

**Microsoft®**



Microsoft®

# Research Faculty Summit 2012

ADVANCING THE STATE OF THE ART

# The Changing Landscape of Parallel Computing

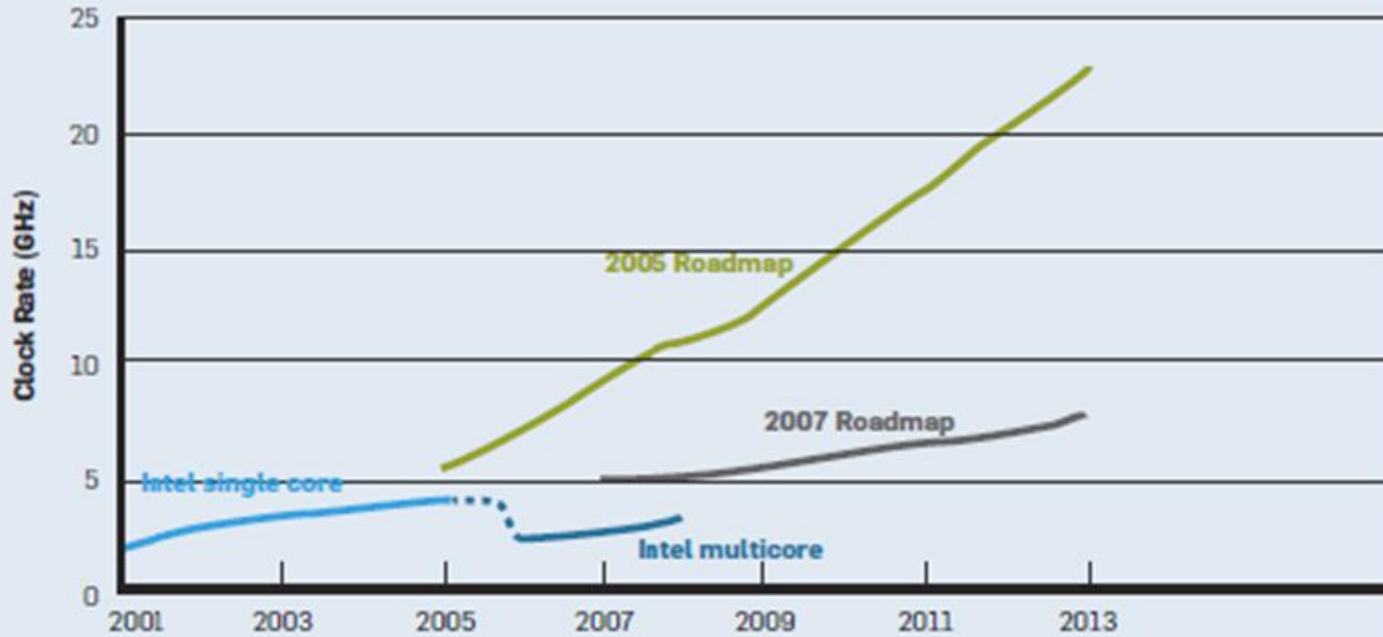
Juan E. Vargas  
Principal Research Manager  
Microsoft

