

Limits of Parallelism

- [LoP1] G.S. Tjaden and M.J. Flynn. Detection and parallel execution of independent instructions. *IEEE Transactions on Computers*, C-19(10):889–895, 1970.
- [LoP2] E. M. Riseman and C. C. Foster. The inhibition of potential parallelism by conditional jumps. *IEEE Transactions on Computers*, C-21(12):1405–1411, 1972.
- [LoP3] A. Nicolau and J. A. Fisher. Measuring the parallelism available for very long instruction word architectures. *IEEE Transactions on Computers*, C-33(11):968–976, 1984.
- [LoP4] D. E. Culler and G. K. Maa. Assessing the benefits of fine-grain parallelism in dataflow programs. In *Proceedings of the 1988 ACM/IEEE Conference on Supercomputing*, pages 60–69, Orlando, FL, 1988.
- [LoP5] N. Jouppi and D. Wall. Available instruction-level parallelism for superscalar and superpipelined machines. In *Proceedings of the Third International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 272–282, Boston, MA, April 1989.
- [LoP6] M. D. Smith, M. Johnson, and M. A. Horowitz. Limits on multiple instruction issue. In *Proceedings of the Third International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 290–302, Boston, MA, 1989.
- [LoP7] D. W. Wall. Limits of instruction-level parallelism. In *Proceedings of the Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 176–188, Santa Clara, CA, July 1991.
- [LoP8] J. Larus. Estimating the Potential Parallelism in Programs. In *Proceedings of the Fourth Workshop on Languages and Compilers for Parallel Computing*, pages 331–349, 1991.
- [LoP9] M. S. Lam and R. P. Wilson. Limits of control flow on parallelism. In *Proceedings of the 19th Annual International Symposium on Computer Architecture*, pages 46–57, Queensland, Australia, 1992.
- [LoP10] K. B. Theobald, G. R. Gao, and L. J. Hendren. On the limits of program parallelism and its smoothability. In *Proceedings of the 25th Annual International Symposium on Microarchitecture*, pages 10–19, Portland, OR, 1992.
- [LoP11] J. T. Oplinger, D. L. Heine, and M. S. Lam. In search of speculative thread-level parallelism. In *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques*, pages 303–313, Newport Beach, CA, October 1999.
- [LoP12] D. Stefanovic and M. Martonosi. Limits and graph structure of available instruction-level parallelism (research note). In *Proceedings of the 6th International Euro-Par Conference on Parallel Processing*, pages 1018–1022, Munich, Germany, August, 2000.
- [LoP13] B. Kreaseck, D. Tullsen and B. Calder. Limits of task-based parallelism in irregular applications. *Proceedings of the International Symposium on High Performance Computing*, pages 43–58, October, 2000.
- [LoP14] K. Scott and J. Davidson. Exploring the limits of sub-word level parallelism. In *Proceedings of the Ninth International Conference on Parallel Architectures and Compilation Techniques*, page 81–91, Philadelphia, PA, October, 2000.

- [LoP15] F. Warg and P. Stenström. Limits on speculative module-level parallelism in imperative and object-oriented programs on cmp platforms. In *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques*, pages 221–230, 2001.
- [LoP16] A. Kejariwal, X. Tian, M. Girkar, S. Kozhukhov, H. Saito, U. Banerjee, A. Nicolau, A. V. Veidenbaum, and C. D. Polychronopoulos. On the performance potential of different types of speculative thread-level parallelism. In *Proceedings of the 20th ACM International Conference on Supercomputing*, Cairns, Australia, 2006.

Exploiting ILP in Presence of Conditionals

Code Motion

- [CM1] C. C. Foster and E. M. Riseman. Percolation of code to enhance parallel dispatching and execution. *IEEE Transactions on Computers*, C-21(12):1411–1415, 1972.
- [CM2] J. A. Fisher. Trace Scheduling: A technique for global microcode compaction. *IEEE Transactions on Computers*, C-30(7):478–490, 1981.
- [CM3] B. Su, S. Ding, and J. Xia. An improvement of trace scheduling for global microcode compaction. In *Proceedings of the 17th Workshop on Microprogramming*, New Orleans, LA, October 1984.
- [CM4] A. Nicolau. Percolation scheduling. In *Proceedings of the 1985 International Conference on Parallel Processing*, August 1985.
- [CM5] A. Aiken and A. Nicolau. Perfect pipelining: A new loop parallelization technique. Technical Report 87-869, Department of Computer Science, Cornell University, 1987.
- [CM6] P. P. Chang and W. W. Hwu. Trace selection for compiling large C application programs to microcode. In *Proceedings of the 21st Workshop on Microprogramming*, pages 21–29, San Diego, CA, 1988.
- [CM7] J. A. Fisher. Global code generation for instruction-level parallelism: Trace Scheduling-2. Technical Report HPL-93-43, Hewlett Packard Laboratories, 1993.
- [CM8] W. M. W. Hwu, S. A. Mahlke, W. Y. Chen, P. P. Chang, N. J. Warter, R. A. Bringmann, R. G. Ouellette, R. E. Hank, T. Kiyohara, G. E. Haab, J. G. Holm, and D. M. Lavery. The superblock: An effective technique for VLIW and superscalar compilation. *The Journal of Supercomputing*, 7(1-2):229–248, 1993.
- [CM9] D. I. August. Hyperblock performance optimizations for ILP processors. Master’s thesis, Department of Computer Science, University of Illinois at Urbana-Champaign, 1996.
- [CM10] Q. Jacobson and J. E. Smith. Trace preconstruction. In *Proceedings of the 27th International Symposium on Computer Architecture*, pages 37–46, Vancouver, BC, Canada, 2000.

Branch Prediction

- [BP1] A. Liles Jr. and B. Wilner. Branch prediction mechanism. *IBM Technical Disclosure Bulletin*, 22(7):3013–3016, 1979.

