

Panel Session:
Will Software Save Moore's Law?

Vivek Sarkar

(vsarkar@us.ibm.com)

Senior Manager, Programming Technologies

IBM T.J. Watson Research Center

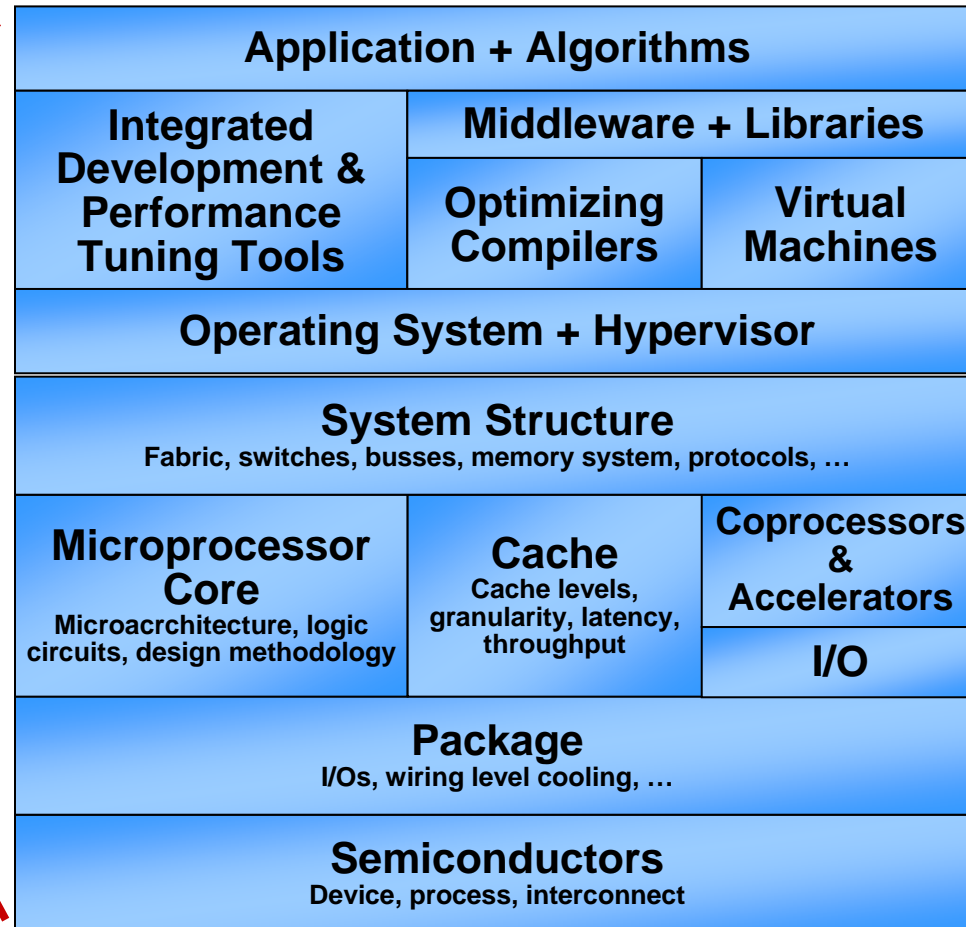


This work has been supported in part by the
Defense Advanced Research Projects Agency (DARPA)
under contract No. NBCH30390004.



System-level Performance of Embedded Applications

Goal: maintain 2x CAGR improvement every 1.5 years in system performance per unit power/area/cost/...



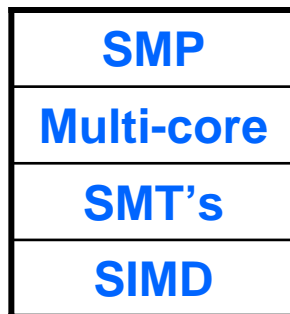
Challenge: Software contribution must exceed 2x every 1.5 years!

