

# BIBLIOGRAPHY

- [1] V. Agarwal, M.S. Hrishikesh, S.W. Keckler and D. Burger. Clock Rate versus IPC: The End of the Road for Conventional Microarchitectures. In *Proc. of the 27th Ann. Intl. Symp. on Computer Architecture*, pp. 248-259, June 2000.
- [2] V. Agarwal, S.W. Keckler, D. Burger. The Effect of Technology Scaling on Microarchitectural Structures. Tech. Report #TR2000-02, Dept. of Computer Sciences, The Univ. of Texas at Austin, May 2001.
- [3] A. Aggarwal and M. Franklin. An Empirical Study of the Scalability Aspects of Instruction Distribution Algorithms for Clustered Processors. In *Proc. of the Intl. Symp. on Performance Analysis of Systems and Software*, pp. 172-179, November 2001.
- [4] A. Aggarwal and M. Franklin. Hierarchical Interconnects for On-chip Clustering. In *Proc. of the Intl. Parallel and Distributed Processing Symposium*, pp. 63-70, April 2002.
- [5] A. Aggarwal and M. Franklin. Instruction Replication: Reducing Delays due to Inter-PE Communication Latency. In *Proc. of the Intl. Conf. on Parallel Architectures and Compilation Techniques*, pp. 46-55, October 2003.
- [6] H. Akkary and M.A. Driscoll, A Dynamic Multithreading Processor. In *Proc. of the 31st. Ann. Intl. Symp. on Microarchitecture*, pp. 226-236, December 1998.
- [7] R.I. Bahar and S. Manne. Power and Energy Reduction Via Pipeline Balancing. In *Proc. of the 28th Ann. Intl. Symp. on Computer Architecture*, pp. 218-229, July 2001.
- [8] R. Balasubramonian, S. Dwarkadas and D. Albonesi. Dynamically Managing the Communication-Parallelism Trade-off in Future Clustered Processors. In *Proc. of the 30th. Ann. Intl. Symp. on Computer Architecture*, pp. 275-286, June 2003.
- [9] S. Banerjia. *Instruction Scheduling and Fetch Mechanisms for Clustered VLIW Processors*. PhD thesis, Dept. of Electrical and Computer Engineering, North Carolina State University, 1998.

- [10] A. Baniasadi, and A. Moshovos. Instruction Distribution Heuristics for Quad-Cluster, Dynamically-Scheduled, Superscalar Processors. In *Proc. of the 33rd. Ann. Intl. Symp. on Microarchitecture*, pp. 337-347, December 2000.
- [11] R. Bhargava, and L.K. John. Improving Dynamic Cluster Assignment for Clustered Trace Cache processors. In *Proc. of the 30th. Ann. Intl. Symp. on Computer Architecture*, pp. 264-274, June 2003.
- [12] M.T. Bohr. Interconnect Scaling - The Real Limiter to High Performance ULSI, in *Proc. of the 1995 IEEE Intl. Electron Devices Meeting*, pp. 241-244, 1995.
- [13] E. Borch, E. Tune, S. Manne and J. Emer. Loose Loops Sink Chips. In *Proc. of the 8th. Intl. Symp. on High-Performance Computer Architecture*, pp. 270-281, February 2002.
- [14] S. Borkar. Design challenges of technology scaling. *IEEE Micro*, 19(4): 23-29, July-August 1999.
- [15] D. Burger, T.M. Austin, S. Bennett. Evaluating Future Microprocessors: The SimpleScalar Tool Set, Technical Report CS-TR-96-1308, University of Wisconsin-Madison, 1996.
- [16] R. Canal, J.-M. Parcerisa, A. Gonzalez, A Cost-Effective Clustered Architecture, in *Proc. of the Intl. Conf. on Parallel Architectures and Compilation Techniques*, pp. 160-168, October 1999.
- [17] R. Canal, J.-M. Parcerisa, A. González. Dynamic Cluster Assignment Mechanisms. In *Proc. of the 6th. Intl. Symp. on High-Performance Computer Architecture*, pp. 132-142, January 2000.
- [18] R. Canal, J.-M. Parcerisa, and A. González. Dynamic Code Partitioning for Clustered Architectures. *International Journal of Parallel Programming*, Kluwer Academic/Plenum Publishers, 29 (1): 59-79, February 2001
- [19] A. Capitano, N. Dutt, and A. Nicolau. Partitioned Register Files for VLIWs: A Preliminary Analysis of Tradeoffs. In *Proc. of the 25th. Ann. Intl. Symp. on Microarchitecture*, pp. 292-300, December 1992.
- [20] R. P. Colwell, R. P. Nix, J.J. O'Donnell, D.P. Papworth, and P.K. Rodman. A VLIW architecture for a trace scheduling compiler. In *Proc. of the 2nd. Intl. Conf. on Architectural Support for Programming Languages and Operating Systems*, October 1987.
- [21] J.-L.Cruz, A.González, M. Valero and N.P. Topham. Multiple-Banked Register File Architectures. In *Proc. of the 27th. Intl. Symp. on Computer Architecture*, pp. 316-324, June 2000.
- [22] G. Desoli. Instruction Assignment for Clustered VLIW DSP Compilers: A New Approach. Tech. Report HPL-98-13, HP Labs, January 1998.