- [BP2] J. E. Smith. A study of branch prediction strategies. In *Proceedings of the Eighth International Symposium on Computer Architecture*, pages 135–148, May 1981.
- [BP3] J. Losq. Generalized history table for branch prediction. *IBM Technical Disclosure Bulletin*, 25(1):99–101, 1982.
- [BP4] J. Lee and A. J. Smith. Branch prediction strategies and branch target buffer design. *IEEE Computer*, 17(1):6–22, 1984.
- [BP5] S. McFarling and J. Hennesey. Reducing the cost of branches. In *Proceedings of the 13th International Symposium on Computer Architecture*, pages 396–403, Tokyo, Japan, 1986.
- [BP6] T.-Y. Yeh and Y. N. Patt. Two-level adaptive training branch prediction. In *Proceedings of the 24th International Symposium of Microarchitecture*, pages 51–61, Albuquerque, NM, 1991.
- [BP7] J. A. Fisher and S. M. Freudenberger. Predicting conditional branch directions from previous runs of a program. In *Proceedings of the Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 85–95, Boston, MA, 1992.
- [BP8] T.-Y. Yeh and Y. N. Patt. Alternative implementations of two-level adaptive branch prediction. In *Proceedings of the 19th International Symposium on Computer Architecture*, pages 124–134, Queensland, Australia, 1992.
- [BP9] S.-T. Pan, K. So, and J. T. Rahmeh. Improving the accuracy of dynamic branch prediction using branch correlation. In *Proceedings of the Fifth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-V)*, pages 76–84, Boston, MA, 1992.
- [BP10] S. McFarling. Combining Branch Predictors. Technical Report TN-36, Western Research Labs, June 1993.
- [BP11] T. Ball and J. Larus. Branch prediction for free. In *Proceedings of the SIGPLAN '93 Conference on Programming Language Design and Implementation*, pages 300–313, Albuquerque, NM, June 1993.
- [BP12] P.-Y. Chang, E. Hao, T.-Y. Yeh, and Y. N. Patt. Branch classification: A new mechanism for improving branch predictor performance. In *Proceedings of the 27th International Symposium of Microarchitecture*, pages 22–31, San Jose, CA, 1994.
- [BP13] B. Calder and D. Grunwald. Fast and accurate instruction fetch and branch prediction. In *Proceedings of the 21th International Symposium on Computer Architecture*, pages 2–11, 1994.
- [BP14] M. Schlansker, V. Kathail, and S. Anik. Height reduction of control recurrences for ILP processors. In *Proceedings of the 27th International Symposium of Microarchitecture*, pages 40–51, San Jose, CA, 1994.
- [BP15] C. Young and M. D. Smith. Improving the accuracy of static branch prediction using branch correlation. *SIGPLAN Notices*, 29(11):232–241, 1994.
- [BP16] P.-Y. Chang, E. Hao, and Y. N. Patt. Alternative implementations of hybrid branch predictors. In *Proceedings of the 28th International Symposium of Microarchitecture*, pages 252–257, Ann Arbor, MI, 1995.

- [BP17] P. P. Chang and U. Banerjee. Profile-guided multi-heuristic branch prediction. In *Proceedings of the International Conference on Parallel Processing*, pages 215–218, Urbana-Champaign, IL, August 1995.
- [BP18] C. Young, N. Gloy, and M. D. Smith. A comparative analysis of schemes for correlated branch prediction. In *Proceedings of the 22nd International Symposium on Computer Architecture*, pages 276–286, S. Margherita Ligure, Italy, 1995.
- [BP19] N. Gloy, M. D. Smith, and C. Young. Performance issues in correlated branch prediction schemes. In *Proceedings of the 28th International Symposium of Microarchitecture*, pages 3–14, Ann Arbor, MI, 1995.
- [BP20] A. Sez nec, S. Jourdan, P. Sainrat, and P. Michaud. Multiple-block ahead branch predictors. In *Proceedings of the Seventh International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-VII)*, pages 116–127, 1996.
- [BP21] S. A. Mahlke, R. E. Hank, J. E. McCormick, D. I. August, and W.-M. W. Hwu. A comparison of full and partial predicated execution support for ILP processors. In *Proceedings of the 22nd International Symposium on Computer Architecture*, pages 138–150, S. Margherita Ligure, Italy, 1995.
- [BP22] M. Evers, P.-Y. Chang, and Y. N. Patt. Using hybrid branch predictors to improve branch prediction accuracy in the presence of context switches. In *Proceedings of the 23th International Symposium on Computer Architecture*, pages 3–11, Philadelphia, PA, 1996.
- [BP23] E. Jacobsen, E. Rotenberg, and J. E. Smith. Assigning confidence to conditional branch predictions. In *Proceedings of the 29th International Symposium of Microarchitecture MICRO-29*, pages 142–152, Paris, France, 1996.
- [BP24] S. A. Mahlke and B. Natarajan. Compiler synthesized dynamic branch prediction. In *Proceedings of the 29th International Symposium of Microarchitecture*, pages 153–164, Paris, France, 1996.
- [BP25] S. Wallace and N. Bagherzadeh. Multiple branch and block prediction. pages 94–105, San Antonio, TX, February 1997.
- [BP26] C.-C. Lee, I-C. K. Chen, and T. N. Mudge. The bi-mode branch predictor. In *Proceedings of the 30th International Symposium of Microarchitecture*, pages 4–13, December 1997.
- [BP27] P. Michaud, A. Sez nec, and R. Uhlig. Trading conflict and capacity aliasing in conditional branch predictors. In *Proceedings of the 24th International Symposium on Computer Architecture*, pages 292–303, 1997.
- [BP28] E. Sprangle, R. S. Chappell, M. Alsup, and Y. N. Patt. The agree predictor: A mechanism for reducing negative branch history interference. *SIGARCH Computer Architecture News*, 25(2):284–291, 1997.
- [BP29] Q. Jacobson, E. Rotenberg, and J. E. Smith. Path-based next trace prediction. In *Proceedings of the 30th International Symposium of Microarchitecture*, pages 14–23, 1997.
- [BP30] M.-D. Tarlescu, K. B. Theobald, and G. R. Gao. Elastic history buffer: A low-cost method to improve branch prediction accuracy. In *Proceedings of the 1997 International Conference on Computer Design*, pages 82–87, Austin, TX, October 1997.
- [BP31] B. Fagin. Partial resolution in branch target buffers. *IEEE Transactions on Computers*, 46(10):1142–1145, 1997.

