PERCS: IBM Effort in HPCS

Mootaz Elnozahy
IBM Austin Research Lab

Overview

- Team
- Design constraints
- Vision
- Technical overview
- Conclusion

The Team

- IBM
  - Austin Research
  - Watson Research
  - Server Group
  - Software Group
  - Microelectronics
- Partners
  - UIUC
  - UT
  - MIT
  - Pittsburgh
  - LANL
  - RPI
  - New Mexico
  - Cornell
  - UC Berkeley
  - Purdue
  - Wisconsin
  - U of Del
  - Vanderbilt
  - Dartmouth

Design Constraints

- Legacy investments
- Looming technology crisis
- HPC customer diversity
- Business model
  - Must do well both on commercial and scientific workloads
- Cost issues
  - Threat of commoditization
- Productivity as a main theme

IBM's Vision

A dynamic system that adapts to application needs

The strategy

- **Aggressive productivity targets**
- **Commercial viability**
- **Link into product cycle toward end of phase 2**

Innovation with Commercial Viability

- **Adaptability is key**
  - Architecture gets closer to application needs, yielding better performance and broader application range
- **Backward compatibility with PowerPC**
  - Leverages existing infrastructure, training and investment
  - Exploits proven ability of the current architecture to perform well for *many* apps
- **Leverage open source**
- **Modular design packaged in different configurations**