July 17, 2012

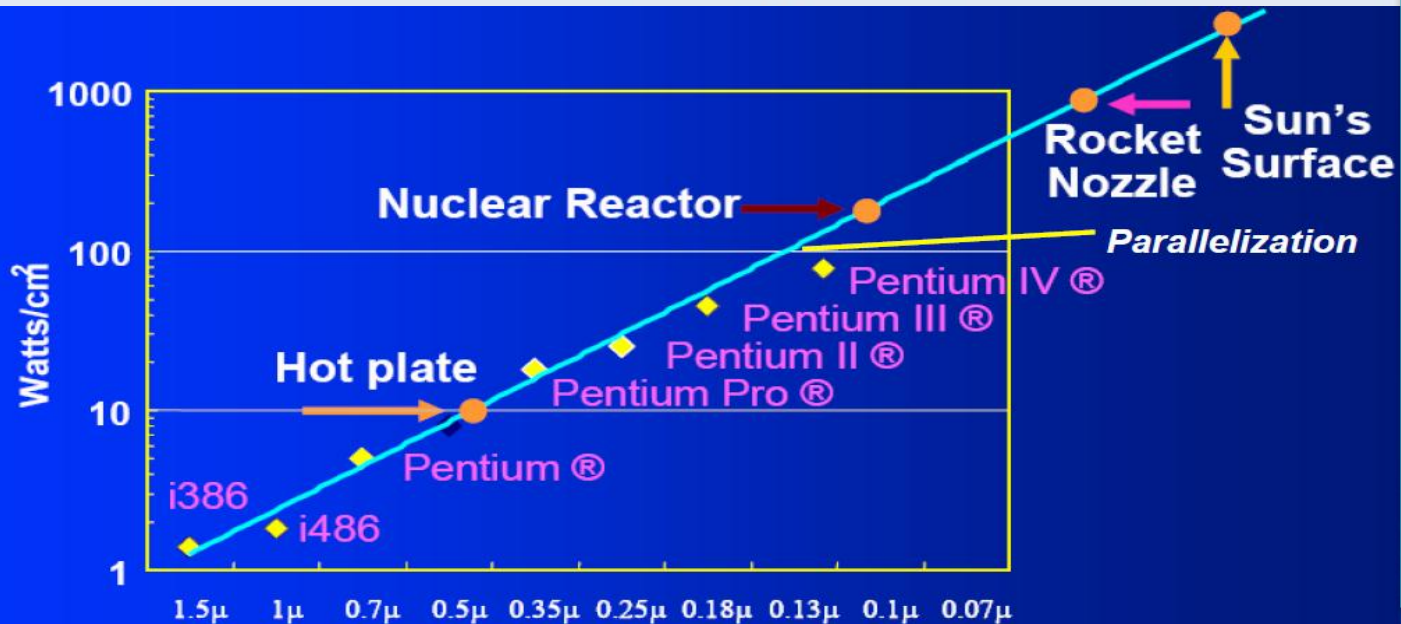


UPCRC = Universal Parallel Computing Research Center(s)