- [BP32] D. I. August, D. A. Connors, J. C. Gyllenhaal, and W. m. W. Hwu. Architectural support for compiler-synthesized dynamic branch prediction strategies: Rationale and initial results. In *Proceedings of the 3rd IEEE Symposium on High-Performance Computer Architecture*, page 84, 1997.
- [BP33] K. Driesen and U. Hözl. The cascaded predictor: economical and adaptive branch target prediction. In *Proceedings of the 31st Annual ACM/IEEE International Symposium on Microarchitecture*, pages 249–258, Dallas, TX, 1998.
- [BP34] S. Reches and S. Weiss. Implementation and analysis of path history in dynamic branch prediction schemes. *IEEE Transactions on Computers*, 47(8):907–912, 1998.
- [BP35] A. N. Eden and T. Mudge. The yags branch prediction scheme. In *Proceedings of the 31st Annual ACM/IEEE International Symposium on Microarchitecture*, pages 69–77, 1998.
- [BP36] J. Stark, M. Evers, and Y. N. Patt. Variable length path branch prediction. In *Proceedings of the Eighth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-VIII)*, pages 170–179, 1998.
- [BP37] A. Farcy, O. Temam, R. Espasa, and T. Juan. Dataflow analysis of branch mispredictions and its application to early resolution of branch outcomes. In *Proceedings of the 31st Annual ACM/IEEE International Symposium on Microarchitecture*, pages 59–68, 1998.
- [BP38] T. Juan, S. Sanjeevan, and J. J. Navarro. Dynamic history-length fitting: a third level of adaptivity for branch prediction. In *Proceedings of the 25th International Symposium on Computer Architecture*, pages 155–166, 1998.
- [BP39] M. Evers, S. J. Patel, R. S. Chappell, and Y. N. Patt. An analysis of correlation and predictability: what makes two-level branch predictors work. In *Proceedings of the 25th International Symposium on Computer Architecture*, pages 52–61, 1998.
- [BP40] K. Skadron, P. S. Ahuja, M. Martonosi, and D. W. Clark. Improving prediction for procedure returns with return-address-stack repair mechanisms. In *Proceedings of the 31st Annual ACM/IEEE International Symposium on Microarchitecture*, pages 259–271, Dallas, TX, 1998.
- [BP41] Y. Chou, J. Fung, and J. P. Shen. Reducing branch misprediction penalties via dynamic control independence detection. In *Proceedings of the 13th ACM International Conference on Supercomputing*, pages 109–118, Rhodes, Greece, 1999.
- [BP42] C.-M. Chen and C.-T. King. Walk-time address adjustment for improving the accuracy of dynamic branch prediction. *IEEE Transactions on Computers*, 48(5):457–469, 1999.
- [BP43] T. H. Heil, Z. Smith, and J. E. Smith. Improving branch predictors by correlating on data values. In *Proceedings of the 32nd Annual ACM/IEEE International Symposium on Microarchitecture*, pages 28–37, 1999.
- [BP44] S. Mantripragada and A. Nicolau. Using profiling to reduce branch misprediction costs on a dynamically scheduled processor. In *Proceedings of the 14th International Conference on Supercomputing*, pages 206–214, Santa Fe, NM, 2000.
- [BP45] R. Rakvic, B. Black, and J. P. Shen. Completion time multiple branch prediction for enhancing trace cache performance. In *Proceedings of the 27th International Symposium on Computer Architecture*, pages 47–58, 2000.

- [BP46] Q. Jacobson and J. E. Smith. Trace preconstruction. In *Proceedings of the 27th International Symposium on Computer Architecture*, pages 37–46, 2000.
- [BP47] D. Jimenez, H. Hanson, and C. Lin. Boolean formula-based branch prediction for future technologies. In *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques*, pages 97–106, 2001.
- [BP48] D. A. Jiménez and C. Lin. Dynamic branch prediction with perceptrons. In *Proceedings of the Seventh International Symposium on High-Performance Computer Architecture*, pages 197–206, Nuevo Leone, Mexico, January 2001.
- [BP49] C. Zilles and G. Sohi. Execution-based prediction using speculative slices. In *Proceedings of the 28th International Symposium on Computer Architecture*, pages 2–13, Göteborg, Sweden, 2001.
- [BP50] A. Ramirez, J. L. Larriba-Pey, and M. Valero. Branch prediction using profile data. In *Proceedings of the International Euro-Par Conference*, pages 386–393, Manchester, UK, August 2001.
- [BP51] D. A. Jiménez and C. Lin. Neural methods for dynamic branch prediction. *ACM Transactions on Computer Systems*, 20(4):369–397, 2002.
- [BP52] T. Li, L. K. John, A. Sivasubramaniam, N. Vijaykrishnan, and J. Rubio. Understanding and improving operating system effects in control flow prediction. In *Proceedings of the 10th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-X)*, pages 68–80, 2002.
- [BP53] R. S. Chappell, F. Tseng, A. Yoaz, and Y. N. Patt. Difficult-path branch prediction using subordinate microthreads. In *Proceedings of the 29th International Symposium on Computer Architecture*, pages 307–317, Anchorage, Alaska, May 2002.
- [BP54] M. Yang, G.-R. Uh, and D. B. Whalley. Efficient and effective branch reordering using profile data. *ACM Transactions on Programming Languages and Systems*, 24(6):667–697, 2002.
- [BP55] D. A. Jiménez. Fast path-based neural branch prediction. In *Proceedings of the 36th Annual ACM/IEEE International Symposium on Microarchitecture*, pages 243–252, San Diego, CA, December 2003.
- [BP56] D. A. Jiménez. Reconsidering complex branch predictors. In *Proceedings of the Ninth International Symposium on High-Performance Computer Architecture*, pages 43–52, Anaheim, CA, February 2003.
- [BP57] Z. Lu, J. Lach, M. R. Stan, and K. Skadron. Alloyed branch history: combining global and local branch history for robust performance. *International Journal of Parallel Programming*, 31(2):137–177, 2003.
- [BP58] L. N. Vintan, M. Sbera, I. Z. Miha, and A. Florea. An alternative to branch prediction: pre-computed branches. *SIGARCH Computer Architecture News*, 31(3):20–29, 2003.
- [BP59] R. Thomas, M. Franklin, C. Wilkerson, and J. Stark. Improving branch prediction by dynamic dataflow-based identification of correlated branches from a large global history. In *Proceedings of the 30th International Symposium on Computer Architecture*, pages 314–323, 2003.
- [BP60] A. Sez nec. Redundant history skewed perceptron predictors: Pushing limits on global history branch predictors. Technical Report 1554, IRISA, 2003.