Hardware contribution reducing to 2x every 3-5 years

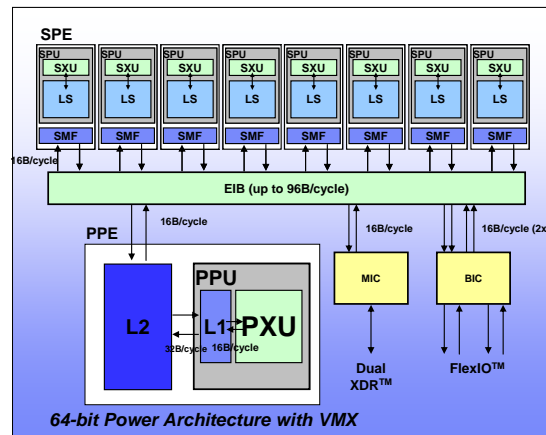
Q1: Can software overcome Wirth's Law in time to save Moore's Law?

Answer: Yes! The only way to save Moore's Law is through parallelism, and the only way to exploit parallelism (beyond ILP) is through software.


Scale-up
Parallelism



Heterogeneous Parallelism
(Co-processors, accelerators)



Scale-out
Parallelism



Parallelism in all the above cases can only be exploited through software: operating system, optimizing compilers, virtual machines, libraries, middleware, algorithms, applications, ...

Q2: What benefits can we expect from the new high-productivity languages under development?

1) **Safety -- eliminate entire classes of errors through static & dynamic safety checks**

- Type errors, initialization errors, pointer errors, array indexing errors – no accesses to inconsistent data
- Branching errors – no branches to unpredictable code locations
- Higher-level semantic errors – deadlock, data races, writes to immutable data, violations of preconditions/postconditions

2) **High level concurrency – integration of threads with a partitioned global address space and scalable synchronization operations**

- Unified model for fine-grain multithreading, one-sided data transfers, remote atomic operations, active messages, ...
- Relaxed memory model for synchronization operations

3) **Optimized implementations on high-end HPC systems**

- Early implementations of PGAS languages (UPC, CAF) point the way

4) **Ecosystem: Integration with Environment, Libraries, and Tools**

- Use of tools integrated into open & extensible development environments, such as Eclipse
- First-class support for components and libraries

Q3: Is it possible to use portable software with the latest hardware technologies? (Graphics processors, Cell, PCA, FPGAs, etc.)

Answer: Yes!
Virtual machines and integrated tools are the key to portability

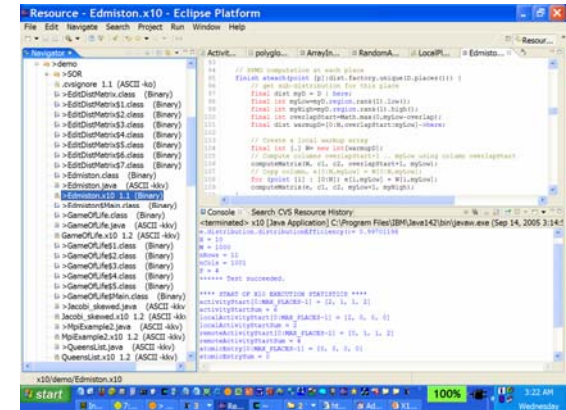
Real-time VM Enhancements --- IBM Metronome project

