Applications Session Summary
How are Applications Evolving?

Workshop on Programming Languages for HPC (HPC WPL)

Jeffrey Vetter

Architectures are Diversifying and their Lifetimes are Shrinking

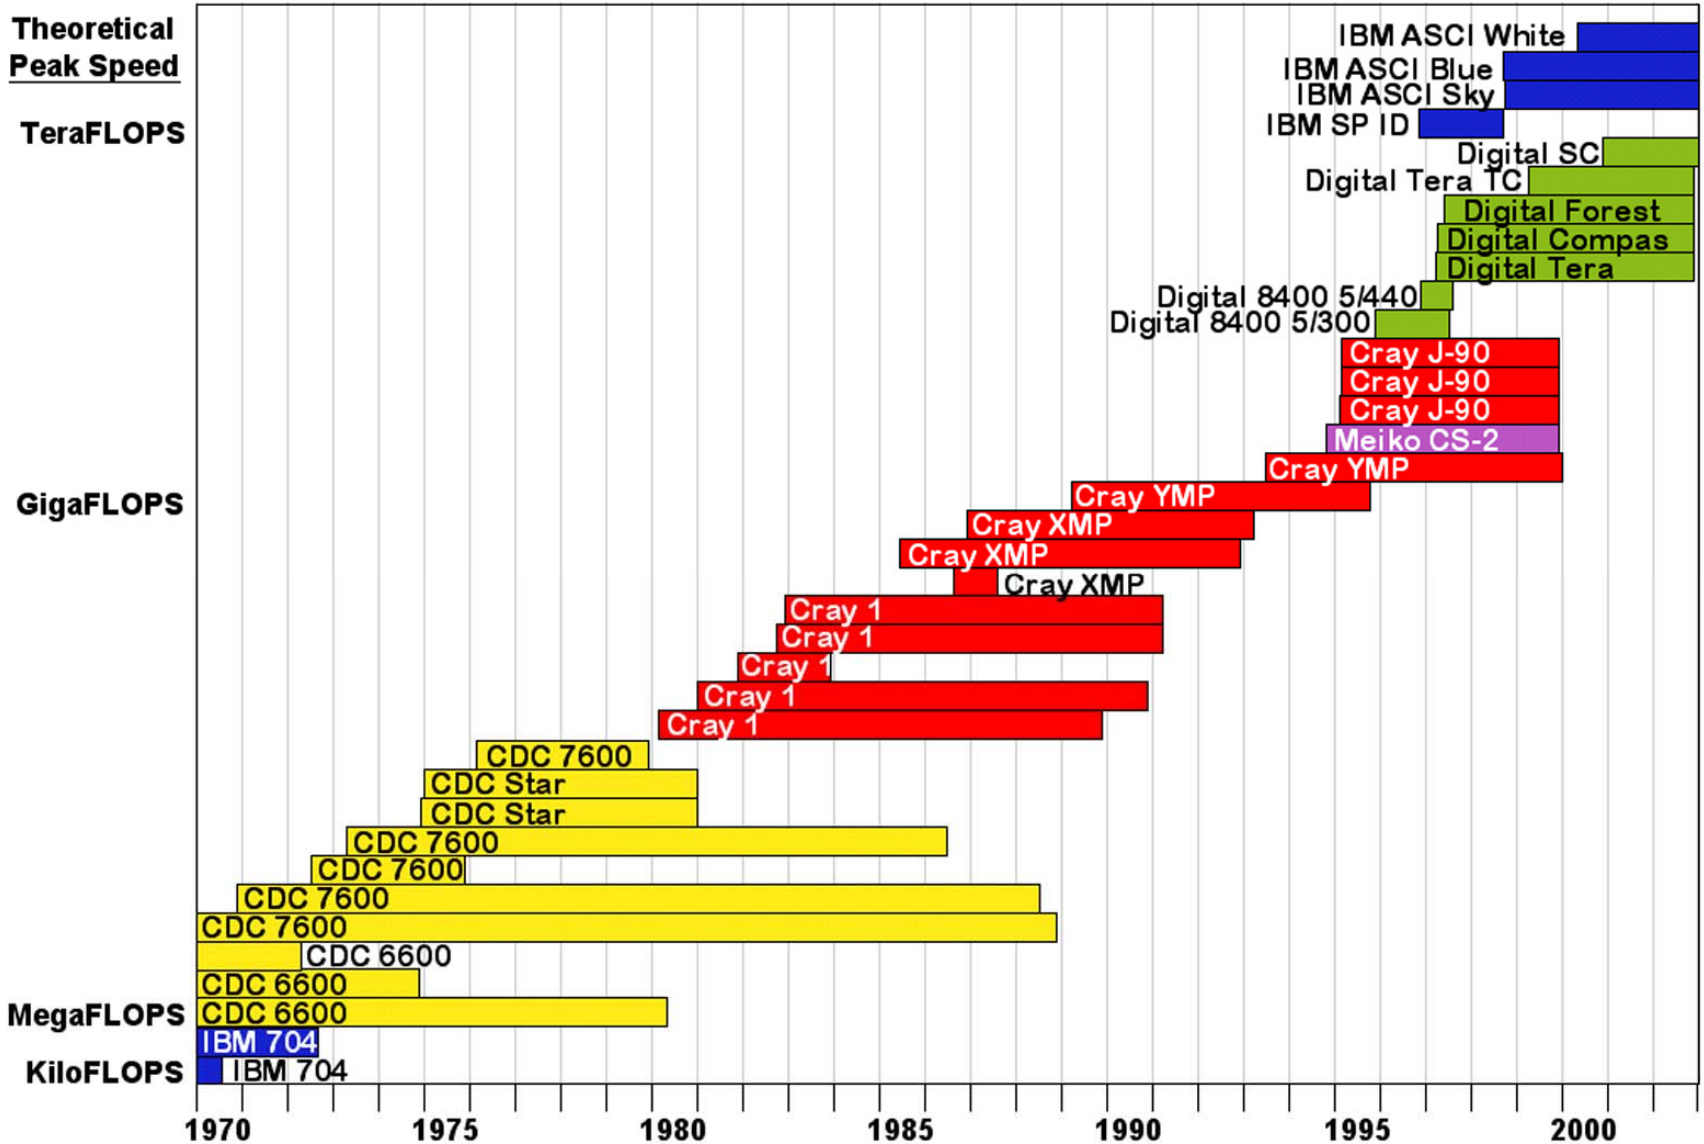

???

