Presidential Young Investigator Award: Cache Memory Design (MIPS-8957278)

Final Report

Principal Investigator: Mark D. Hill

July 1995

1. Summary

This research targeted the design and evaluation of the memory systems for high-performance uniprocessors and shared-memory multiprocessors. Memory system design is important, because it largely determines a computer's sustained performance. In particular, we investigated *caches*, which are small, fast (relative to main memory) buffers that hold recently-used blocks of data. Insofar as programs reuse data and instructions, caches create the illusion of low-latency, high-bandwidth memory. Using caches in multiprocessors presents the additional challenges of making them transparent to applications (*cache coherence*) and specifying shared-memory semantics (*memory consistency model*).

Two themes run through this research. First, memory system design is a quantitative enterprise, in which new design scenarios often require new evaluation techniques. Second, memory system design is a software-hardware tradeoff, in which selected assistance from operating systems, compilers, and even application programmers can be tremendously useful. Contributions include:

- improved the design and evaluation methods for multi-megabyte level-two caches,
- refined memory consistency model options,
- investigated the use of multiple page sizes and subblocking in translation lookaside buffers and page tables, and
- seeded the **Wisconsin Wind Tunnel Project's** work on the design and evaluation of multiprocessor memory systems.

2. References

- Improved the design and evaluation methods for multi-megabyte level-two caches: [KeH92], [WHK91], [KHW94], and **Richard Kessler's** Ph.D. Thesis [Kes91] (first employment: Cray Research).
- Refined memory consistency model options: [AdH90b], [AdH90a], [AHM91], [GAG92], [AdH93] and **Sarita Adve's** Ph.D. Thesis [Adv93] (first employment: Rice University).
- Investigated the use of multiple page sizes and subblocking in translation lookaside buffers and page tables: [TKH92], [TaH94], [TaH95], and **Madhu Talluri's** Ph.D. Thesis [Tal95] (first employment: Sun Microsystems).
- Seeded the **Wisconsin Wind Tunnel Project's** work on the design and evaluation of multiprocessor memory systems: [HLR93], [RHL93], [WCF93], [MuH94], [FLR94], [MSH95], [WoH95], [THK95], [HLW95], and URL `http://www.cs.wisc.edu/~wwt` [HLW94].
- Other work on uniprocessor memory systems: [PnH90], [KHW91], [HLL93], [GHP93], and [PoH93].
- Other work on multiprocessor memory systems: [HiL90], [AAH91], and [Hil92].

[AdH90a] S. V. ADVE and M. D. HILL, Implementing Sequential Consistency In Cache-Based Systems, *Proc. International Conference on Parallel Processing*(August 1990), I-47-50.

[AdH90b] S. V. ADVE and M. D. HILL, Weak Ordering - A New Definition, *Proc. 17th Annual Symposium on Computer Architecture, Computer Architecture News*, 18, 2 (June 1990), 2-14, ACM.

[AAH91] S. V. ADVE, V. S. ADVE, M. D. HILL and M. K. VERNON, Comparison of Hardware and Software Cache Coherence Schemes, *Proc. 18th Annual Symposium on Computer Architecture, Computer Architecture News*, 19, 2 (June 1991), 298-308, ACM.

[AHM91] S. V. ADVE, M. D. HILL, B. P. MILLER and R. H. B. NETZER, Detecting Data Races on Weak Memory Systems, *Proc. 18th Annual Symposium on Computer Architecture, Computer Architecture News*, 19, 2 (June 1991), 234-243, ACM.

[Adv93] S. V. ADVE, Designing Memory Consistency Models for Shared-Memory Multiprocessors, Ph.D. Thesis, Computer Sciences Technical Report 1198, University of Wiscosin, Madison (November 1993).

[AdH93] S. V. ADVE and M. D. HILL, A Unified Formalization of Four Shared-Memory Models, *IEEE Trans. on Parallel and Distributed Systems*, 4, 6 (June 1993), 613-624.

[FLR94] B. FALSAFI, A. LEBECK, S. REINHARDT, I. SCHOINAS, M. D. HILL, J. LARUS, A. ROGERS and D. WOOD, Application-Specific Protocols for User-Level Shared Memory, *Proc. Supercomputing '94*, Washington (November 1994), 380-389.

[GHP93] J. D. GEE, M. D. HILL, D. N. PNEVMATIKATOS and A. J. SMITH, *Cache Performance of the SPEC92 Benchmark Suite*, IEEE Micro, (August 1993).

[GAG92] K. GHARACHORLOO, S. V. ADVE, A. GUPTA, J. L. HENNESSY and M. D. HILL, Programming for Different Memory Consistency Models, *Journal of Parallel and Distributed Computing*, 15, 4 (August 1992), 399-407.

[HiL90] M. D. HILL and J. R. LARUS, Cache Considerations for Multiprocessor Programmers, *Communications of the ACM*, 33, 8 (August 1990), 97-102.

[Hil92] M. D. HILL, What is Scalability?, in *Scalable Shared Memory Multiprocessors*, Kluwer Academic Publishers (1992), 89-96.