- [23] J. Duato, S. Yalamanchili, L. Ni, *Interconnection Networks, An Engineering Approach*, IEEE Computer Society Press, 1997.
- [24] P.K. Dubey, K. O'Brien, K.M. O'Brien and C. Barton, Single-Program Speculative Multithreading (SPSM) Architecture: Compiler-Assisted Fine-Grained Multithreading, in *Proc. of the Intl. Conf. on Parallel Architectures and Compilation Techniques*, pp. 109-121, 1995.
- [25] R.J. Eickemeyer and S. Vassiliadis. A load-instruction unit for pipelined processors. *IBM Journal of Research and Development*, 37(4): 547-564, July 1993.
- [26] J.R. Ellis. *Bulldog: A Compiler for VLIW Architectures*. PhD thesis, Yale University, 1985.
- [27] K.I. Farkas, N.P. Jouppi and P. Chow. Register File Design Considerations in Dynamically Scheduled Processors, in *Proc. of the 2nd. Intl. Symp. on High-Performance Computer Architecture*, pp. 40-51, 1996.
- [28] K.I.Farkas. *Memory-system Design Considerations for Dynamically-scheduled Microprocessors*, PhD thesis, Department of Electrical and Computer Engineering, Univ. of Toronto, Canada, January 1997.
- [29] K.I. Farkas, P. Chow, N.P. Jouppi, and Z. Vranesic. The Multicluster Architecture: Reducing Cycle Time Through Partitioning. In *Proc. of the 30th. Ann. Intl. Symp. on Microarchitecture*, pp. 149-159, December 1997.
- [30] M.M. Fernandes, J.Llosa and N.Topham, Distributed Modulo Scheduling, in *Proc. of the 5th. Intl. Symp. on High-Performance Computer Architecture*, pp. 130-134, January 1999.
- [31] B. Fields, S. Rubin, and R. Bodík. Focusing Processor Policies via Critical-Path Prediction. In *Proc. of the 28th. Ann. Intl. Symp. on Computer Architecture*, pp. 74-85, June 2001.
- [32] M.J. Flynn, P. Hung, and K. Rudd. Deep-Submicron Microprocessor Design Issues. *IEEE Micro*, 19(4): 11-22, July/August 1999.
- [33] M. Franklin and G.S. Sohi. The Expandable Split Window Paradigm for Exploiting Fine-Grain Parallelism. In *Proc. of the 19th Ann. Intl. Symp. on Computer Architecture*, pp.58-67, May 1992.
- [34] M. Franklin. *The Multiscalar Architecture*. PhD. Thesis, Technical Report TR 1196, Computer Sciences Department, University of Wisconsin-Madison, November 1993.
- [35] J. Fridman and Z. Greenfield. The TigerSHARC DSP Architecture. *IEEE Micro*, 20(1): 66-76, January-February 2000.
- [36] D.H. Friendly, S.J. Patel, and Y.N. Patt. Putting the Fill Unit to Work: Dynamic Optimizations for Trace Cache Processors. In *Proc. of the 31st. Ann. Intl. Symp. on Microarchitecture*, pp. 173-181, November 1998.

