

Trends in HPC and HPEC Convergence

Richard Games
rg@mitre.org

Thesis

- Over the last 10 years R&D investments have made *high performance embedded computing* for national security applications more like mainstream *high performance computing*

- Over the next 10 years R&D investments will make mainstream *high performance computing* for national security applications more like *high performance embedded computing*

Thesis

- Over the last 10 years R&D investments have made *high performance embedded computing* for national security applications more like mainstream *high performance computing*
 - DARPA Touchstone
 - DARPA Embedded Systems
 - OSD High Performance Embedded Computing Software Initiative
- Over the next 10 years R&D investments will make mainstream *high performance computing* for national security applications more like *high performance embedded computing*
 - DARPA Adaptive Computing
 - DARPA Data Intensive Systems
 - DARPA Polymorphic Computing Architectures
 - DARPA High Productivity Computing Systems

Thesis

- Over the last 10 years R&D investments have made *high performance embedded computing* for national security applications more like mainstream *high performance computing*
 - DARPA Touchstone
 - DARPA Embeddable Systems
 - **OSD High Performance Embedded Computing Software Initiative**
- Over the next 10 years R&D investments will make mainstream *high performance computing* for national security applications more like *high performance embedded computing*
 - DARPA Adaptive Computing
 - DARPA Data Intensive Systems
 - DARPA Polymorphic Computing Architectures
 - **DARPA High Productivity Computing Systems**



HPEC Software Initiative

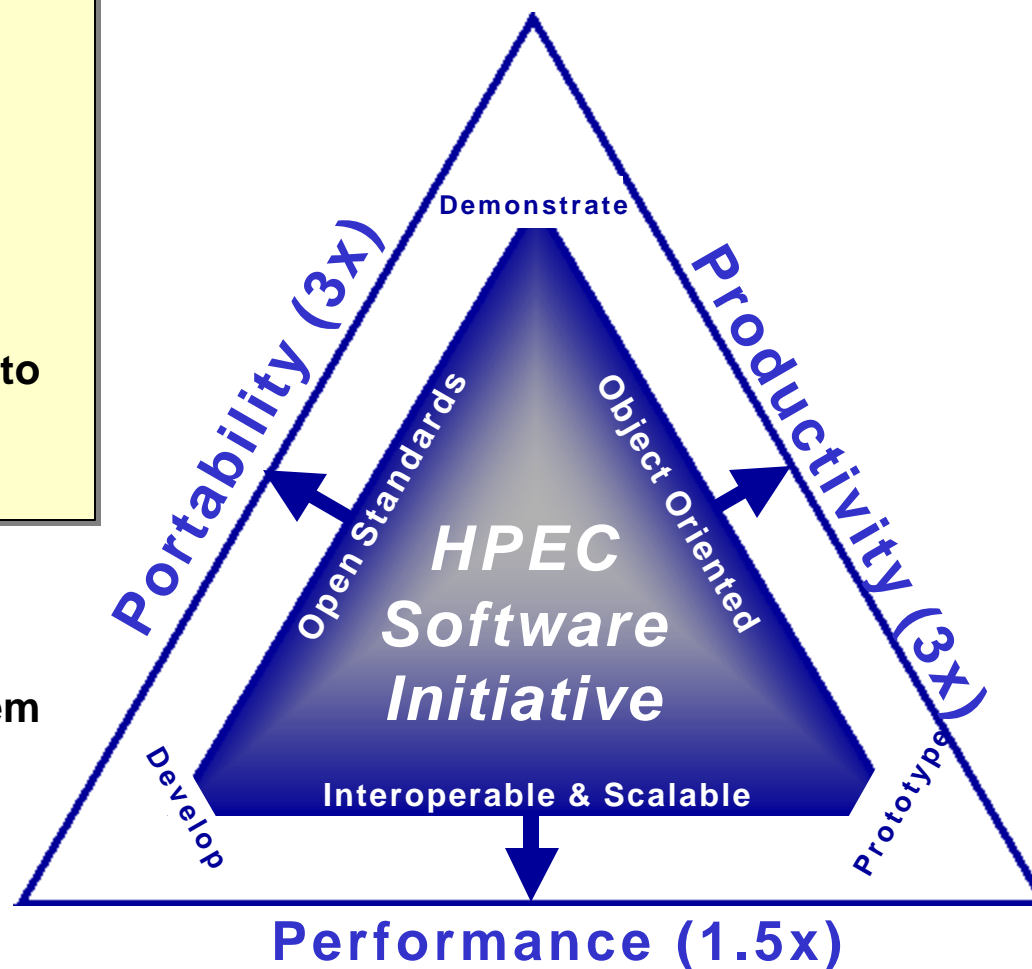
Program Goals

- Develop and integrate software technologies for embedded parallel systems to address portability, productivity, and performance
- Engage acquisition community to promote technology insertion
- Deliver quantifiable benefits