- [BP61] H. Akkary, S. T. Srinivasan, and K. Lai. Recycling waste: exploiting wrong-path execution to improve branch prediction. In *Proceedings of the 17th ACM International Conference on Supercomputing*, pages 12–21, 2003.
- [BP62] A. Sezneq. Revisiting the perceptron predictor. Technical Report 1620, IRISA, May 2004.
- [BP63] A. Falcon, J. Stark, A. Ramirez, K. Lai, and M. Valero. Better branch prediction through prophet/critic hybrids. *IEEE Micro*, 25(1):80–89, 2005.
- [BP64] W. C. Krehling, D. Whalley, M. W. Bailey, X. Yuan, G.-R. Uh, and R. van Engelen. Branch elimination by condition merging. *Software—Practice and Experience*, 35(1):51–74, 2005.
- [BP65] D. A. Jimenez. Piecewise linear branch prediction. *SIGARCH Computer Architecture News*, 33(2):382–393, 2005.
- [BP66] H. Kim, O. Mutlu, J. Stark, and Y. N. Patt. Wish branches: Combining conditional branching and predication for adaptive predicated execution. In *Proceedings of the 38th Annual ACM/IEEE International Symposium on Microarchitecture*, pages 43–54, Barcelona, Spain, November 2005.
- [BP67] V. Desmet, H. Veerle, and K. De Bosschere. 2FAR: A 2bcgskew predictor fused by an alloyed redundant history skewed perceptron branch predictor. *Journal of Instruction-Level Parallelism*, 7(4):1–11, 2005.
- [BP68] E. Ipek, S. A. McKee, M. Schulz, and S. Ben-David. Perceptron based branch prediction: Performance of some design options. Technical Report CSL-TR-2005-1043, Cornell Computer Systems Lab, 2005.
- [BP69] D. Tarjan and K. Skadron. Merging path and gshare indexing in perceptron branch prediction. *ACM Transactions on Architecture and Code Optimization*, 2(3):280–300, 2005.
- [BP70] D. A. Jiménez. Improved latency and accuracy for neural branch prediction. *ACM Transactions on Computer Systems*, 23(2):197–218, 2005.
- [BP71] C. Lai, S.-L. Lu, Y. Chen, and T. Chen. Improving branch prediction accuracy with parallel conservative correctors. In *Proceedings of the 2nd Conference on Computing Frontiers*, pages 334–341, 2005.

Predication

- [P1] J. R. Allen, K. Kennedy, C. Porterfield, and J. Warren. Conversion of control dependence to data dependence. In *Conference Record of the Tenth Annual ACM Symposium on the Principles of Programming Languages*, Austin, TX, January 1983.
- [P2] J. Park and M. Schlansker. On predicated execution. Technical Report 58–91, Hewlett Packard Laboratories, 1991.
- [P3] S. A. Mahlke, D. C. Lin, W. Y. Chen, R. E. Hank, and R. A. Bringmann. Effective compiler support for predicated execution using the hyperblock. In *Proceedings of the 25th International Symposium of Microarchitecture*, pages 45–54, 1992.
- [P4] N. J. Warter, S. A. Mahlke, W.-M. W. Hwu, and B. R. Rau. Reverse if-conversion. In *Proceedings of the ACM SIGPLAN 1993 Conference on Programming Language Design and Implementation*, pages 290–299, Albuquerque, NM, 1993.

- [P5] N. J. Warter, D. M. Lavery, and W. W. Hwu. The benefit of predicated execution for software pipelining. In *Proceedings of the 23rd International Conference on System Sciences*, Hawaii, USA, January 1993.
- [P6] S. A. Mahlke, R. E. Hank, R. A. Bringmann, J. C. Gyllenhaal, D. M. Gallagher, and W. W. Hwu. Characterizing the impact of predicated execution on branch prediction. In *Proceedings of the 27th International Symposium of Microarchitecture*, pages 217–227, 1994.
- [P7] G. S. Tyson. The effects of predicated execution on branch prediction. In *Proceedings of the 27th International Symposium of Microarchitecture*, pages 196–206, San Jose, CA, 1994.
- [P8] R. Johnson and M. Schlansker. Analysis techniques for predicated code. In *Proceedings of the 29th International Symposium of Microarchitecture*, pages 100–113, Paris, France, 1996.
- [P9] D. M. Gillies, D. c. Roy Ju, R. Johnson, and M. Schlansker. Global predicate analysis and its application to register allocation. In *Proceedings of the 29th International Symposium of Microarchitecture*, pages 114–125, Paris, France, 1996.
- [P10] D. I. August, W. m. W. Hwu, and S. A. Mahlke. A framework for balancing control flow and predication. In *Proceedings of the 30th International Symposium of Microarchitecture*, pages 92–103, Research Triangle Park, NC, 1997.
- [P11] D. I. August, W.-M. W. Hwu, and S. A. Mahlke. The partial reverse if-conversion framework for balancing control flow and predication. *International Journal of Parallel Programming*, 27(5):381–423, 1999.
- [P12] D. Milicev and Z. Jovanovic. Predicated software pipelining technique for loops with conditions. In *Proceedings of the 12th International Parallel Processing Symposium*, Orlando, FL, 1998.
- [P13] A. Eichenberger, W. Meleis, and S. Maradani. An integrated approach to accelerate data and predicate computations in hyperblocks. In *Proceedings of the 33rd Annual ACM/IEEE International Symposium on Microarchitecture*, pages 101–111, Monterey, CA, 2000.
- [P14] W. Chuang, B. Calder, and J. Ferrante. Phi-predication for light-weight if-conversion. In *Proceedings of the International Symposium on Code Generation and Optimization*, pages 179–190, San Francisco, CA, 2003.
- [P15] M. Smelyanskiy, S. A. Mahlke, E. S. Davidson, and H.-H. S. Lee. Predicate-aware scheduling: a technique for reducing resource constraints. In *Proceedings of the International Symposium on Code Generation and Optimization*, pages 169–178, San Francisco, CA, 2003.

Speculative Execution

Early Work

- [EW1] P. Hudak and R. M. Keller. Garbage collection and task deletion in distributed applicative processing systems. In *Proceedings of the 1982 ACM symposium on LISP and Functional Programming*, pages 168–178, Pittsburgh, PA, 1982.
- [EW2] F. W. Burton. Speculative computation, parallelism and functional programming. *IEEE Transactions on Computers*, 34(12):1190–1193, 1985.

- [EW3] K. Ueda. Guarded horn clauses. In *Proceedings of the 4th Conference Logic Programming*, pages 168–179, 1986.
- [EW4] K. Clark and S. Gregory. PARLOG: Parallel programming in logic. *ACM Transactions on Programming Languages and Systems*, 8(1):1–49, 1986.
- [EW5] J. Miller. Multischeme: A parallel processing system based on mit scheme. Technical Report TR-402, Laboratory for Computer Science, M.I.T., 1987.
- [EW6] R. Soley. On the efficient exploitation of speculation under dataflow paradigms of control. Technical Report TR-443, Laboratory for Computer Science, M.I.T., 1989.
- [EW7] R. Goldman and R. P. Gabriel. Qlisp: Parallel processing in Lisp. *IEEE Software*, 6(4):51–59, 1989.
- [EW8] A. Partridge and A. Dekker. Speculative parallelism in a distributed graph reduction machine. In *Proceedings of the 22nd Annual Hawaii International Conference on System Sciences, Volume 2*, pages 771–779. 1989.
- [EW9] R. B. Osborne. Speculative Computation in Multilisp. In *Proceedings of the 1990 ACM Conference on LISP and Functional Programming*, pages 198–208, Nice, France, June, 1990.
- [EW10] E. Lusk, R. Butler, T. Disz, R. Olson, R. Overbeek, R. Stevens, D. H.D. Warren, A. Calderwood, P. Szeredi, S. Haridi, P. Brand, M. Carlsson, A. Ciepielewski, and B. Hausman. The Aurora or-parallel Prolog system. *New Generation Computing*, 7(2-3):243–271, 1990.