- [37] F.Gabbay and A.Mendelson. Speculative Execution Based on Value Prediction, Tech. Report #1080, Univ. Technion, Israel, 1996.
- [38] J.González and A.González. Memory Address Prediction for Data Speculation. Tech. Report #UPC-DAC-1996-50, Univ. Politècnica de Catalunya, Spain, 1996.
- [39] J.González and A.González. Speculative Execution Via Address Prediction and Data Prefetching. In *Proc. of the 11th. International Conference on Supercomputing*, pp. 196-203, July 1997.
- [40] J.González and A.González. The Potential of Data Value Speculation to Boost ILP. In *Proc. of the 12th. International Conference on Supercomputing*, pp. 21-28, 1998.
- [41] L. Gwennap. Intel's MMX Speeds Multimedia Instructions, *Microprocessor Report*, 10(3), March 1996.
- [42] L. Gwennap. Digital 21264 Sets New Standard. *Microprocessor Report*, 10 (14), October 1996.
- [43] R. Ho, K.W. Mai, M.A. Horowitz. The Future of Wires. *Proceedings of the IEEE*, 89(4): 490-504, April 2001.
- [44] M.S. Hrishikesh, N.P. Jouppi, K.I. Farkas, D. Burger, S.W. Keckler, and P. Shivakumar. The Optimal Logic Depth per Pipeline Stage Is 6 to 8 FO4 Inverter Delays. In *Proc. of the 29th Ann. Intl. Symp. on Computer Architecture*, pp. 14-24, May 2002.
- [45] *The International Technology Roadmap for Semiconductors*. Semiconductor Industry Association. 1999.
- [46] *The International Technology Roadmap for Semiconductors*. Semiconductor Industry Association. 2002.
- [47] A. Iyer and D. Marculescu. Power and Performance of Globally Asynchronous Locally Synchronous Processors. In *Proc. of the 29th Ann. Intl. Symp. on Computer Architecture*, pp.158-168, May 2002.
- [48] Q. Jacobson, E. Rotenberg, and J.E. Smith. Path-Based Next Trace Prediction. In *Proc. of the 30th. Ann. Intl. Symp. on Microarchitecture*, pp. 14-23, December 1997.
- [49] J. Johnson. Expansion Caches for Superscalar Processors. Tech. Report CSL-TR-94-630, Computer Science Laboratory, Stanford University, June 1994.
- [50] M. Johnson. *Superscalar Microprocessor Design*. Ed. Prentice-Hall, Englewood Cliffs, New Jersey, 1991.