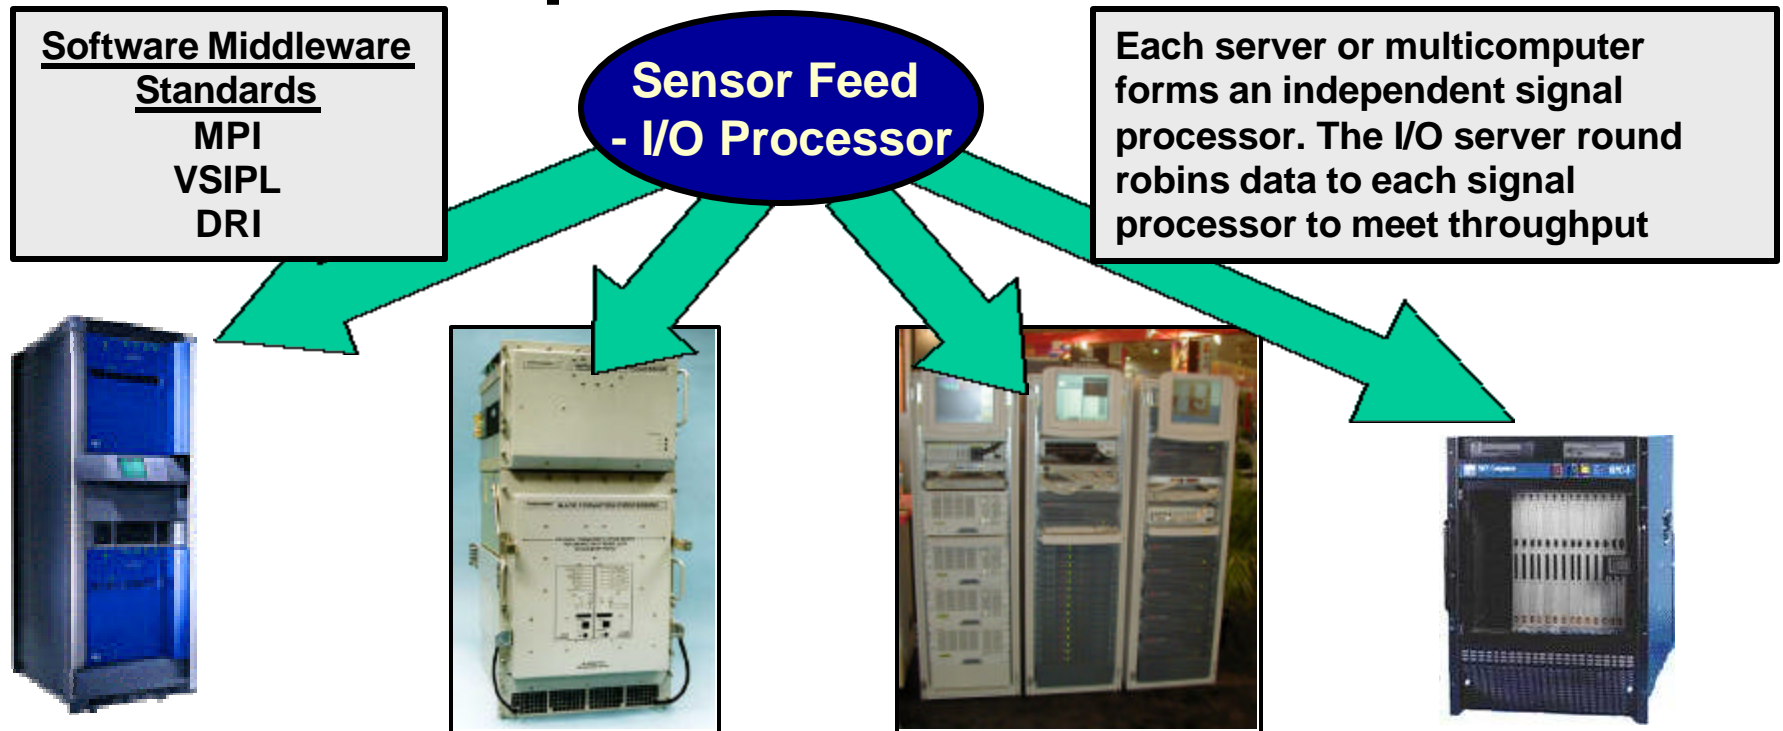
Portability: reduction in lines-of-code to change port/scale to new system