Source: LLNL

Commodity Platforms Set the Baseline

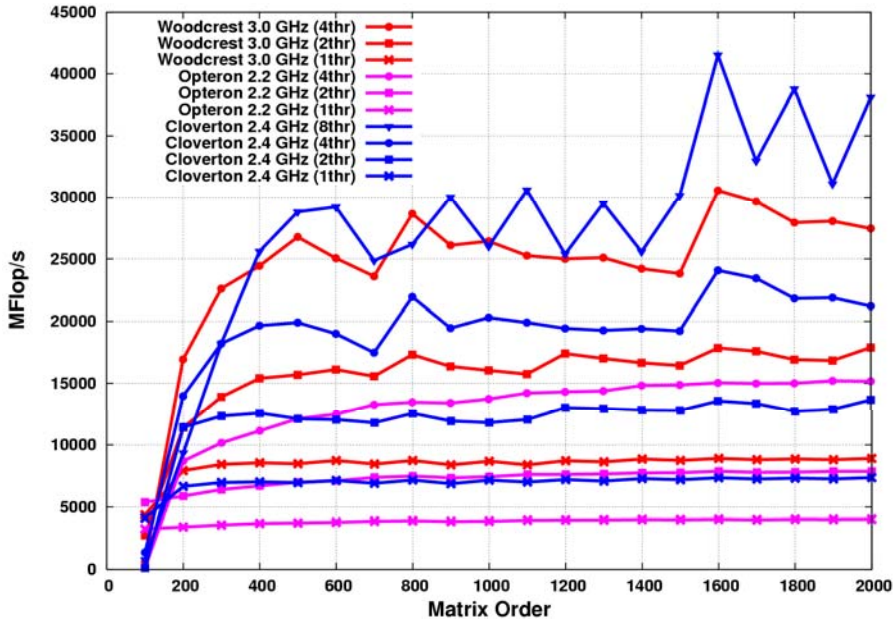

DGEMM

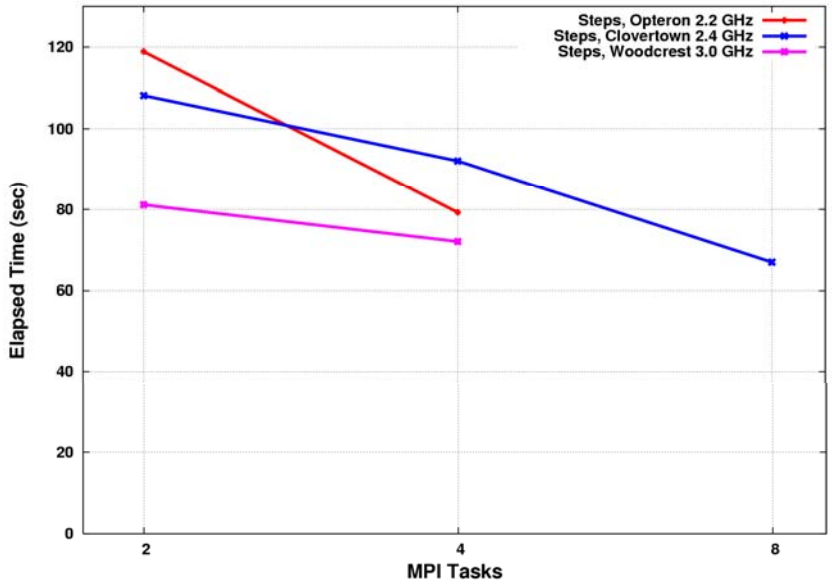

POP

Architectures that Match Application Requirements can offer Impressive/Astounding Performance Benefits

- Numerous FPGA results on integer, logic, flop applications
- Numerous results on GPU
  - Geo-registration MD
  - GPGPU …
- HPCC RandomAccess on Cray X1E

Molecular Dynamics	
System	**Seconds**
Cell PPE	0.425
MTA2 w/32 procs	~0.035
2.2GHz Opteron	0.125
Cell w/ 8 SPEs	0.013
GPU (7900GT)	0.012

