

Simultaneous Multithreading: Multiplying Alpha Performance

Dr. Joel Emer
Principal Member Technical Staff
Alpha Development Group
Compaq Computer Corporation

COMPAQ

www.compaq.com

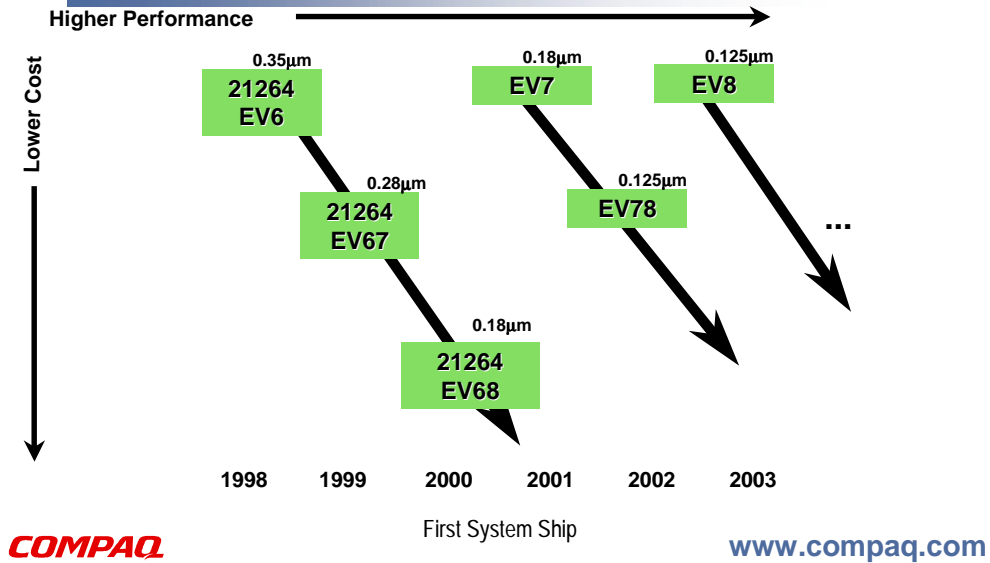
Outline

- ◆ Alpha Processor Roadmap
- ◆ Motivation for Introducing SMT
- ◆ Implementation of an SMT CPU
- ◆ Performance Estimates
- ◆ Architectural Abstraction

COMPAQ

www.compaq.com

Alpha Microprocessor Overview



EV8 Technology Overview

- ◆ Leading edge process technology – 1.2-2.0GHz
 - 0.125µm CMOS
 - SOI-compatible
 - Cu interconnect
 - low-k dielectrics
- ◆ Chip characteristics
 - ~1.2V V_{dd}
 - ~250 Million transistors
 - ~1100 signal pins in flip chip packaging

EV8 Architecture Overview

- ◆ Enhanced out-of-order execution
- ◆ 8-wide superscalar
- ◆ Large on-chip L2 cache
- ◆ Direct RAMBUS interface
- ◆ On-chip router for system interconnect
- ◆ Glueless, directory-based, ccNUMA for up to 512-way SMP
- ◆ 4-way simultaneous multithreading (SMT)

COMPAQ

www.compaq.com


Goals

- ◆ Leadership single stream performance
- ◆ Extra multistream performance with multithreading
 - Without major architectural changes
 - Without significant additional cost