Post 1990's

- [P90-1] K. Kurihara, D. Chaiken, and A. Agarwal. Latency tolerance through multi-threading in large-scale multiprocessors. In *Proceedings of the International Symposium on Shared Memory Multiprocessing*, pages 91–101, Tokyo, Japan, 1991.
- [P90-2] A. K. Uht, V. Sindagi, and K. Hall. Disjoint eager execution: An optimal form of speculative execution. In *Proceedings of the 28th International Symposium of Microarchitecture*, pages 313–325, December 1995.
- [P90-3] J. Greiner and G. E. Blelloch. A provably time-efficient parallel implementation of full speculation. In *Proceedings of the Twenty-third Annual ACM Symposium on the Principles of Programming Languages*, pages 309–321, St. Petersburg Beach, FL, 1996.
- [P90-4] T. H. Heil and J. E. Smith. Selective dual path execution. Technical report, University of Wisconsin - Madison, November 1996.
- [P90-5] M. Lipasti and J. Shen. Approaching 10 IPC via superspeculation. Technical Report CMU-MIG-1, Carnegie Mellon University, 1997.
- [P90-6] M. H. Lipasti and J. P. Shen. Superspeculative microarchitecture for beyond AD 2000. *IEEE Computer*, 30(9):59–66, 1997.
- [P90-7] M. Farrens, T. H. Heil, J. E. Smith, and G. Tyson. Restricted dual path execution. Technical Report CSE-97-18, University of California at Davis, November 1996.
- [P90-8] G. Tyson, K. Lick, and M. Farrens. Limited dual path execution. Technical Report CSE-TR-346-97, University of Michigan at Ann Arbor, 1997.

- [P90-9] P. S. Ahuja, K. Skadron, M. Martonosi, and D. W. Clark. Multipath execution: opportunities and limits. In *Proceedings of the 12th ACM International Conference on Supercomputing*, pages 101–108, Melbourne, Australia, 1998.
- [P90-10] S. Wallace, B. Calder, and D. M. Tullsen. Threaded multiple path execution. In *Proceedings of the 25th International Symposium on Computer Architecture*, pages 238–249, Barcelona, Spain, 1998.
- [P90-11] V. Krishnan and J. Torrellas. Hardware and software support for speculative execution of sequential binaries on a chip-multiprocessor. In *Proceedings of the 12th ACM International Conference on Supercomputing*, pages 85–92, Melbourne, Australia, July, 1998.
- [P90-12] K. Hiraki, J. Tamatsukuri, and T. Matsumoto. Speculative execution model with duplication. In *Proceedings of the 12th ACM International Conference on Supercomputing*, pages 321–328, Melbourne, Australia, 1998.
- [P90-13] T. N. Vijaykumar and G. S. Sohi. Task selection for a multiscalar processor. In *Proceedings of the 31st Annual ACM/IEEE International Symposium on Microarchitecture*, pages 81–92, Dallas, TX, November 1998.
- [P90-14] D. Patel and L. Rauchwerger. Principles of speculative run-time parallelization. In *Proceedings of the 11th International Workshop on Languages and Compilers for Parallel Computing*, pages 323–337, Chapel Hill, NC, August 1998.
- [P90-15] Y. Zhang, L. Rauchwerger, and J. Torrellas. Hardware for speculative run-time parallelization in distributed shared-memory multiprocessors. In *Proceedings of the 4th International Symposium on High-Performance Computer Architecture*, pages 162–173, Las Vegas, NV, January 1998.
- [P90-16] V. Krishnan and J. Torrellas. A Chip-Multiprocessor Architecture with Speculative Multithreading. *IEEE Transactions on Computers*, 48(9):866–880, 1999.
- [P90-17] Y. Zhang, L. Rauchwerger, and J. Torrellas. Hardware for speculative parallelization of partially-parallel loops in DSM multiprocessors. In *Proceedings of the Fifth International Symposium on High-Performance Computer Architecture*, pages 135–141, Orlando, FL, January 1999.
- [P90-18] D. Patel and L. Rauchwerger. Implementation issues of loop-level speculative run-time parallelization. In *Proceedings of 8th International Conference of Compiler Construction*, pages 183–197, Amsterdam, The Netherlands, 1999.
- [P90-19] M. Cintra, J. F. Martínez, and J. Torrellas. Architectural support for scalable speculative parallelization in shared-memory multiprocessors. In *Proceedings of the 27th International Symposium on Computer Architecture*, pages 13–24, Vancouver, BC, Canada, 2000.
- [P90-20] J. G. Steffan, C. B. Colohan, A. Zhai, and T. C. Mowry. A scalable approach to thread-level speculation. In *Proceedings of the 27th International Symposium on Computer Architecture*, pages 1–12, Vancouver, BC, Canada, 2000.
- [P90-21] P. Marcuello and A. González. A quantitative assessment of thread-level speculation techniques. In *Proceedings of the 14th International Parallel and Distributed Processing Symposium*, pages 595–604, Cancun, Mexico, May 2000.