Productivity: reduction in overall lines-of-code

Performance: computation and communication benchmarks



Common Imagery Processor Experiment Overview



Shared Memory Servers

Distributed Memory Multicomputers

Distributed memory CIP software would allow insertion of embedded multicomputer or commodity clusters as a signal processor in CIP system

HPEC-SI Middleware

Development

VSIPL++

- MAPPING (data parallelism)
- Early binding (computations)
- Compatibility (backward/forward)
- Local Knowledge (accessing local data)
- Extensibility (adding new functions)
- Remote Procedure Calls (CORBA)
- C++ Compiler Support
- Test Suite
- Adoption Incentives (vendor, integrator)


Applied Research

Parallel VSIPL++

- MAPPING (task/pipeline parallel)
- Reconfiguration (for fault tolerance)
- Threads
- Reliability/Availability
- Data Permutation (DRI functionality)
- Tools (profiles, timers, ...)
- Quality of Service



From the Small to the Big




Earth Simulator

Earth Simulator pictures from www.es.jamstec.go.jp/esc/eng
JSTARS pictures courtesy of Northrop Grumman

Efficiency: The Big Picture

Losses accumulate from the point of electricity generation, through distribution, and *finally during utilization by the end user*

Efficiency of Electrical Power Generation



One Ton of Coal Generated

Year	Energy
1891	150 kWh
1914	550 kWh
1920	630 kWh
1939	1566 kWh
2002	3000 kWh

Electricity distribution efficiency: 92%

ASCI Q: 24 - 30 Tflop/s (peak)
 3 megawatts to run plus
 2 megawatts to cool
 (energy for 5000 homes)

Computing efficiency

Gflop/s

Percent peak

Gflop/s/Watt

...

MITRE

www.electricity.org.uk/uk_inde/viron/env_19.html
www.parcon.uci.edu/paper/energy.htm



High Productivity Computing Systems




Goals:





- Provide a new generation of economically viable high productivity computing systems for the national security and industrial user community (2007 – 2010)

Impact:

- **Performance** (efficiency): critical national security applications by a factor of 10X to 40X
- **Productivity** (time-to-solution)
- **Portability** (transparency): insulate research and operational application software from system
- **Robustness** (reliability): apply all known techniques to **protect against outside attacks**, hardware faults, & programming errors



HPCS Program Focus Areas



Applications:

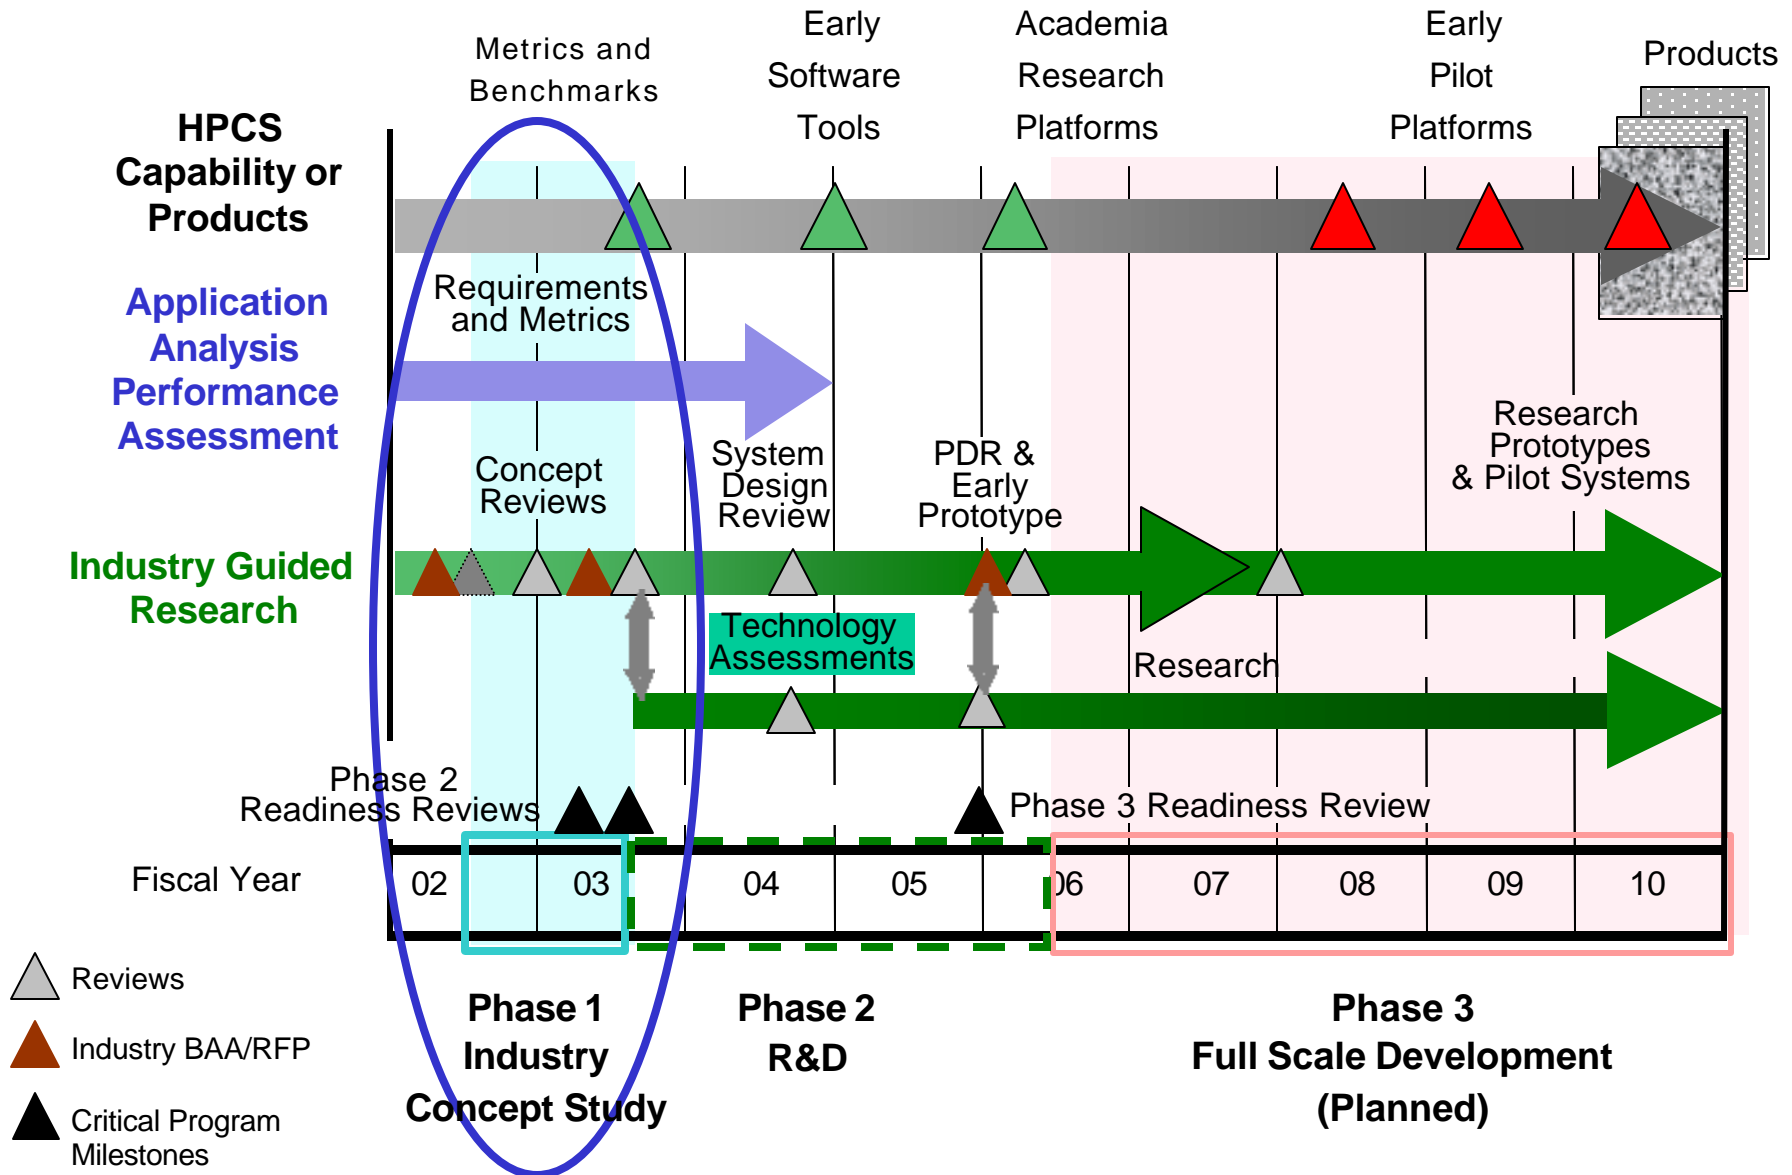

- Intelligence/surveillance, reconnaissance, cryptanalysis, weapons analysis, airborne contaminant modeling and biotechnology

Fill the Critical Technology and Capability Gap


Today (late 80's HPC technology).....to.....Future (Quantum/Bio Computing)



HPCS Program Phases 1-3