COMPAQ

www.compaq.com

Instruction Issue




Reduced function unit utilization due to dependencies

COMPAQ

www.compaq.com

Superscalar Issue




Superscalar leads to more performance, but lower utilization

COMPAQ

www.compaq.com

Predicated Issue




Adds to function unit utilization, but results are thrown away

COMPAQ

www.compaq.com

Chip Multiprocessor




Limited utilization when only running one thread

COMPAQ

www.compaq.com

Fine Grained Multithreading




Intra-thread dependencies still limit performance

COMPAQ

www.compaq.com

Simultaneous Multithreading

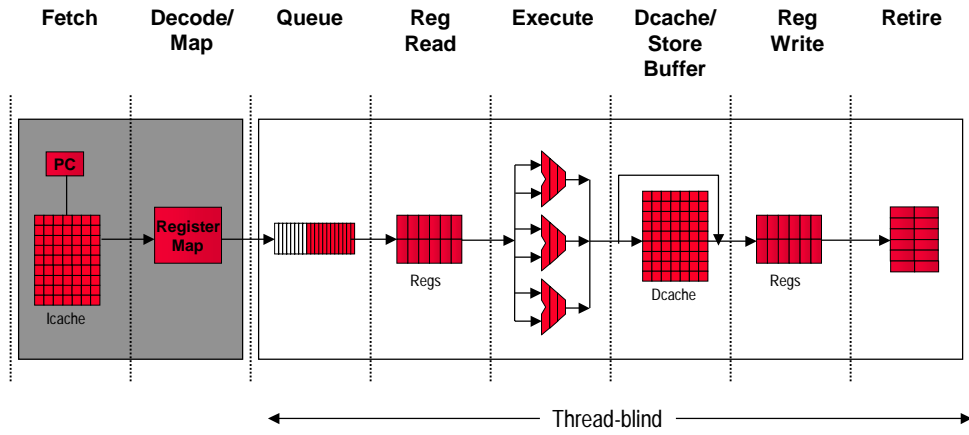


Maximum utilization of function units by independent operations

COMPAQ

www.compaq.com

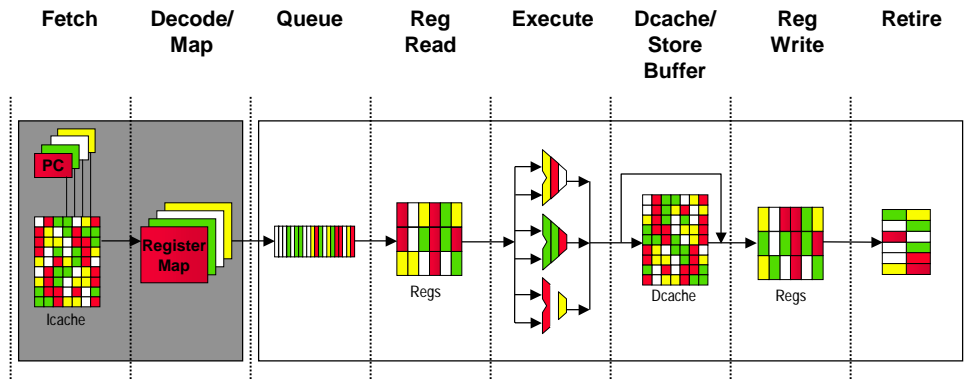
Basic Out-of-order Pipeline



COMPAQ

www.compaq.com

SMT Pipeline



COMPAQ

www.compaq.com

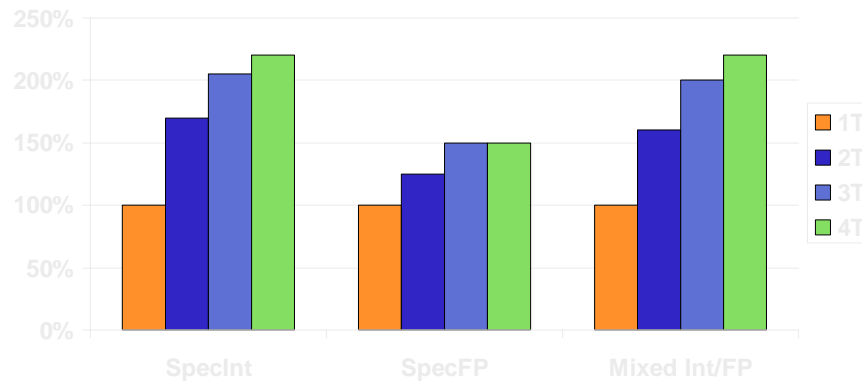
Changes for SMT

- ◆ Basic pipeline – unchanged
- ◆ Replicated resources
 - Program counters
 - Register maps
- ◆ Shared resources
 - Register file (size increased)
 - Instruction queue
 - First and second level caches
 - Translation buffers
 - Branch predictor