# UPCRC Background and Motivation

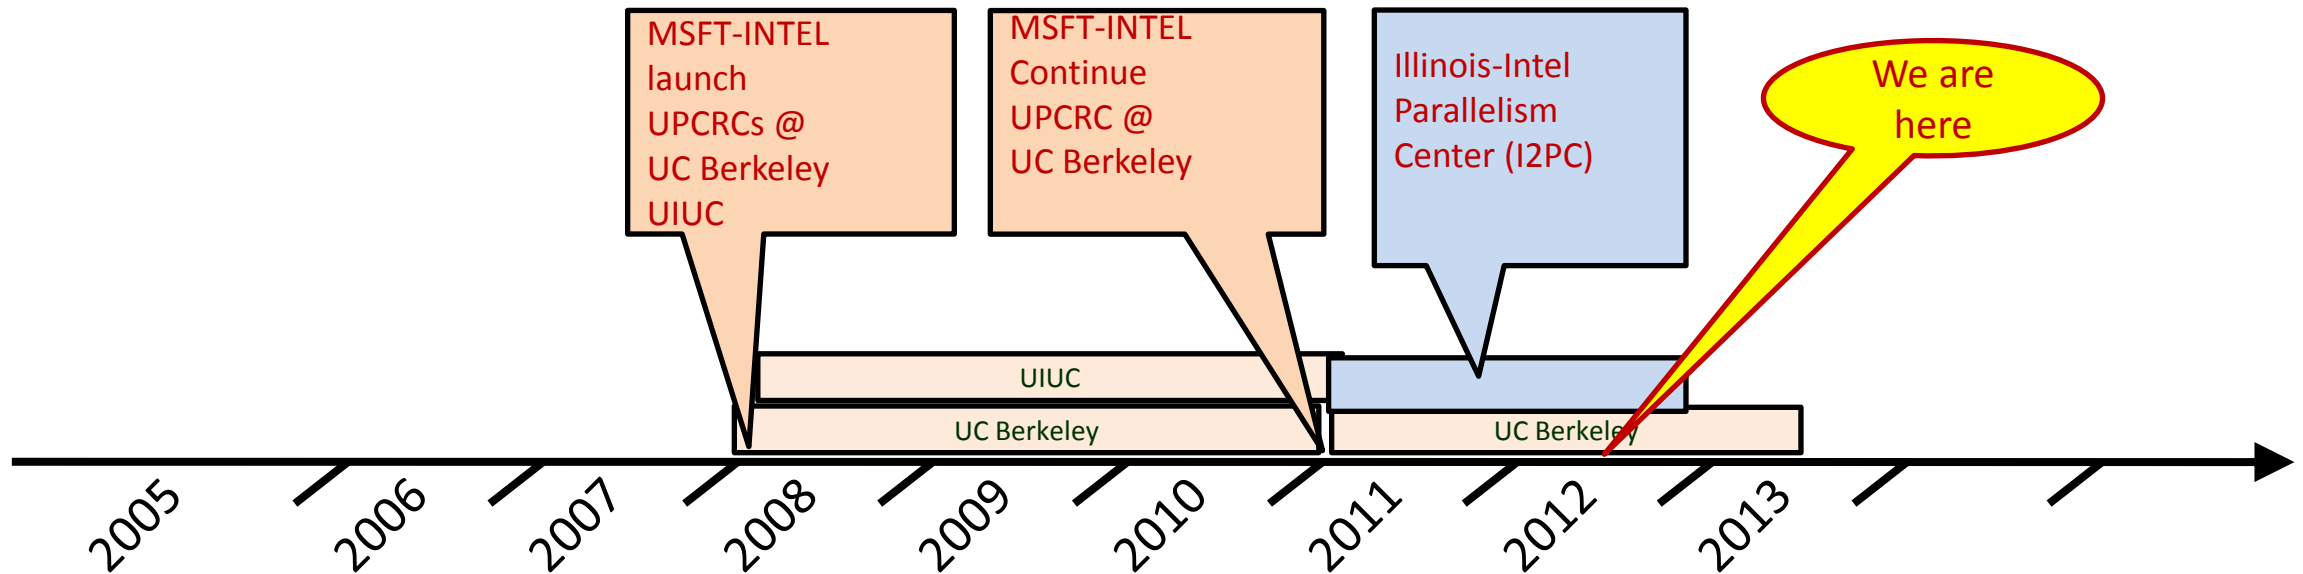


- For decades, computing devices improved performance by increasing CPU clock rate
- Clock rate increases limited by heat and power consumption
- Concurrency and parallelism are the most promising approaches to increase performance
- Microsoft and Intel launched in 2008 the **UPCRC** program to sponsor research into parallelism at the client



# UPCRC Goals and Timeline

- Conduct innovative research into exploiting the computational power of parallelism and concurrency at the client
- Develop new types of programming models, languages, and tools that could be used to build applications that exploit concurrency and parallelism
- Identify architectures that could support the new generations of programming models, languages, tools, and applications
- Understand how the increased numbers of (possibly heterogeneous) processing cores could be combined into a single system
- Understand how future operating systems will manage such systems



# UPCRC Major Accomplishments

**Disruptive Effect and Mindshare:** The program has gained critical mass and it is a leader in academic research on parallelism.

## Academic Metrics


**UIUC:** 18 faculty, 45+ graduate students, 2 MSR interns, 20+ visits from/to MSFT, strong participation at the MSFT Faculty Summits, 42 UPCRC seminars, 60+ journal publications, 100+ conference presentations, outreach activities, including courses and summer schools with 750+ attendees.

**UC Berkeley:** 12 faculty, 120 graduate students, 4 post-docs, 10 MSR interns, 300+ publications in journals & conferences, 22 Books or chapters, 64 Awards, 20+ visits to/from MSFT, strong participation at the MSFT Academic Summits, 20+ lectures at MSFT, outreach activities, 1100+ attendees to summer camps and 1086 attendees to retreats.