HPC VM Enhancements --- IBM PERCS/X10 project

Commodity Virtual Machines



Eclipse as a foundation for common tools




PERCS
5

HPEC 2005 Panel Session, Sep 2005

V. Sarkar

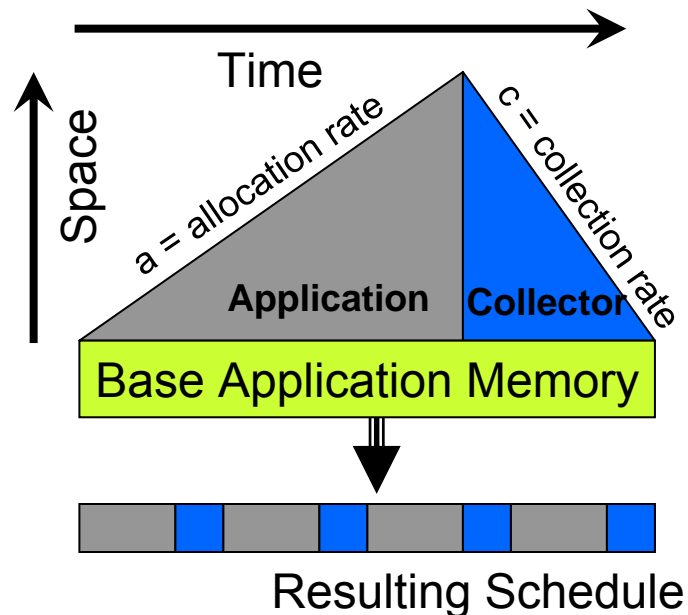




IBM Metronome project: Portable Real-time Garbage Collection

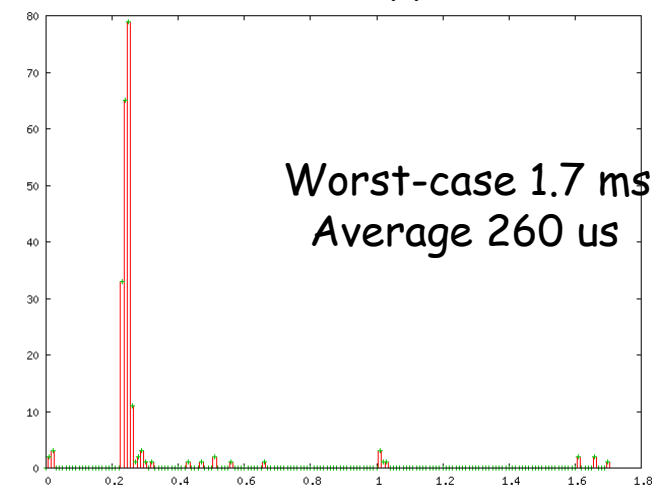
David Bacon, Perry Cheng, David Grove, V.T. Rajan, Martin Vechev

- **Garbage collection is fundamental to Java's value proposition**
 - Safety, reliability, programmer productivity
 - But also causes the most non-determinism (100 ms – 10 s latencies)
 - RTSJ standard does not support use of garbage collection for real-time
- **Metronome is our hard real-time garbage collector**
 - Worst-case 2 ms latencies; high throughput and utilization
 - 100x better than competitors' best garbage collection technology