- [51] S.W. Keckler, D. Burger, C.R. Moore, R. Nagarajan, K. Sankaralingam, V. Agarwal, M.S. Hrishikesh, N. Ranganathan, P. Shivakumar. A Wire-Delay Scalable Microprocessor Architecture for High Performance Systems. In *Proc. of the 2003 IEEE Intl. Solid-State Circuits Conference*, paper no. 9.6, February 2003.
- [52] G.A. Kemp, and M.Franklin. PEWs: A Decentralized Dynamic Scheduler for ILP Processing. In *Proc. of Intl. Conf. on Parallel Processing*, pp. 239-246, August 1996.
- [53] R.E. Kessler. The Alpha 21264 Microprocessor. *IEEE Micro*, 19(2):24-36, 1999.
- [54] H-S.Kim and J.E.Smith. An Instruction Set and Microarchitecture for Instruction Level Distributed Processing. In *Proc. of the 29th Ann. Intl. Symp. on Computer Architecture*, pp. 71-87, May 2002
- [55] K. Krewell. Intel Embraces Multithreading, *Microprocessor Report*, Sept. 2001, pp. 1-2.
- [56] C. Lee, M. Potkonjak and W. H. Mangione-Smith, Mediabench: A Tool for Evaluating and Synthesizing Multimedia and Communications Systems, In *Proc. of the 30th IEEE/ACM Intl. Symposium on Microarchitecture*, pp. 330-335, December 1997.
- [57] M.H.Lipasti, C.B.Wilkerson, and J.P.Shen. Value Locality and Load Value Prediction. In *Proc. of the 7th. Intl. Conf. on Architectural Support for Programming Languages and Operating Systems*, pp.138-147, October 1996.
- [58] M. Lipasti. *Value Locality and Speculative Execution*. PhD thesis, Carnegie Mellon University, April 1997.
- [59] P. Marcuello, A. González and J. Tubella, Speculative Multithreaded Processors, in *Proc of the 12th ACM Intl. Conf. on Supercomputing*, pp 77-84, July 1998.
- [60] P. Marcuello and A. González, Clustered Speculative Multithreaded Processors, *Proc. of the 13th ACM Intl. Conf. on Supercomputing*, pp. 365-372, June 1999.
- [61] D. Matzke. Will Physical Scalability Sabotage Performance Gains?. *IEEE Computer* 30(9): 37-39, September 1997.
- [62] Mediabench Home Page. URL: <http://www.cs.ucla.edu/~leec/mediabench/>
- [63] S. Melvin and Y. Patt. Performance Benefits of Large Atomic Units in Dynamically Scheduled Machines. in *Proc. of the 3rd. Intl. Conference on Supercomputing*, pp. 427-432, June 1989.
- [64] R. Nagarajan, K. Sankaralingam, D. Burger and S.W. Keckler. A Design Space Evaluation of Grid Processor Architectures. In *Proc of the 34th. Ann. Intl. Symp. on Microarchitecture*, pp. 40-51, 2001
- [65] E. Nystrom and A.E. Eichenberger, Effective Cluster Assignment for Modulo Scheduling, In *Proc. of the 31st. Ann. Symp. on Microarchitecture*, pp. 103-114, November 1998

- [66] P.S. Oberoi and G.S. Sohi. Parallelism in the Front-End. In *Proc. of the 30th. Ann. Intl. Symp. on Computer Architecture*, pp. 230-240, June 2003
- [67] E. Ozer, S. Banerjia, T.M. Conte. Unified Assign and Schedule: A New Approach to Scheduling for Clustered Register File Microarchitectures. in *Proc. of the 31st. Intl. Symp. on Microarchitecture*, pp. 308-315, November 1998.
- [68] S. Palacharla, and J.E. Smith. Decoupling Integer Execution in Superscalar Processors. In *Proc. of the 28th. Ann. Symp. on Microarchitecture*, pp. 285-290, November 1995.
- [69] S. Palacharla, N.P. Jouppi, and J.E. Smith. Quantifying the Complexity of Superscalar Processors. Tech. Report CS-TR-961328, Univ. of Wisconsin-Madison, November 1996.
- [70] S. Palacharla, N.P. Jouppi, and J.E. Smith. Complexity-Effective Superscalar Processors. In *Proc. of the 24th. Intl. Symp. on Computer Architecture*, pp. 206-218, June 1997.
- [71] S.Palacharla. *Complexity-Effective Superscalar Processors*. Ph.D. thesis, Univ. of Wisconsin-Madison, 1998.
- [72] J.-M. Parcerisa and A. González, The Latency Hiding Effectiveness of Decoupled Access/Execute Processors. In *Proc. of the 24th. Euromicro Conference*, Vasteras, Sweden, pp. 293-300, August 1998.
- [73] J.-M. Parcerisa and A. González, Reducing Wire Delay Penalty through Value Prediction, In *Proc. of the 33rd. Intl. Symp. on Microarchitecture*, pp. 317-326, December 2000.
- [74] J.-M. Parcerisa, A. González and J.E. Smith. A Clustered Front-End for Superscalar Processors. Tech. Report #UPC-DAC-2002-29, Computer Architecture Dept., Univ. Politècnica de Catalunya, Spain, July 2002.
- [75] J.-M. Parcerisa, J. Sahuquillo, A. González and J. Duato. Efficient Interconnects for Clustered Microarchitectures. In *Proc. of the 11th Intl. Conf. on Parallel Architectures and Compilation Techniques*, pp. 291-300, September 2002.
- [76] S. Patel, M. Evers, and Y. Patt. Improving Trace Cache Effectiveness with Branch Promotion and Trace Packing. In *Proc. of the 25th. Ann. Intl. Symp. on Computer Architecture*, pp. 262-271, June 1998.
- [77] A. Peleg and U. Weiser. Dynamic Flow Instruction Cache Memory Organized around Trace Segments Independent of Virtual Address Line. U.S. Patent Number 5,381,533, January 1995.
- [78] P. Racunas and Y.P. Patt. Partitioned First-Level Cache Design for Clustered Microarchitectures. In *Proc. of the 17th Ann. Intl. Conf. on Supercomputing*, pp. 22-31, June 2003