- [P90-22] D. Bruening, S. Devabhaktuni, and S. Amarasinghe. Softspec: Software-based speculative parallelism. In *Proceedings of the 3rd ACM Workshop on Feedback-Directed and Dynamic Optimization*, 2000.
- [P90-23] C.-L. Ooi, S. W. Kim, I. Park, R. Eigenman, B. Falsafi, and T. N. Vijaykumar. Multiplex: Unifying conventional and speculative thread-level parallelism on a chip multiprocessor. In *Proceedings of the 15th ACM International Conference on Supercomputing*, pages 368–380, Sorrento, Italy, 2001.
- [P90-24] M. Prvulovic, M. J. Garzarán, L. Rauchwerger, and J. Torrellas. Removing architectural bottlenecks to the scalability of speculative parallelization. In *Proceedings of the 28th International Symposium on Computer Architecture*, pages 204–215, Göteborg, Sweden, 2001.
- [P90-25] J.-M. Parcerisa and A. Gonzalez. Improving latency tolerance of multithreading through decoupling. *IEEE Transactions on Computers*, 50(10):1084–1094, 2001.
- [P90-26] J. D. Collins, H. Wang, D. M. Tullsen, C. J. Hughes, Y.-F. Lee, D. M. Lavery, and J. P. Shen. Speculative precomputation: long-range prefetching of delinquent loads. In *Proceedings of the 28th International Symposium on Computer Architecture*, pages 14–25, Göteborg, Sweden, 2001.
- [P90-27] J. D. Collins, D. M. Tullsen, H. Wang, and J. P. Shen. Dynamic speculative precomputation. In *Proceedings of the 34th Annual ACM/IEEE International Symposium on Microarchitecture*, pages 306–317, Austin, TX, December 2001.
- [P90-28] J. L. Aragón, J. González, A. González, and J. E. Smith. Dual path instruction processing. In *Proceedings of the 16th ACM International Conference on Supercomputing*, pages 220–229, New York, NY, 2002.
- [P90-29] M. Agarwal, S. K. Nandy, J. v. Eijndhoven, and S. Balakrishnan. Speculative trace scheduling in VLIW processors. In *Proceedings of the 20th International Conference on Computer Design*, Freiburg, Germany, 2002.
- [P90-30] P. Marcuello and A. González. Thread spawning schemes for speculative multithreading. In *Proceedings of the Eighth International Symposium on High-Performance Computer Architecture*, pages 55–64, Boston, MA, February 2002.
- [P90-31] M. J. Garzarán, M. Prvulovic, V. Viñals, J. M. Llabería, L. Rauchwerger and J. Torrellas. Using Software Logging to Support Multi-Version Buffering in Thread-Level Speculation. In *Proceedings of the 12th International Conference on Parallel Architectures and Compilation Techniques*, pages 170–181, New Orleans, LA, September, 2003.
- [P90-32] P. S. Chen, M. Y. Hung, Y. S. Hwang, R. D. Ju, and J. K. Lee. Compiler support for speculative multithreading architecture with probabilistic point-to analysis. In *Proceedings of the 9th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 25–36, San Diego, CA, 2003.
- [P90-33] F. Warg and P. Stenström. Improving speculative thread-level parallelism through module run-length prediction. In *Proceedings of the International Parallel and Distributed Processing Symposium*, pages 12–230, Nice, France, 2003.
- [P90-34] M. Cintra and D. R. Llanos. Toward efficient and robust software speculative parallelization on multiprocessors. In *Proceedings of the 9th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 13–24, San Diego, CA, 2003.

- [P90-35] T. A. Johnson, R. Eigenmann, and T. N. Vijaykumar. Min-cut program decomposition for thread-level speculation. In *Proceedings of the SIGPLAN '04 Conference on Programming Language Design and Implementation*, pages 59–70, Washington DC, USA, 2004.
- [P90-36] A. Bhowmik and M. Franklin. A general compiler framework for speculative multithreaded processors. *IEEE Transactions on Parallel and Distributed Systems*, 15(8):713–724, 2004.
- [P90-37] J. Whaley and C. Kozyrakis. Heuristics for profile-driven method-level speculative parallelization. In *Proceedings of the 33rd International Conference on Parallel Processing*, Oslo, Norway, June 2005.
- [P90-38] J. Renau, J. Tuck, W. Liu, L. Ceze, K. Strauss and J. Torrellas. Tasking with out-of-order spawn in TLS chip multiprocessors: Microarchitecture and compilation. In *Proceedings of the 19th ACM International Conference on Supercomputing*, pages 179–188, Cambridge, MA, June, 2005.
- [P90-39] C. G. Quiñones, C. Madriles, J. Sánchez, P. Marcuello, A. González, and D. M. Tullsen. Mitosis compiler: An infrastructure for speculative threading based on pre-computation slices. In *Proceedings of the SIGPLAN '05 Conference on Programming Language Design and Implementation*, pages 269–279, Chicago, IL, 2005.
- [P90-40] M. J. Garzarán, M. Prvulovic, M. Llabería J, V. Viñals, L. Rauchwerger, and J. Torrellas. Tradeoffs in buffering speculative memory state for thread-level speculation in multiprocessors. *ACM Transactions on Architecture and Code Optimization*, 2(3):247–279, 2005.
- [P90-41] W. Liu and J. Tuck and L. Ceze and W. Ahn and K. Strauss and J. Renau and J. Torellas. POSH: A TLS compiler that exploits program structure. In *Proceedings of the 11th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, March, 2006.

Control Speculation

- [CS1] D. N. Pnevmatikatos, M. Franklin, and G. S. Sohi. Control flow prediction for dynamic ILP processors. In *Proceedings of the 26th International Symposium of Microarchitecture*, pages 153–163, Austin, TX, 1993.
- [CS2] N. J. Warter, P. P. Chang, S. A. Mahlke, W. Y. Chen, and W.-M. W. Hwu. Three architectural models for compiler-controlled speculative execution. *IEEE Transactions on Computers*, 44(4):481–494, 1995.
- [CS3] Q. Jacobson, S. Bennett, N. Sharma, and J. E. Smith. Control flow speculation in multiscalar processors. In *Proceedings of the 3rd International Symposium on High-Performance Computer Architecture*, pages 218–229, San Antonio, TX, February 1997.
- [CS4] J. Tubella and A. González. Control speculation in multithreaded processors through dynamic loop detection. In *Proceedings of the 4th International Symposium on High-Performance Computer Architecture*, pages 14–23, Las Vegas, NV, February 1998.
- [CS5] D. Grunwald, A. Klauser, S. Manne, and A. Pleszkun. Confidence estimation for speculation control. In *Proceedings of the 25th International Symposium on Computer Architecture*, pages 122–131, Barcelona, Spain, 1998.

- [CS6] K. Luo, M. Franklin, S. S. Mukherjee, and A. Seznev. Boosting smt performance by speculation control. In *Proceedings of the 15th International Parallel & Distributed Processing Symposium*, page 2, San Francisco, CA, 2001.
- [CS7] J. González and A. González. Control-flow speculation through value prediction. *IEEE Transactions on Computers*, 50(12):1362–1376, 2001.