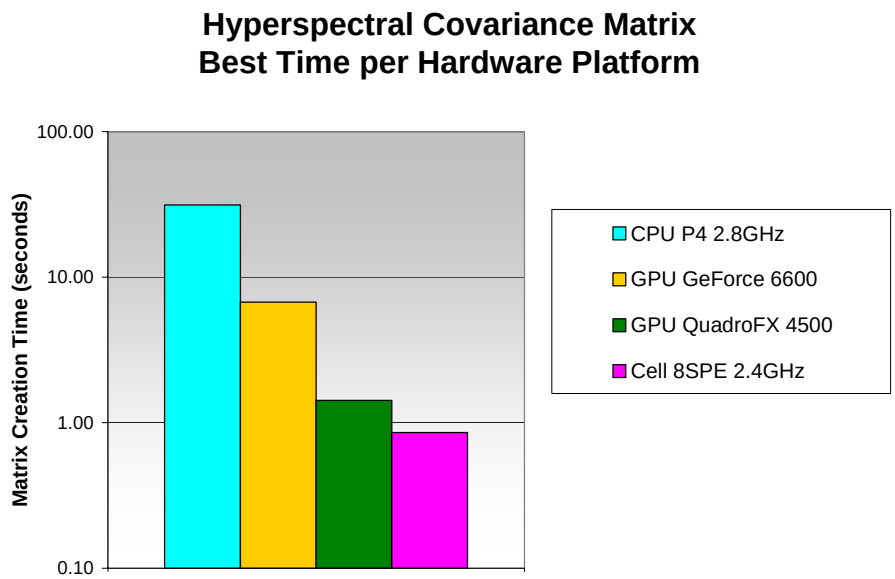

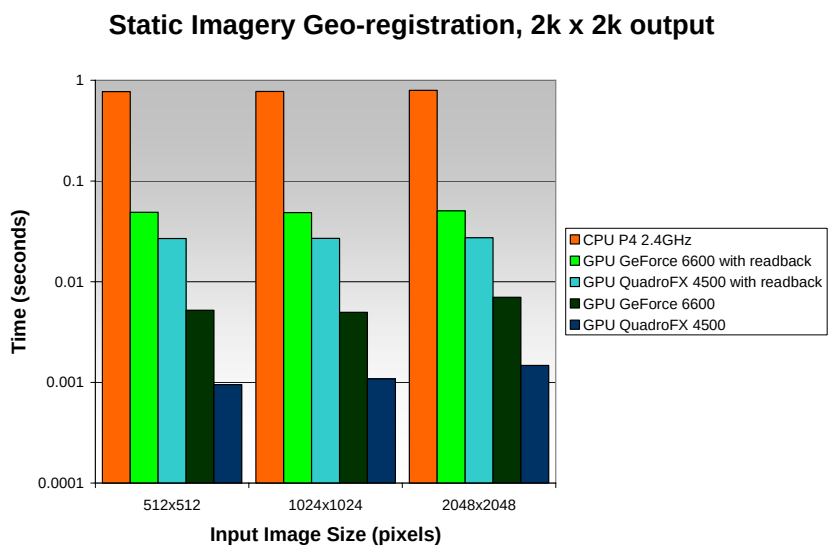

Programming Models/Languages

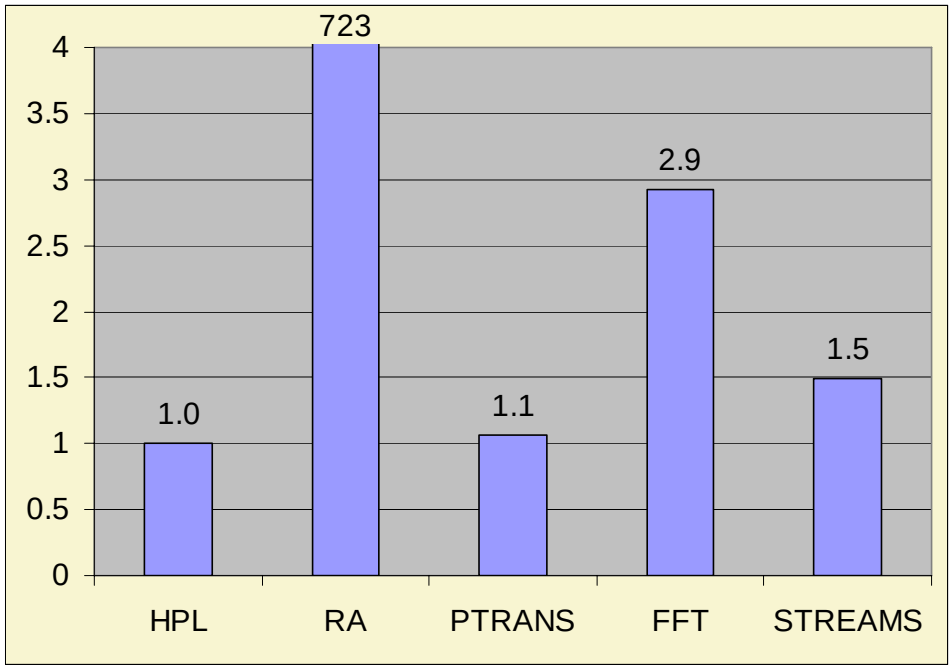

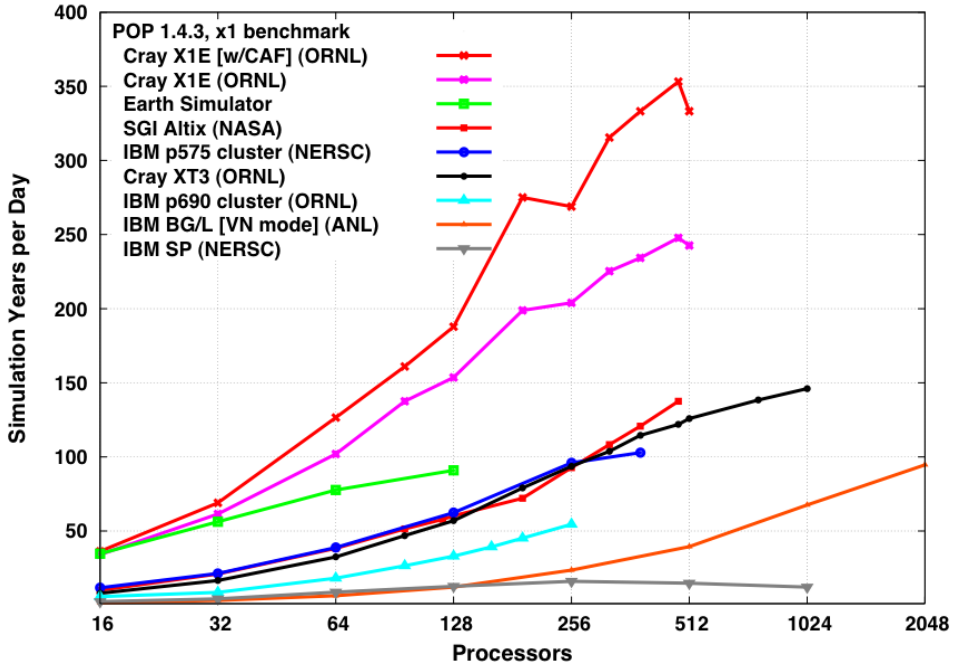

Productivity

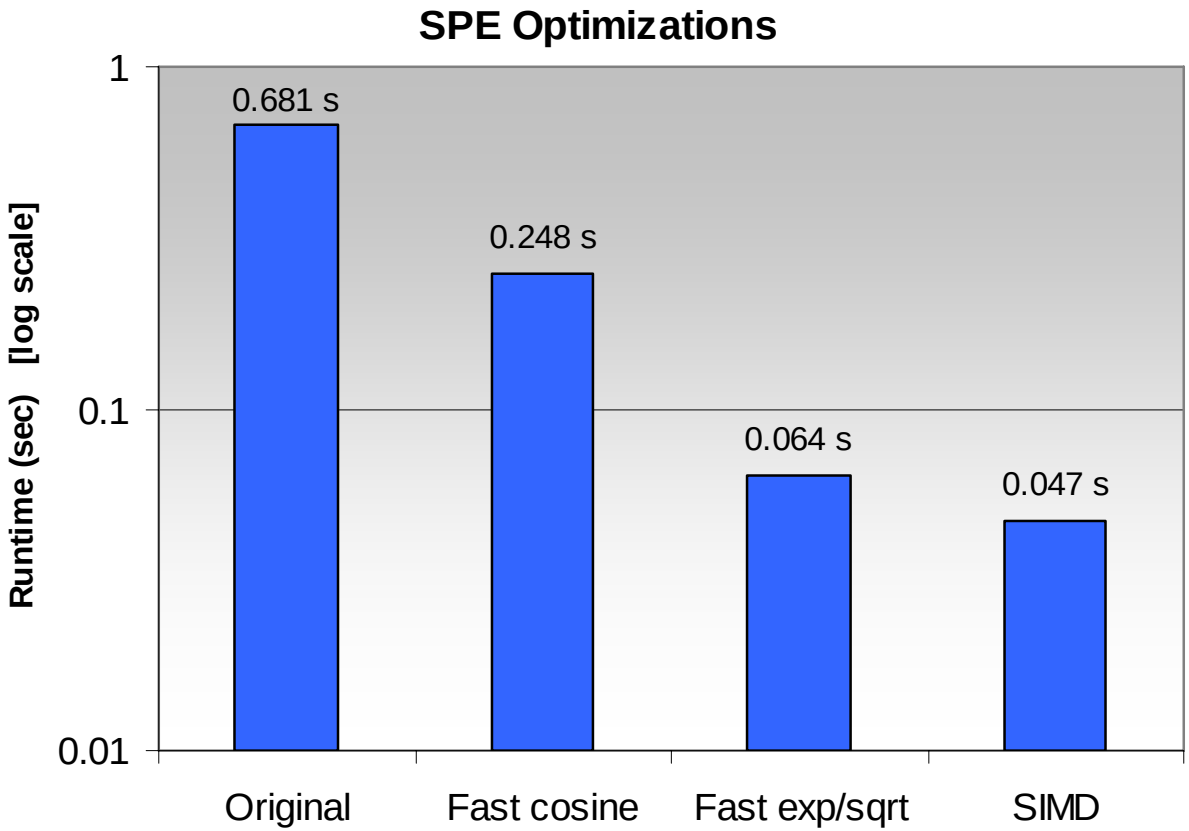

Ackley's function from Genetic Algorithm benchmark

Challenges for New Architectures, Languages

- Performance prediction
  - Should I purchase this new architecture, language?
  - What will be the performance improvement on my application workload?
  - Is it working as we expect?
  - How can I optimize my application on this target?
- Productive software systems
  - Do I have to rewrite my application for each architecture?
  - How stable is the functionality across systems?
  - How stable is the performance across systems?

Notes

- Lucas
  - Apps
    - New apps, remember old
  - Execution time productivity
    - Benchmarks on real systems
    - Easier than development time
  - Productivity
  - Languages