# UPCRC/Berkeley

- *Let compelling applications drive research agenda*
- Identify common programming patterns
- Productivity versus efficiency programmers
- Autotuning and software synthesis
- Build-in correctness + power/performance diagnostics
- OS/Architecture support applications, provide flexible primitives not pre-packaged solutions
- FPGA simulation of new parallel architectures
- Co-located integrated collaborative center


*Above all, no preconceived big idea  
- see what works driven by application needs.*



# UPCRC/UIUC

## Make parallel programming easy

- Easy to write correct programs
- Easy to tune for performance, when needed
- Focus on easy parallel programming patterns
- Use sophisticated environment to support a simple programming model atop complex HW and SW
- End goal: **Better applications**



## The Changing Landscape of Parallel Computing I

- 12:30 – 01:00 **Josep Torrellas (UIUC)** : “Research on Parallelism at UIUC: UPCRC and I2PC Illinois”
- 01:00 – 01:30 **Tim Mattson (Intel)**: “Evolution of Parallel Patterns from Design Tool to development tool”
- 01:30 – 02:00 **Dave Patterson (UC Berkeley)** : “UC Berkeley Parallel Computing Laboratory”

## The Changing Landscape of Parallel Computing II

**This session will continue @ 03:00 PM**, with a presentation from Microsoft, followed by a panel about the educational impact of the UPCRC:

- 03:00 – 03:30 **Burton Smith (Microsoft)**: “Operating System Architecture for Parallel and Distributed Computing”
- 03:30 – 04:30 **Panel: Educational Impact of the UPCRC**
  - **David Padua (UIUC)**: “What to cover in an introductory course on parallel programming?”
  - **Jim Demmel (UC Berkeley)**: “Teaching parallel computing nationwide and beyond via NSF/XSEDE”
  - **David Patterson (UC Berkeley)**: “ Teaching Software Engineering to 35,000 students using a Massive Open Online Course (MOOC)”



# UPCRC Workshop, Wednesday, July 18

Microsoft Research (Bldg 99), Redmond Campus

08:30-09:00	<b>Breakfast (Bldg 99 Atrium)</b>
09:00-10:30	<b>Programming Systems</b> David Padua (UIUC): “Abstractions for Parallel Programming” Armando Fox, Shoaib Kamil (UC Berkeley): “SEJITS” David Callahan (Microsoft): “ C++ AMP”
11:00–12:00	<b>Tools for Parallel Testing and Debugging</b> Darko Marinov (UIUC): “UIUC Testing Tools” Danny Dig (UIUC): “Refactoring” Koushik Sen(UCB): “Active Testing and Concurrir” Sunny Chatterjee (Microsoft): “Fighting concurrency bugs with advanced static analysis technology”
12:00-01:00	Lunch
01:00-03:00	<b>Applications</b> John Hart (UIUC): “Avascholar” Minh Do (UIUC ): “3D-Reconstruction” Gerald Friedman (UCB): “PyCASP: Scalable Multimedia Content Analysis on Parallel Platforms Using Python” Leo Meyerovich (UCB): Parallel Components for an Energy-Efficient Web Browser Ras Bodik (UCB): “Program Synthesis for Systems Biology”
03:30-04:30	<b>Architecture</b> Joseph Torrellas (UIUC): “The Bulk Architecture” Krste Asanovic (UCB): “21st Century Computer Architecture Research” Tim Mattson (Intel)
04:30–05:30	<b>Panel : UPCRC : Can Industry &amp; Academia Collaborations be Effective?</b> Dave Patterson (UCB), Burton Smith, Jim Larus (Microsoft), Tim Mattson (Intel), Josep Torrellas (UIUC)

# The Changing Landscape of Parallel Computing I

Josep Torrellas is a professor of CS at UIUC. He is the director of the Center for Programmable Extreme Scale Computing, and the Illinois-Intel Parallelism Center (I2PC). He is a Fellow of IEEE and ACM. He was a Willett Faculty Scholar and the chair of the IEEE Technical Committee on Computer Architecture (2005–2010). He has graduated 30 Ph.D. students, who are now leaders in academia and industry. Torrellas' research is on shared-memory parallel computer architectures, thread-level speculation and speculative synchronization; cache hierarchy organizations for high-performance sequential consistency; embedded-ring snoopy cache-coherence protocols; incremental in-memory check-pointing; popular models of process variation and wearout; and new techniques for software debugging, including deterministic replay of parallel programs. He has received nine best-paper awards. His current research on Bulk Multicore Architecture for parallel programming productivity is funded by I2PC. He also leads The Thrifty-Runnemedes Extreme Scale Architecture for energy and power efficiency, funded by DARPA in collaboration with Intel. Torrellas received a Ph.D. from Stanford University.



