Memorial & Memorial Pools


Wall Street


Grand Central Station


Little Italy


Federal Hall


Trinity Church


Statue of Liberty
Express Ticket


Statue of Liberty Pedestal
(Unique Access)


Wall Street Bull


Ticker-Tape Parade


Empire State Building
Express Ticket