- [79] N. Ranganathan and M. Franklin. An Empirical Study of Decentralized ILP Execution Models. In *Proc. of the 8th. Intl. Conf. on Architectural Support for Programming Languages and Operating Systems*, pp. 272-281, October 1998.
- [80] N. Ranganathan and M. Franklin. The PEWs microarchitecture: reducing complexity through data-dependence based decentralization. *Microprocessors and Microsystems*, Elsevier B.V., 22(6): 333-343, November 1998.
- [81] E. Rotenberg, S. Bennet, and J.E. Smith. Trace Cache: A Low Latency Approach to High Bandwidth Instruction Fetching. In *Proc. of the 29th. Ann. Intl. Symp. on Microarchitecture*, pp 24-34, December 1996.
- [82] E. Rotenberg, Q. Jacobson, Y. Sazeides and J.E. Smith. Trace Processors. In *Proc. of the 30th Ann. Intl. Symp. on Microarchitecture*, pp. 138-148, December 1997.
- [83] E. Rotenberg. *Trace Processors: Exploiting Hierarchy and Speculation*. Ph.D. thesis, Univ. of Winsconsin-Madison, 1999.
- [84] J. Sánchez and A. González. Modulo Scheduling for a Fully-Distributed Clustered VLIW Architecture. In *Proc. of the 33rd. Ann. Intl. Symp. on Microarchitecture*, pp. 124-133, December 2000.
- [85] K. Sankaralingam, R. Nagarajan, S.W. Keckler, and D. Burger. A Technology-Scalable Architecture for Fast Clocks and High ILP. Tech. Report #TR2001-02, Dept. of Computer Sciences, Univ. of Texas, Austin, 2002.
- [86] K. Sankaralingam, V.A. Singh, S.W. Keckler and D. Burger. Routed Inter-ALU Networks for ILP Scalability and Performance. In *Proc. of the 21st. Intl. Conf. on Computer Design*, pp. 170-177, October 2003.
- [87] S.S. Sastry, S. Palacharla, and J.E. Smith, Exploiting Idle Floating-Point Resources For Integer Execution, in *Proc. of the Intl. Conf. on Programming Language Design and Implementation*, pp. 118-129, Montreal, June 1998.
- [88] Y. Sazeides, S. Vassiliadis, and J.E. Smith. The Performance Potential of Data Speculation and Collapsing. In *Proc. of the 29th Ann. Intl. Symp. on Microarchitecture*, pp. 238-247, December 1996.
- [89] A. Sez nec, E. Toullec, and O. Rochecouste. Register Write Specialization Register Read Specialization: A Path to Complexity-Effective Wide-Issue Superscalar Processors. In *Proc. of the 35th. Ann. Intl. Symp. on Microarchitecture*, pp. 383-394, November 2002.
- [90] J.E. Smith, Decoupled Acces/Execute Computer Architectures. *ACM Transactions on Computer Systems*, 2(4): 289-308, November 1984.