COMPAQ

www.compaq.com

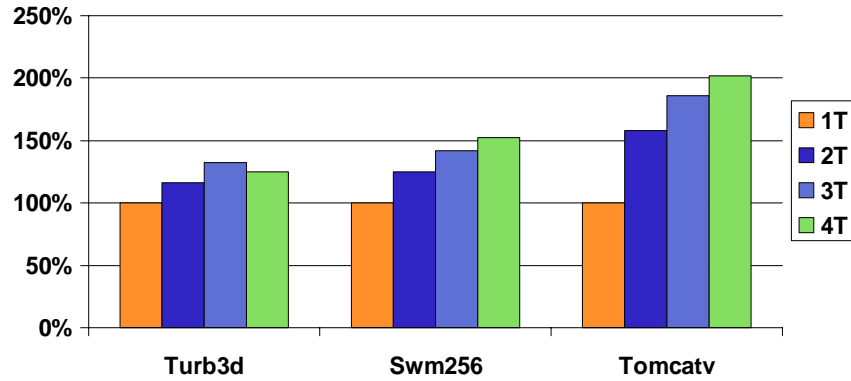
Multiprogrammed workload



COMPAQ

www.compaq.com

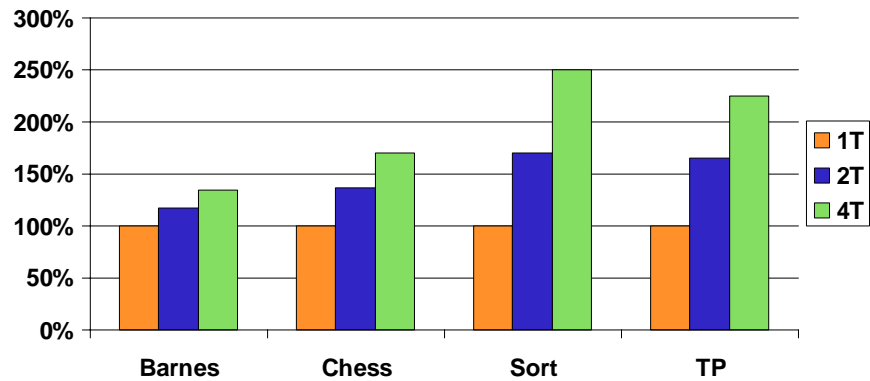
Decomposed SPEC95 Applications



COMPAQ

www.compaq.com

Multithreaded Applications




COMPAQ

www.compaq.com

Architectural Abstraction

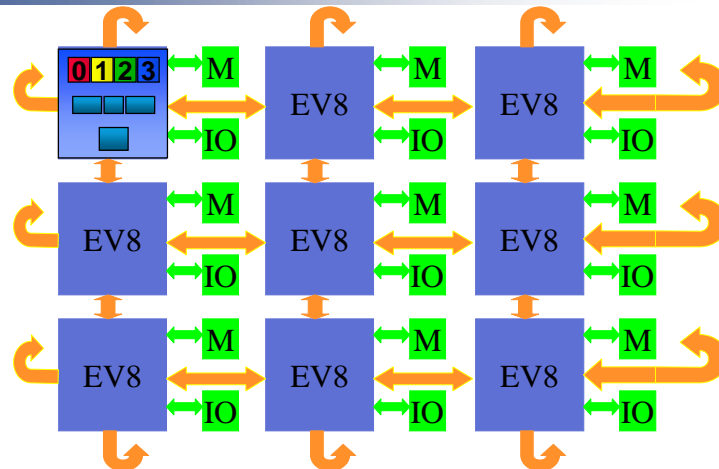
- ◆ 1 CPU with 4 Thread Processing Units (TPUs)
- ◆ Shared hardware resources



COMPAQ

www.compaq.com

System Block Diagram



COMPAQ

www.compaq.com

Quiescing Idle Threads

- ◆ Problem:
Spin looping thread consumes resources
- ◆ Solution:
Provide quiescing operation that allows a TPU to sleep until a memory location changes

COMPAQ

www.compaq.com

Summary

- ◆ Alpha will maintain single stream performance leadership
- ◆ SMT will significantly enhance multistream performance
 - Across a wide range of applications,
 - Without significant hardware cost, and
 - Without major architectural changes

COMPAQ

www.compaq.com

References

- ◆ "*Simultaneous Multithreading: Maximizing On-Chip Parallelism*" by Tullsen, Eggers and Levy in ISCA95.
- ◆ "*Exploiting Choice: Instruction Fetch and Issue on an Implementable Simultaneous Multithreaded Processor*" by Tullsen, Eggers, Emer, Levy, Lo and Stamm in ISCA96.
- ◆ "*Converting Thread-Level Parallelism to Instruction-Level Parallelism via Simultaneous Multithreading*" by Lo, Eggers, Emer, Levy, Stamm and Tullsen in ACM Transactions on Computer Systems, August 1997.
- ◆ "Simultaneous Multithreading: A Platform for Next-Generation Processors" by Eggers, Emer, Levy, Lo, Stamm and Tullsen in IEEE Micro, October, 1997.

COMPAQ

www.compaq.com