Data Dependence Speculation

- [DDS1] A. Nicolau. Run-time disambiguation: coping with statically unpredictable dependencies. *IEEE Transactions on Computers*, 38(5):633–678, 1989.
- [DDS2] A. S. Huang, G. Slavenburg, and J. P. Shen. Speculative disambiguation: A compilation technique for dynamic memory disambiguation. In *Proceedings of the 21st Annual International Symposium on Computer Architecture*, pages 200–210, Chicago, IL, 1994.
- [DDS3] D. M. Gallagher, W. Y. Chen, S. A. Mahlke, J. C. Gyllenhaal, and W. m. W. Hwu. Dynamic memory disambiguation using the memory conflict buffer. In *Proceedings of the Sixth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 183–193, San Jose, CA, 1994.
- [DDS4] M. Franklin. Multi-version caches for multiscalar processors. In *Proceedings of 1st International Conference on High Performance Computing*, 1995.
- [DDS5] L. Rauchwerger and D. Padua. Parallelizing while loops for multiprocessor systems. In *Proceedings of the 9th International Parallel Processing Symposium*, pages 347–356, Santa Barbara, CA, April, 1995.
- [DDS6] L. Rauchwerger, N. M. Amato and D. Padua. Run-Time Methods for Parallelizing Partially Parallel Loops. In *Proceedings of the 9th ACM International Conference on Supercomputing*, pages 137–146, Barcelona, Spain, July, 1995.
- [DDS7] L. Rauchwerger and D. Padua. The LRPD test: Speculative run-time parallelization of loops with privatization and reduction parallelization. In *Proceedings of the ACM SIGPLAN 1995 Conference on Programming Language Design and Implementation*, pages 218–232, La Jolla, CA, 1995.
- [DDS8] M. Franklin and G. S. Sohi. ARB: A hardware mechanism for dynamic reordering of memory references. *IEEE Transactions on Computers*, 45(5):552–571, 1996.
- [DDS9] Y. Sazeides, S. Vassiliadis, and J. E. Smith. The performance potential of data dependence speculation & collapsing. In *Proceedings of the 29th Annual ACM/IEEE International Symposium on Microarchitecture*, pages 238–247, Paris, France, 1996.
- [DDS10] A. Moshovos, S. E. Breach, T. N. Vijaykumar, and G. S. Sohi. Dynamic speculation and synchronization of data dependences. In *Proceedings of the 24th Annual International Symposium on Computer Architecture*, pages 181–193, Denver, CO, 1997.
- [DDS11] B.-C. Cheng, D. A. Connors, and W.-M. W. Hwu. Compiler-directed early load-address generation. In *Proceedings of the 31st annual ACM/IEEE International Symposium on Microarchitecture*, pages 138–147, Dallas, TX, 1998.

- [DDS12] G. Reinman and B. Calder. Predictive techniques for aggressive load speculation. In *Proceedings of the 31st annual ACM/IEEE International Symposium on Microarchitecture*, pages 127–137, Dallas, TX, 1998.
- [DDS13] T. Sato. Data dependence speculation using data address prediction and its enhancement with instruction reissue. In *Proceedings of the 24th Euromicro Conference*, pages 285–292, August 1998.
- [DDS14] G. Z. Chrysos and J. S. Emer. Memory dependence prediction using store sets. In *Proceedings of the 25th International Symposium on Computer Architecture*, pages 142–153, Barcelona, Spain, 1998.
- [DDS15] S. Onder and R. Gupta. Dynamic memory disambiguation in the presence of out-of-order store issuing. In *Proceedings of the 32nd Annual ACM/IEEE International Symposium on Microarchitecture*, pages 170–176, Haifa, Israel, 1999.
- [DDS16] A. Moshovos and G. S. Sohi. Read-after-read memory dependence prediction. In *Proceedings of the 32nd Annual ACM/IEEE International Symposium on Microarchitecture*, pages 177–185, Haifa, Israel, 1999.
- [DDS17] A. Roth, A. Moshovos, and G. S. Sohi. Improving virtual function call target prediction via dependence-based pre-computation. In *Proceedings of the 13th ACM International Conference on Supercomputing*, pages 356–364, Rhodes, Greece, 1999.
- [DDS18] M. Bekerman, A. Yoaz, F. Gabbay, S. Jourdan, M. Kalaev, and R. Ronen. Early load address resolution via register tracking. In *Proceedings of the 27th International Symposium on Computer Architecture*, pages 306–315, Vancouver, BC, 2000.
- [DDS19] B.-K. Chung, J. Zhang, J.-K. Peir, S.-C. Lai, and K. Lai. Direct load: dependence-linked dataflow resolution of load address and cache coordinate. In *Proceedings of the 34th Annual ACM/IEEE International Symposium on Microarchitecture*, pages 76–87, Austin, Texas, 2001.
- [DDS20] T. N. Vijaykumar, S. Gopal, J. E. Smith, and G. Sohi. Speculative versioning cache. *IEEE Transactions on Parallel and Distributed Systems*, 12(12):1305–1317, 2001.
- [DDS21] D. Ortega, E. Ayguadé, and M. Valero. Dynamic memory instruction bypassing. In *Proceedings of the 17th ACM International Conference on Supercomputing*, pages 316–325, San Francisco, CA, 2003.

Data Value Speculation

- [DVS1] F. Gabbay and A. Mendelson. Speculative execution based on value prediction. Technical Report EE Department TR # 1080, Technion–Israel Institute of Technology, November 1996.
- [DVS2] M. H. Lipasti, C. B. Wilkerson, and J. P. Shen. Value locality and load value prediction. *SIGPLAN Notices*, 31(9):138–147, 1996.
- [DVS3] F. Gabbay and A. Mendelson. Can program profiling support value prediction? In *Proceedings of the 30th International Symposium of Microarchitecture*, pages 270–280, Research Triangle Park, North Carolina, United States, 1997.
- [DVS4] M. H. Lipasti and J. P. Shen. The performance potential of value and dependence prediction. In *Proceedings of the Third International Euro-Par Conference on Parallel Processing*, pages 1043–1052, Passau, Germany, August 1997.