- [91] J.E. Smith and G.S. Sohi. The Microarchitecture of Superscalar Processors. *Proceedings of the IEEE*, 83(12):1609-1624, December 1995.
- [92] J.E. Smith. Instruction-Level Distributed Processing. *Computer*, IEEE Computer Society, 34(4): 59-65, April 2001.
- [93] G.S. Sohi, S.E. Breach and T.N. Vijaykumar. Multiscalar Processors. In *Proc. of the 22nd. Intl. Symp. on Computer Architecture*, pp. 414-425, June 1995.
- [94] S.P. Song, M. Denman, and J. Chang. The PowerPC 604 RISC Microprocessor. *IEEE Micro*, 14(5): 8-17, October 1994.
- [95] E. Sprangle, D. Carmean. Increasing Processor Performance by Implementing Deeper Pipelines. In *Proc. of the 29th Ann. Intl. Symp. on Computer Architecture*, pp. 25-34, May 2002.
- [96] V. Srinivasan, D. Brooks, M. Gschwind, P. Bose, V. Zyuban, P.N. Strensky, P.G. Emma. Optimizing Pipelines for Power and Performance. In *Proc. of the 35th. Ann. Intl. Symp. on Microarchitecture*, pp. 333-344, November 2002.
- [97] Standard Performance Evaluation Corporation, *SPEC Newsletter*, September 1995.
- [98] J. Stark, M.D. Brown and Y.N. Patt. On Pipelining Dynamic Instruction Scheduling Logic. In *Proc. of the 33th. Ann. Intl. Symp. on Microarchitecture*, pp. 57-66, Dec. 2000
- [99] J.M. Tendler, S. Dodson, S. Fields, H. Le, and B. Sinharoy, *POWER4 System Microarchitecture*, Technical white paper, IBM server group web site, October 2001
- [100] A. Terechko, E.L. Thenaff, M. Garg, J.van Eijndhoven, and H. Corporaal. Inter-cluster Communication Models for Clustered VLIW Processors. In *Proc. of the 9th. Intl. Symp. on High-Performance Computer Architecture*, 2003.
- [101] Texas Instrument Inc. TMS320C62x/67x CPU and Instruction Set Reference Guide, 1998.
- [102] T.N. Theis. The Future of Interconnection Technology. *IBM Journal of Research and Development*, 44(3):379-389, May 2000.
- [103] R. Tomasulo. An Efficient Algorithm for Exploiting Multiple Arithmetic Units. *IBM Journal of Research and Development*, 11(1):25-33, January 1967.
- [104] N.P. Topham, A. Rawsthorne, C.E. McLean, M.J.R.G. Mewissen and P. Bird. Compiling and Optimizing for Decoupled Architectures. In *Proc. of the Supercomputing '95*, San Diego, December 1995.
- [105] M. Tremblay, J. Chan, S. Chaundrhy, A.W. Conigliaro, S.S. Tse, The MAJC Architecture: A Synthesis of Parallelism and Scalability, *IEEE Micro* 20(6): 12-25, November/December 2000.



- [106] J-Y. Tsai and P-C. Yew, The Superthreaded Architecture: Thread Pipelining with Run-Time Data Dependence Checking and Control Speculation, in *Proc. of the Intl. Conf. on Parallel Architectures and Compilation Techniques*, pp. 35-46, 1996.
- [107] E. Tune, D. Liang, D.M. Tullsen, and B. Calder. Dynamic Prediction of Critical Path Instructions. In *Proc. of the 7th. Intl. Symp. on High-Performance Computer Architecture*, pp. 185-195, January 2001.
- [108] S. Vajapeyam and T. Mitra, Improving Superscalar Instruction Dispatch and Issue by Exploiting Dynamic Code Sequences. In *Proc. of the 24th. Intl. Symp. on Computer Architecture*, pp. 1-12, June 1997.
- [109] E. Waingold, M. Taylor, D. Srikrishna, V. Sarkar, W. Lee, V. Lee, J. Kim, M. Frank, P. Finch, R. Barua, J. Babb, S. Amarasinghe, A. Agarwal. Baring It All to Software: Raw Machines. *IEEE Computer*, pp. 86-93, September 1997.
- [110] K.C. Yeager. The MIPS R10000 Superscalar Microprocessor. *IEEE Micro*, 16(2): 28-41, April 1996.
- [111] A. Yoaz, M. Erez, R. Ronnen and S. Jourdan. Speculation Techniques for Improving Load Related Instruction Scheduling. In *Proc. of the 26th. Ann. Intl. Symp. on Computer Architecture*, pp. 42-53, 1999
- [112] V. Zyuban. *Inherently Lower-Power High-Performance Superscalar Architectures*, Ph.D. thesis, Univ. of Notre Dame, January 2000.
- [113] V.V.Zyuban and P.M. Kogge. Inherently Lower-Power High-Performance Superscalar Architectures. *IEEE Transactions on Computers* 50(3): 268-285, March 2001