    - FORTRAN forever ☺
  - Heterogeneous computing
  - Data driven simulation: sensors, databases

Notes (2)

- Hendrickson
  - Unstructured, adaptive
  - static v. dynamic
  - Fine-grained parallelism and synchronization
  - Nested parallelism
  - Latency dominated operations
  - Global address spaces
  - Poor locality
  - Correct architecture matters
    - 4 MTA-2 processors = 32k BGL processors

Notes (3)

- Heroux / Trilinos
  - Not much in software has changed in 10 years
  - Applications are built to abstract interfaces to hide differences across platforms
  - Libraries/interfaces are more important than language
    - Can hide accelerators under the abstraction?
  - Same conclusion as 1995
    - Not much different than SMP

Notes (4)

- Perumalla / PDES
  - Largely non-fp computation
  - Irregular and adaptive computation
  - Fine-grained parallelism and synchronization
  - Small computational intensity
  - Scalability
  - Speculation

Summary

- Surprisingly, the diversity of architectures over the past 15 years has forced a convergence of software → standards, libraries, APIs
- Entirely new architectures may force a reexamination of software stacks and applications
  - Fine grained parallelism
  - Global address spaces
  - Irregular and adaptive computation
- How will we know
  - Need a better mapping between applications and architectures/languages
  - We need tools to estimate performance/benefit beforehand
- When will the benefit outweigh the switching costs?
  - Multicore
  - Extreme scaling
  - Entirely new, heterogeneous architectures

Acknowledgements and More Info

- This research was sponsored by the Office of Mathematical, Information, and Computational Sciences, Office of Science, U.S. Department of Energy under Contract No. DE-AC05-00OR22725 with UT-Batelle, LLC. Accordingly, the U.S. Government retains a non-exclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

- http://ft.ornl.gov
- vetter@ornl.gov