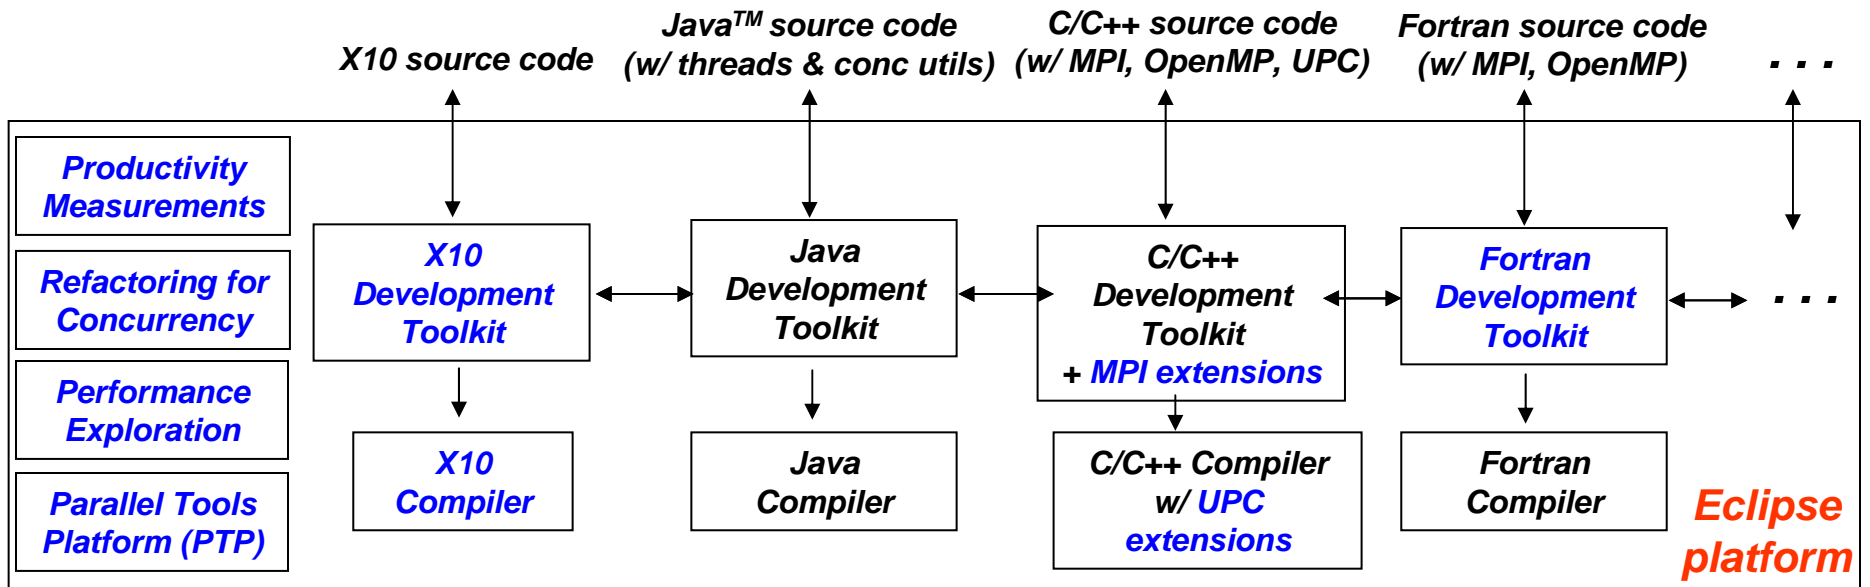
JACS
PERCS

Garbage Collection Pause Times
(Customer application)

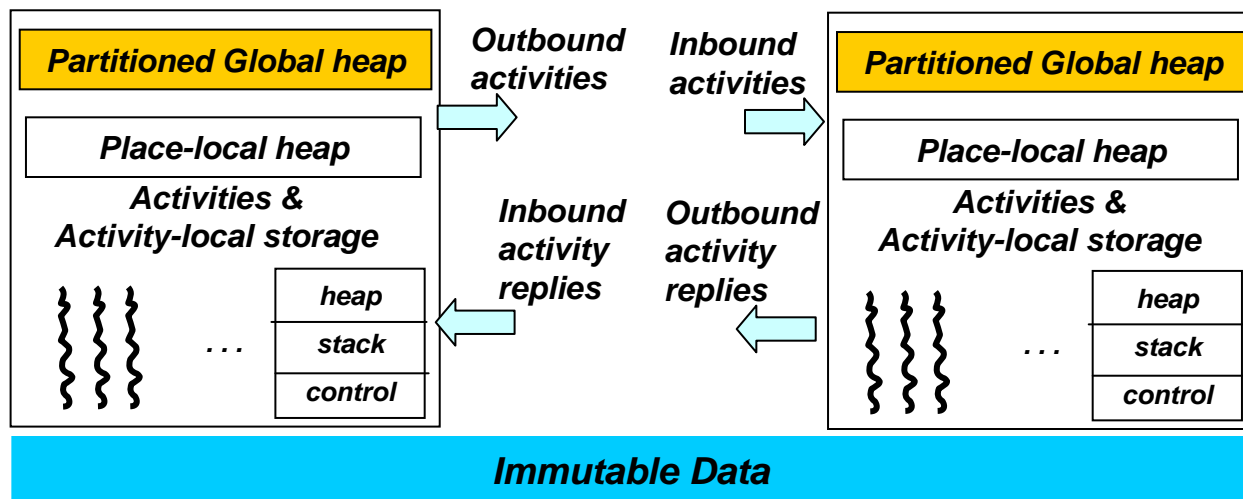


PERCS Programming Model, Tools and Compilers

(PERCS = Productive Easy-to-use Reliable Computer Systems)



Text in blue identifies exploratory PERCS contributions



X10 Threaded Partitioned Global Address Space model (T-PGAS)