# LIST OF FIGURES

## Chapter 2

Figure 2-1: Reference Clustered Architecture .....	23
Figure 2-2: Example of Map Table for 4 clusters.....	24
Figure 2-3: Performance improvements of splitting issue queues for copy/regular instructions, for various lengths of the issue queue and copy queue .....	25

## Chapter 3

Figure 3-1: Block diagram of a cost-effective clustered architecture.....	33
Figure 3-2: Example of a RDG.....	34
Figure 3-3: Static versus dynamic LdSt slice steering.....	36
Figure 3-4: LdSt slice versus Br slice steering: communications per instruction .....	38
Figure 3-5: LdSt slice versus Br slice steering: performance.....	38
Figure 3-6: LdSt slice versus Br slice steering: distribution of the workload imbalance, metric I2 (SpecInt95 average).....	39
Figure 3-7: Non-slice Balance steering versus slice steering: performance.....	39
Figure 3-8: Non-slice Balance steering versus Slice steering: number of communications per dynamic instruction (SpecInt95 average) .....	40
Figure 3-9: Non-slice Balance steering versus slice steering: distribution of the workload imbalance, metric I2 (SpecInt95 average) .....	40
Figure 3-10: Hardware support for the Slice Balance steering.....	41
Figure 3-11: Slice Balance steering: performance.....	42
Figure 3-12: Slice Balance steering: distribution of the workload imbalance, metric I2 (SpecInt95 average) .....	42
Figure 3-13: Priority Slice Balance steering: performance .....	43
Figure 3-14: General Balance steering .....	44
Figure 3-15: Performance comparison among all the proposed steering schemes.....	45
Figure 3-16: General Balance steering versus FIFO-based steering [70].....	46
Figure 3-17: Register replication on a cost-effective 2-clusters architecture .....	47
Figure 3-18: Performance of the SpecFP95 on a cost-effective 2-clusters architecture.....	48

## Chapter 4

Figure 4-1: Performance of Modulo, Simple RMB and Balanced RMB steering schemes on a four-cluster architecture.....	56
Figure 4-2: Average number of communications per dynamic instruction .....	56
Figure 4-3: Average NREADY workload imbalance .....	56
Figure 4-4: Distribution function of the NREADY workload imbalance (perl) .....	57
Figure 4-5: IPC of the Advanced RMB vs. Balanced RMB steering schemes on a four-cluster architecture.....	58
Figure 4-6: Average number of communications per dynamic instruction .....	58
Figure 4-7: Average NREADY workload imbalance .....	58
Figure 4-8: Distribution function of the NREADY workload imbalance (perl) .....	59
Figure 4-9: IPC of the Priority RMB vs. Advanced RMB steering schemes on a four-cluster architecture.....	60
Figure 4-10: IPC of the Priority RMB steering scheme with Accurate Rebalancing, and without it, for four clusters .....	61
Figure 4-11: The Advanced RMB versus the best slice-based steering schemes, on a cost-effective architecture (IPC speedups over a conventional superscalar) .....	62
Figure 4-12: Performance of the AR-Priority RMB, FIFO-based and MOD3 steering schemes, for four clusters .....	63
Figure 4-13: Performance sensitivity to the communication latency.....	64
Figure 4-14: Sensitivity to the interconnect bandwidth for four clusters (AR-Priority RMB steering) .....	66
Figure 4-15: Overall performance of the RMB steering schemes .....	67
Figure 4-16: Average number of communications per instruction .....	67
Figure 4-17: Average NREADY workload imbalance .....	67