Bounding the HPCS Challenges



Why applications with limited memory reuse perform inefficiently today

STREAMS ADD: Computes $A + B$ for long vectors A and B


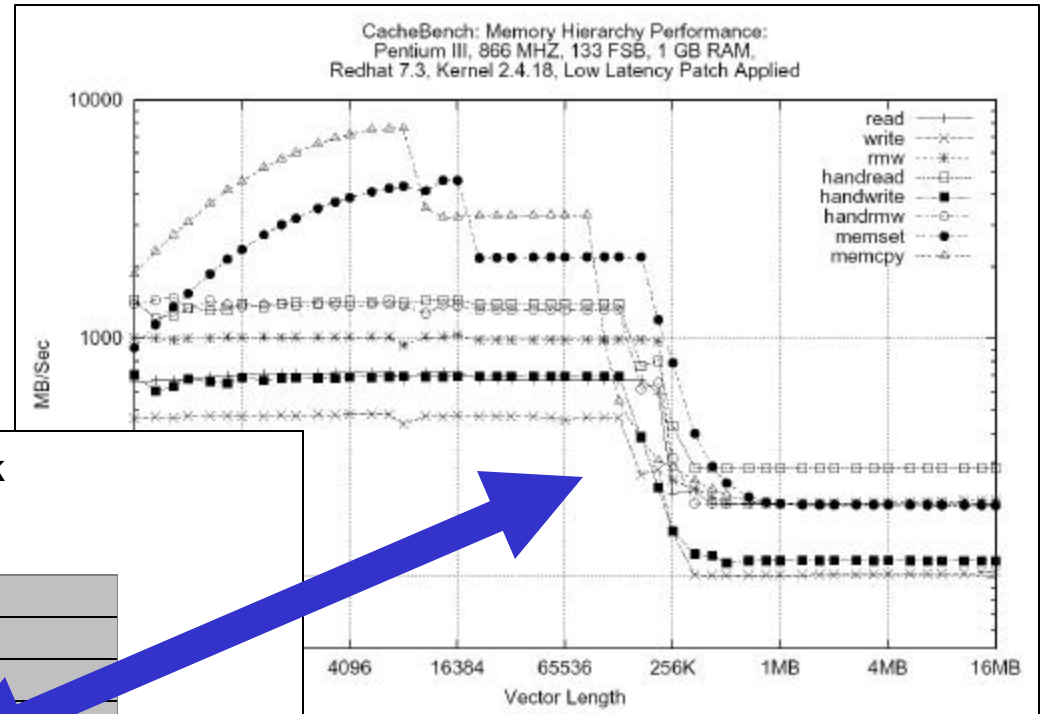


- New microprocessor generations “reset” performance to around 6% of peak
- Performance degrades to 1% - 3% of peak as clock speed increases