Suggested PERCS Roadmap

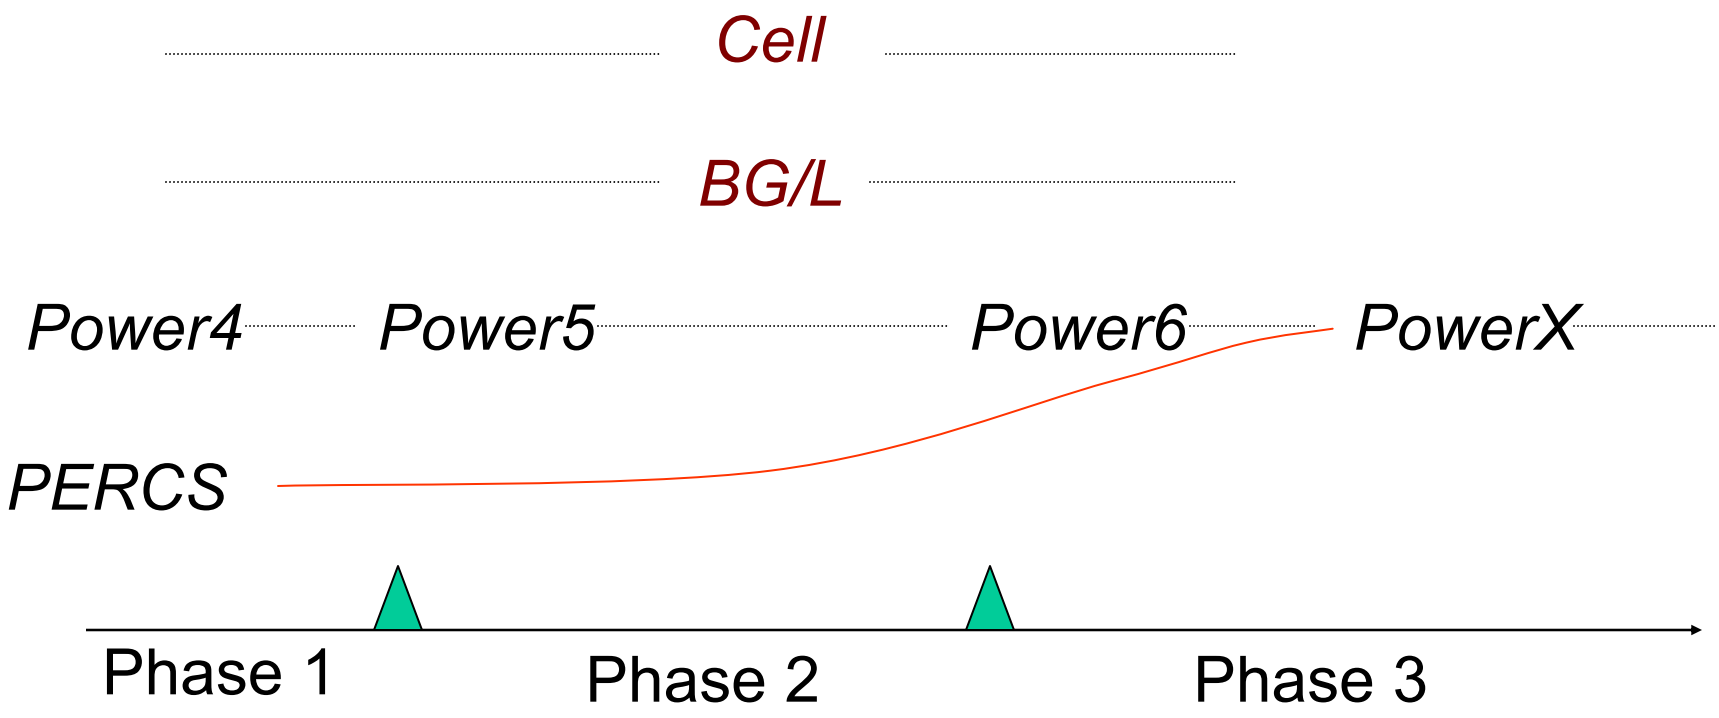

Technical Overview

Scope

- **Application focus**
  - **Commercial**
  - **Security**
  - **HPC**
  - **Bioinformatics**
  - **Data streaming**
  - **New 2010-apps ??**
- **Integrated solution**

HPCS PERCS

IBM

Productivity Metrics

- **A theory for productivity that**
  - Reflects the importance of time-to-solution
  - Incorporates
    - Software development
    - Maintenance costs
    - Hardware costs
    - Tradeoffs among the three
  - Uses $ as a common denominator
- **Measurable system metrics:**
  - Characterize the productivity of programming environments and execution platforms
    - Experiments with programmers
  - Weighted according to application set and customer goals

Architecture Innovations

- **Adaptation**
  - Vector/stream processor morphs
  - Memory-in-processor morph (PIM-like programming model with practical hardware)
  - SMT and conventional caches for commercial apps
- **Proactive memory architecture**
  - Embedding intelligence across memory hierarchy for better performance and lower latency
- **Leverage IBM's technology advantages**
  - Aggressive hardware design
  - New revolutionary packaging and device technologies subject to practicality & cost

IBM

Programming Model Work

- **Support for newer programming languages**
  - UPC, StreamIt, and domain-specific language
- **Morphogenic software process**
  - Bridging the gap between domain experts and programmers
- **New and revised abstractions**
  - Enclaves, atomic actions and asynchronous calls
- **Aggressive compiler support**
- **Integrated development environments, visual tools & component-based software**
  - Integrate best practices from commercial into HPC

Infrastructure Work

- **K-42 operating system**
  - Design for scalability from scratch
- **High-level automated verification**
  - Hardware and software
- **Design for low-power, high-performance circuits**
- **Continuous performance monitoring and automatic tuning**
- **Robustness:**
  - New programmer-transparent efficient checkpointing
  - Self-healing autonomic middleware

Getting It Right! PERCSim

- **Feedback to hardware & software designers**
  - Test-bed for quick evaluation of “what if” questions across all system levels
- **Pre-hardware software development**
  - Tuning and evaluation well before design freeze
- **Speeds up verification and enhance methodology**
- **Execution-driven**
  - Better represents scalable multiprocessors compared to traditional tracing (e.g. synchronization traffic)
- **Power management support**