## Chapter 5

Figure 5-1: IPC of baseline architectures, without value prediction.....	73
Figure 5-2: Impact of using value prediction on IPC.....	73
Figure 5-3: (a) Average inter-cluster communications ratio (b) Average workload imbalance.....	74
Figure 5-4: Speedups of value prediction, with PRMB and VPB steering schemes .....	75
Figure 5-5: Sensitivity of value prediction speedups to (a) communication latency, and (b) latency of the register read stage.....	76
Figure 5-6: Timing diagram of two instructions I1 and I2, where I2 mispredicts the value produced by I1, and must re-issue non-speculatively (the arrows show wakeup signals in case of hit and miss).....	77
Figure 5-7: Impact of value predictor table size for 4 clusters (a) prediction rate and accuracy (b) IPC.....	78

## Chapter 6

Figure 6-1: Sample timing of a communication between two dependent instructions I1 and I2, steered to clusters c1 and c2 respectively (solid arrows mean wakeup signals, hollow arrows mean data signals, and transmission time is 2 cycles in both cases).....	84
Figure 6-2: Router schemes for synchronous and asynchronous point-to-point interconnects .....	85
Figure 6-3: Four-cluster topologies .....	87
Figure 6-4: Additional topologies for 8 clusters. A black dot at the end of a link means that there is more than one link that can be followed for the next hop if the corresponding node is not the destination. Messages can always be routed through solid links but dashed links are only used for their last hop.....	91
Figure 6-5: Average contention delays of 1-hop and 2-hops messages, with synchronous, partially asynchronous and ideal ring interconnects.....	93
Figure 6-6: Comparing the IPC of 4-cluster interconnects.....	94
Figure 6-7: Comparing the IPC of 8-cluster interconnects.....	96
Figure 6-8: Effectiveness of the Accurate Rebalancing (AR) and Topology Aware (TA) techniques applied to the baseline PRMB steering scheme, for a four-cluster ring and an eight-cluster torus: (a) IPC, (b) Communications rate and (c) Workload imbalance.....	97

## Chapter 7

Figure 7-1: Accuracy vs. latency for a centralized and a 4-clustered predictors.....	102
Figure 7-2: Partitioning an 8-way processor's front-end into four clusters.....	103
Figure 7-3: Branch predictor (gshare) and instruction cache interleaving .....	105
Figure 7-4: Pipeline timing of conventional and clustered front-ends. The diagram shows the fetch of 3 blocks of instructions B0, B1 and B2, at addresses PC0, PC1 and PC2 respectively.....	106
Figure 7-5: Clustered and Baseline (centralized) integer pipelines. Fetch address generation loop is (1) in both pipelines, but it becomes (2) in the clustered one when switching predictor cluster, producing a bubble.....	106
Figure 7-6: Pipeline diagrams for three clustered steering logic schemes .....	108
Figure 7-7: Block diagram of the Dependence Check and Overriding (stage 5) for one cluster. Signals with subscript 1 belong to the current block of 8 instructions. Signals with subscript 0 belong to the previous block.....	109
Figure 7-8: IPC of the baseline and the clustered front-end schemes .....	112
Figure 7-9: Impact of the predictor bank switch penalty. A clustered predictor is compared to a centralized one.....	113
Figure 7-10: Performance impact of the interleaving bits, for a clustered front-end .....	113
Figure 7-11: Impact of using 2-cycle outdated renaming information by the steering heuristic: two clustered front-ends are compared, a naive scheme and our approach with assignment overriding (speed-ups are relative to a clustered front-end with replicated steering logic that uses updated information) .....	114



# LIST OF TABLES

## Chapter 2

Table 2-1: Default main architecture parameters .....	26
---	----

## Chapter 3

Table 3-1: Benchmarks and their inputs.....	35
Table 3-2: Machine parameters (split into cluster 1 and cluster 2 if not common).....	35

## Chapter 6

Table 6-1: Rules to secure a link in the source cluster src (D refers to distance in cycles).....	88
Table 6-2: Rules to secure a link in the source cluster src (D refers to distance in cycles).....	90
Table 6-3: Queue length distribution (for djpeg).....	9