Long FFTs are Inefficient

CacheBench

Direct correlation between memory bandwidth from various levels of the memory hierarchy and the performance of real applications




For the CacheBench benchmark see icl.cs.utk.edu/projects/llcbench/cachebench.html (Phil Mucci)
For FFTW software see www.fftw.org (Matteo Frigo & Steven G. Johnson)

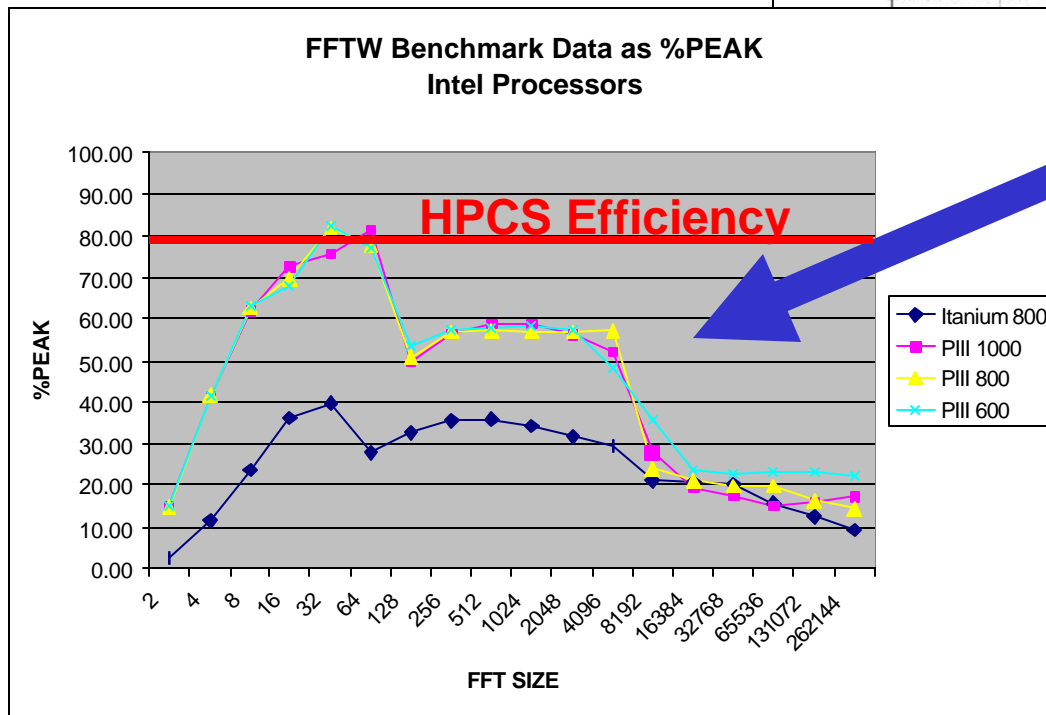
Long FFTs are Inefficient

CacheBench

Direct correlation between memory bandwidth from various levels of the memory hierarchy and the performance of real applications



FFTW Benchmark Data as %PEAK
Intel Processors



HPCS Challenge
Cost effective signal processing in software for high throughput streaming applications

For the CacheBench benchmark see icl.cs.utk.edu/projects/llcbench/cachebench.html (Phil Mucci)
For FFTW software see www.fftw.org (Matteo Frigo & Steven G. Johnson)

Value Proposition: Metrics

Producer

- Sells computers
- Sells support
- Profit
- Market share
- Stockholder's equity
- Reputation
- Peak rates
- Customer satisfaction
- Deliver solutions
- Novel technology
- ...

Consumer

- Has national security mission
- Needs a computer to process data or calculate answers
- In time--time to solution
- Fits (size, weight, power)
- Easy to program--idea to sol'n
- Affordable--life-cycle, facilities and support costs
- Efficient--sustained rates
- Reliable
- Evolvable
- ...

Thesis

- Over the last 10 years R&D investments have made *high performance embedded computing* for national security applications more like mainstream *high performance computing*
 - Looking for good HPEC-SI demonstrations
- Over the next 10 years R&D investments will make mainstream *high performance computing* for national security applications more like *high performance embedded computing*
 - Looking for good HPCS challenge problems

Richard Games, rg@mitre.org