Example: SPEC SDET

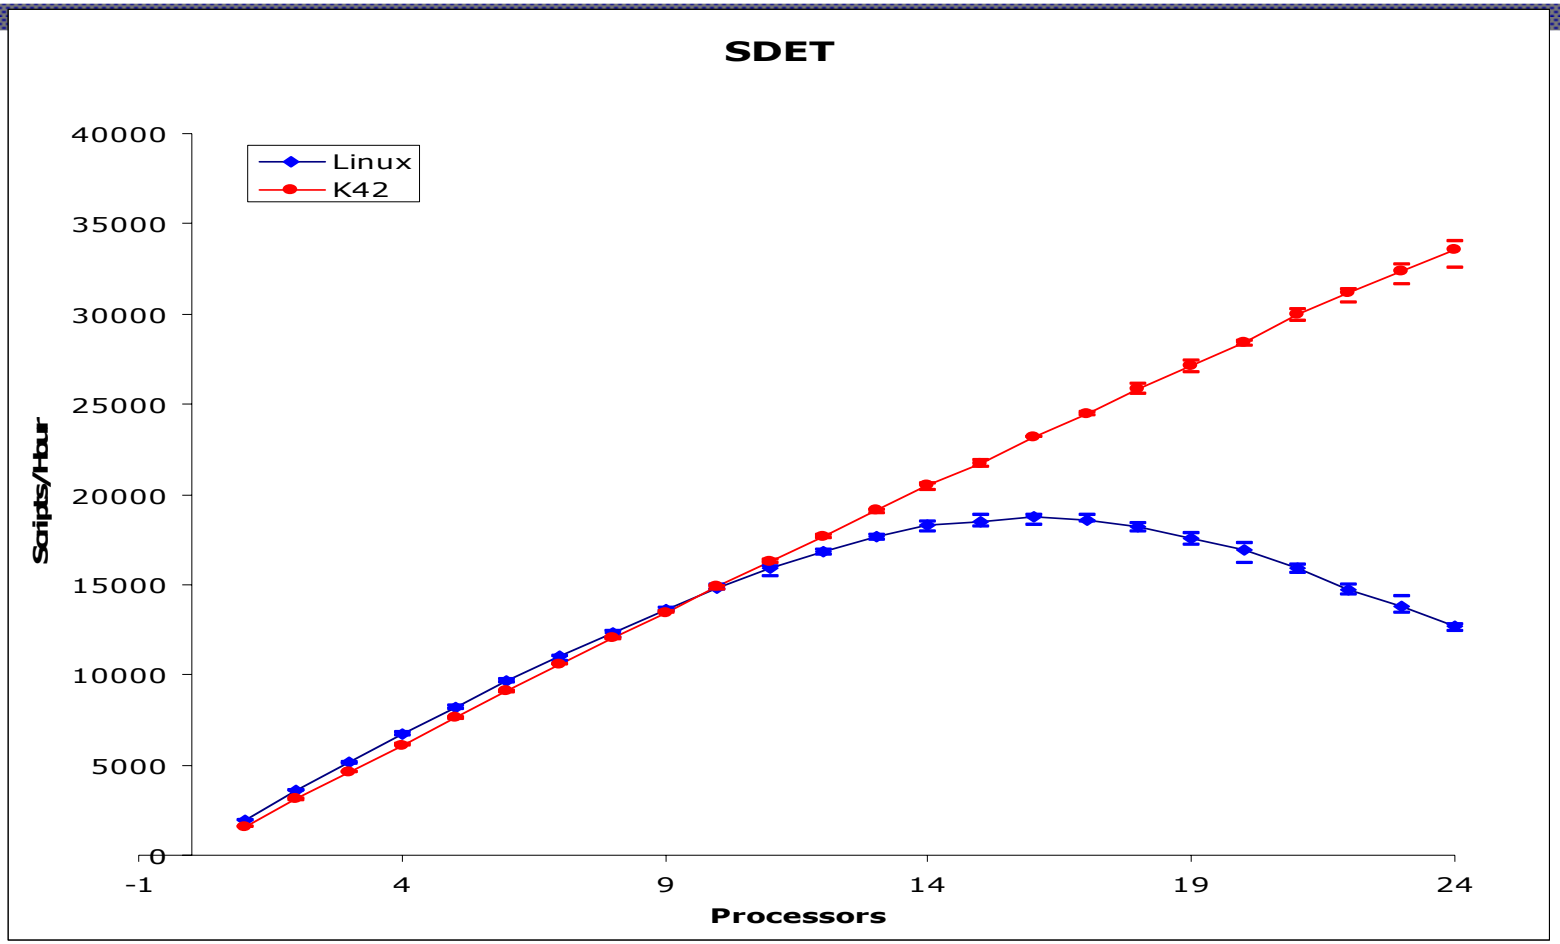

Summary

- **An ambitious vision for adapting systems to applications**
  - Solve productivity problems of HPC community
  - Explore technologies otherwise deemed too risky
  - Economic viability
- **Breadth and depth of IBM's R&D behind the effort**
  - Record of innovation with *reliability & delivery*
- **HPCS will have a strong impact on IBM and universities**
  - We hope to change status quo