[HLL93] M. D. HILL, J. R. LARUS, A. R. LEBECK, M. TALLURI and D. A. WOOD, Wisconsin Architectural Research Tool Set, Vol. 21 (August 1993).

[HLR93] M. D. HILL, J. R. LARUS, S. K. REINHARDT and D. A. WOOD, Cooperative Shared Memory: Software and Hardware for Scalable Multiprocessors, *ACM Trans. on Computer Systems*, 11, 4 (November 1993), 300-318.

[HLW94] M. D. HILL, J. R. LARUS and D. A. WOOD, The Wisconsin Wind Tunnel Project: An Annotated Bibliography, *ACM SIGARCH Computer Architecture News*, 22, 5 (December 1994), 19-26.

[HLW95] M. D. HILL, J. R. LARUS and D. A. WOOD, Tempest: A Substrate for Portable Parallel Programs, *Proc. Compcon*(March 1995), 327-332.

[Kes91] R. E. KESSLER, Analysis of Multi-Megabyte Secondary CPU Cache Memories, Computer Sciences Technical Report #1032, Univ. of Wisconsin (July 1991).

[KeH92] R. E. KESSLER and M. D. HILL, Page Placement Algorithms for Large Real-Index Caches, *ACM Trans. on Computer Systems*, 10, 4 (November 1992), 338-359.

[KHW94] R. E. KESSLER, M. D. HILL and D. A. WOOD, A Comparison of Trace-Sampling Techniques for Multi-Megabyte Caches, *IEEE Trans. on Computers*, 43, 6 (June 1994), 664-675.

[KHW91] Y. H. KIM, M. D. HILL and D. A. WOOD, Implementing Stack Simulation for Highly-Associative Memories, *Proc. ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems*(May 1991), 212-213.

[MuH94] S. S. MUKHERJEE and M. D. HILL, An Evaluation of Directory Protocols for Medium-Scale Shared-Memory Multiprocessors, *Proc. of the International Conference on Supercomputing*

(ICS), Manchester, England (July 11-15, 1994), 64-74.

[MSH95] S. MUKHERJEE, S. SHARMA, M. D. HILL, J. LARUS, A. ROGERS and J. SALTZ, Efficient Support for Irregular Applications on Distributed-Memory Machines, *Proc. 5th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*(to appear July 1995).

[PnH90] D. N. PNEVMATIKATOS and M. D. HILL, Cache Performance of the Integer SPEC Benchmarks on a RISC, *Computer Architecture News*, 18, 2 (June 1990), 53-68.

[PoH93] A. F. POUR and M. D. HILL, Performance Implications of Tolerating Cache Faults, *IEEE Trans. on Computers*, 42, 3 (March 1993), 257-267.

[RHL93] S. K. REINHARDT, M. D. HILL, J. R. LARUS, A. R. LEBECK, J. C. LEWIS and D. A. WOOD, The Wisconsin Wind Tunnel: Virtual Prototyping of Parallel Computers, *Proc. ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems*(May 1993), 48-60.

[TKH92] M. TALLURI, S. KONG, M. D. HILL and D. A. PATTERSON, Tradeoffs in Supporting Two Page Sizes, *Proc. 19th Annual Symposium on Computer Architecture, Computer Architecture News*, 20, 2 (May 1992), 415-424, ACM.

[TaH94] M. TALLURI and M. D. HILL, Surpassing the TLB Performance of Superpages with Less Operating System Support, *Proc. Sixth International Conference on Architectural Support for Programming Languages and Operating Systems*, San Jose (October 1994), 171-182.

[TaH95] M. TALLURI and M. D. HILL, A New Page Table for 64-bit Address Spaces, *To Appear in Symposium on Operating System Principles*(December 1995).

[Tal95] M. TALLURI, Use of Superpages and Subblocking in the Address Translation Hierarchy, Computer Sciences Technical Report #???, Univ. of Wisconsin (expected August 1995).

[THK95] F. TRAENKLE, M. D. HILL and S. KIM, Solving Microstructure Electrostatics on a Proposed Parallel Computer, *Computers and Chemical Engineering*, 19 (1995), 743-757. Also Univ. of Wisconsin Computer Sciences Technical Report #1223 (April 1994).

[WHK91] D. A. WOOD, M. D. HILL and R. E. KESSLER, A Model for Estimating Trace-Sample Miss Ratios, *Proc. ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems*(May 1991), 79-89.

[WCF93] D. A. WOOD, S. CHANDRA, B. FALSAFI, M. D. HILL, J. R. LARUS, A. R. LEBECK, J. C. LEWIS, S. S. MUKHERJEE, S. PALACHARLA and S. K. REINHARDT, Mechanisms for Cooperative Shared Memory, (Also appeared as Univ. of Wisconsin Computer Sciences Technical Report #1142, March 1993.) *Proc. 20th Annual Symposium on Computer Architecture, Computer Architecture News*(May 1993), 156-167.

[WoH95] D. A. WOOD and M. D. HILL, Cost-Effective Parallel Computing, *IEEE Computer*, 28, 2 (February 1995), 69-72.