Tim Mattson is an old-fashioned parallel programmer specializing in scientific computing—quantum chemistry, exploration geophysics, and bioinformatics. He was part of the team that created the first TFLOP computer, ASCI Red; the OpenMP API for shared memory programming; the OpenCL programming language for heterogeneous platforms; Intel's first TFLOP chip, the 80-core Terascale chip; and Intel's 48-core, Single Chip Cloud Computer research processor. Mattson has published extensively, including the books *Patterns for Parallel Programming*, with Beverly Sanders and Berna Massingill (Addison-Wesley, 2004); *Introduction to Concurrency in Programming Languages*, with Matthew J. Sottile and Craig E Rasmussen (CRC Press, 2009); and *OpenCL Programming Guide*, with Aaftab Munshi, Ben Gaster, James Fung, and Dan Ginsburg (Addison-Wesley, 2011).



David Patterson joined the University of California, Berkeley in 1977. He is director of the Parallel Computing Laboratory. In the past, he served as director of the Reliable Adaptive Distributed Systems Laboratory, as chair of Berkeley's Computer Science Division, chair of the Computing Research Association, and president of the ACM. His most successful projects have been Reduced Instruction Set Computers, Redundant Arrays of Inexpensive Disks, and Network of Workstations. All three research projects helped lead to multibillion-dollar industries. This research led to many papers and six books, with the most recent being *Engineering Long-Lasting Software: An Agile Approach Using SaaS and Cloud Computing*, co-authored with Armando Fox.



# The Changing Landscape of Parallel Computing II

Burton J. Smith, a Microsoft technical fellow, works with various groups within the company to help address challenges associated with the emergence of many-core systems and the increasing importance of distributed services. Before joining Microsoft in 2005, he co-founded Cray Inc., formerly Tera Computer, where he variously served as its chief scientist, a member of the board of directors, and its chairman until 1999. Before that, Smith spent six years with Denelcor, Inc. and three years at the Institute for Defense Analyses Supercomputing Research Center. In 2003, Smith received the Seymour Cray Award from the IEEE Computer Society and was elected to the National Academy of Engineering. He received the Eckert-Mauchly Award in 1991, given jointly by IEEE and ACM, and was elected a fellow of each organization in 1994. He was elected Fellow of the American Academy of Arts and Sciences in 2010. Smith attended the University of New Mexico, where he earned a B.S.E.E. degree, and the MIT, where he earned S.M., E.E., and Sc.D. degrees.



Jim Demmel received his B.S. in Mathematics from the California Institute of Technology in 1975 and his Ph.D. in Computer Science from the University of California, Berkeley in 1983. After spending six years on the faculty of the Courant Institute at New York University, he joined the Computer Science Division and Mathematics Department at Berkeley in 1990, where he holds a joint appointment.



David Padua is the Donald Biggar Willet Professor of Engineering at UIUC, where he has been since 1985. He has been associate director of the Center for Supercomputing Research and Development, a member of the Center for Simulation of Advanced Rockets' Science Steering Committee, and chair of the College of Engineering Faculty Advisory Committee. He has served as a program-committee member, program chair, or general chair for more than 60 conferences and workshops. He served on the editorial board of the *IEEE Transactions of Parallel and Distributed Systems* and the *ACM Transactions on Programming Languages and Systems*. He was editor-in-chief of the *International Journal of Parallel Programming (IJPP)* and steering-committee chair of ACM SIGPLAN's Principles and Practice of Parallel Programming. He is member of the editorial boards of the *Journal of Parallel and Distributed Computing* and *IJPP*, and editor-in-chief of the *Encyclopedia of Parallel Computing* (Springer-Verlag, 2012). His areas of interest include compilers, programming tools, machine organization, and parallel computing. He is a fellow of the IEEE and the ACM.



David Patterson joined the University of California, Berkeley in 1977. He is director of the Parallel Computing Laboratory. In the past, he served as director of the Reliable Adaptive Distributed Systems Laboratory, as chair of Berkeley's Computer Science Division, chair of the Computing Research Association, and president of the ACM. His most successful projects have been Reduced Instruction Set Computers, Redundant Arrays of Inexpensive Disks, and Network of Workstations. All three research projects helped lead to multibillion-dollar industries. This research led to many papers and six books, with the most recent being *Engineering Long-Lasting Software: An Agile Approach Using SaaS and Cloud Computing*, co-authored with Armando Fox.



# *Microsoft*