- [DVS5] Y. Sazeides and J. E. Smith. The predictability of data values. In *Proceedings of the 30th Annual ACM/IEEE International Symposium on Microarchitecture*, pages 248–258, Research Triangle Park, NC, December 1997.
- [DVS6] K. Wang and M. Franklin. Highly accurate data value prediction using hybrid predictors. In *Proceedings of the 30th Annual ACM/IEEE International Symposium on Microarchitecture*, pages 281–290, Research Triangle Park, NC, December 1997.
- [DVS7] F. Gabbay and A. Mendelson. Using value prediction to increase the power of speculative execution hardware. *ACM Transactions on Computer Systems*, 16(3):234–270, 1998.
- [DVS8] S. Jourdan, R. Ronen, M. Bekerman, B. Shomar, and A. Yoaz. A novel renaming scheme to exploit value temporal locality through physical register reuse and unification. In *Proceedings of the 31st Annual ACM/IEEE International Symposium on Microarchitecture*, pages 216–225, Dallas, TX, 1998.
- [DVS9] J. González and A. González. The potential of data value speculation to boost ILP. In *Proceedings of the 12th International Conference on Supercomputing*, pages 21–28, Melbourne, Australia, 1998.
- [DVS10] J. Steffan and T. Mowry. The potential for using thread-level data speculation to facilitate automatic parallelization. In *Proceedings of the 4th International Symposium on High-Performance Computer Architecture*, pages 2–13, February 1998.
- [DVS11] C.-Y. Fu, M. D. Jennings, S. Y. Larin, and T. M. Conte. Value speculation scheduling for high performance processors. In *Proceedings of the Eighth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 262–271, San Jose, CA, 1998.
- [DVS12] B. Calder, G. Reinman, and D. M. Tullsen. Selective value prediction. In *Proceedings of the 26th International Symposium on Computer Architecture*, pages 64–74, Atlanta, Georgia, 1999.
- [DVS13] T. Nakra, R. Gupta, and M. L. Soffa. Global context-based value prediction. In *Proceedings of the Fifth International Symposium on High-Performance Computer Architecture*, pages 4–12, Orlando, FL, January 1999.
- [DVS14] M. Burtscher and B. G. Zorn. Exploring last n value prediction. pages 66–76, Newport Beach, CA, October 1999.
- [DVS15] J. Huang and D. Lilja. Exploiting basic block value locality with block reuse. In *Proceedings of the Fifth International Symposium on High-Performance Computer Architecture*, pages 106–114, Orlando, FL, January 1999.
- [DVS16] Y. Sazeides and J. E. Smith. Limits of data value predictability. *International Journal of Parallel Programming*, 27(4):229–256, 1999.
- [DVS17] P. Marcuello, J. Tubella, and A. González. Value prediction for speculative multithreaded architectures. In *Proceedings of the 32nd Annual ACM/IEEE International Symposium on Microarchitecture*, pages 230–236, Haifa, Israel, 1999.
- [DVS18] E. Larson and T. Austin. Compiler controlled value prediction using branch predictor based confidence. In *Proceedings of the 33rd Annual ACM/IEEE International Symposium on Microarchitecture*, pages 327–336, Monterey, CA, 2000.

- [DVS19] K. M. Lepak, G. B. Bell, and M. H. Lipasti. Silent stores and store value locality. *IEEE Transactions on Computers*, 50(11):1174–1190, 2001.
- [DVS20] Y. Wu, D.-Y. Chen, and J. Fang. Better exploration of region-level value locality with integrated computation reuse and value prediction. In *Proceedings of the 28th International Symposium on Computer Architecture*, pages 98–108, Göteborg, Sweden, 2001.
- [DVS21] C.-H. Liao and J.-J. Shieh. Exploiting speculative value reuse using value prediction. In *Proceedings of the Seventh Asia-Pacific Conference on Computer Systems Architecture*, pages 101–108, Melbourne, Victoria, Australia, 2002.
- [DVS22] M. Burtscher and B. G. Zorn. Hybrid load-value predictors. *IEEE Transactions on Computers*, 51(7):759–774, 2002.
- [DVS23] C. Y. Fu, J. T. Bodine, and T. M. Conte. Modeling value speculation: An optimal edge selection problem. *IEEE Transactions on Computers*, 52(3):277–292, 2003.
- [DVS24] H. Zhou and T. M. Conte. Enhancing memory level parallelism via recovery-free value prediction. In *Proceedings of the 17th ACM International Conference on Supercomputing*, pages 326–335, San Francisco, CA, 2003.
- [DVS25] H. Zhou, J. Flanagan, and T. M. Conte. Detecting global stride locality in value streams. In *Proceedings of the 30th International Symposium on Computer Architecture*, pages 324–335, San Diego, CA, 2003.
- [DVS26] N. Tuck and D. M. Tullsen. Multithreaded value prediction. In *Proceedings of the 11th International Symposium on High-Performance Computer Architecture*, pages 5–15, San Francisco, CA, February 2005.

Multithreaded Processors

- [MP1] R. H. Halstead Jr. and T. Fujita. MASA: A multithreaded processor architecture for parallel symbolic computing. In *Proceedings of the 15th Annual International Symposium on Computer Architecture*, pages 443–451, Honolulu, Hawaii, 1988.
- [MP2] R. S.-Barrera, D. Culler, and T. von Eicken. Analysis of multithreaded architectures for parallel computing. In *Proceedings of the Second Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 169–178, Island of Crete, Greece, 1990.
- [MP3] R. S. Nikhil, G. M. Papadopoulos, and Arvind. T: A multithreaded massively parallel architecture. In *Proceedings of the 19th Annual International Symposium on Computer Architecture*, pages 156–167, Queensland, Australia, 1992.
- [MP4] M. Franklin and G. S. Sohi. The expandable split window paradigm for exploiting fine-grain parallelism. In *Proceedings of the 19th International Symposium on Computer Architecture*, pages 58–67, Gold Coast, Australia, May 1992.
- [MP5] G. S. Sohi, S. Breach, and T. N. Vijaykumar. Multiscalar processors. In *Proceedings of the 22nd International Symposium on Computer Architecture*, pages 414–425, Ligure, Italy, 1995.

- [MP6] P. Dubey, K. O'Brien, K. M. O'Brien, and C. Barton. Single-program speculative multithreading (spsm) architecture: Compiler assisted fine-grained multithreading. In *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques*, pages 109–121, 1995.
- [MP7] J. Y. Tsai and P. C. Yew. The superthreaded architecture: Thread pipelining with run-time data dependence checking and control speculation. In *Proceedings of International Conference on Parallel Architecture and Compiler Techniques (PACT'96)*, pages 35–46, 1996.
- [MP8] H. Akkary and M. A. Driscoll. A dynamic multithreading processor. In *Proceedings of the 31st Annual ACM/IEEE International Symposium on Microarchitecture*, pages 226–236, Dallas, TX, November, 1998.
- [MP9] P. Marcuello, A. González, and J. Tubella. Speculative multithreaded processors. In *Proceedings of the 12th International Conference on Supercomputing*, pages 77–84, Melbourne, Australia, 1998.
- [MP10] P. Marcuello and A. González. Clustered speculative multithreaded processors. In *Proceedings of the 13th International Conference on Supercomputing*, pages 365–372, 1999.
- [MP11] L. Barroso, K. Gharachorloo, R. McNamara, A. Nowatzky, S. Qadeer, B. Sano, S. Smith, R. Stets, and B. Verghese. Piranha: A scalable architecture based on single-chip multiprocessing. In *Proceedings of the 27th International Symposium on Computer Architecture*, pages 12–14, Vancouver, Canada, 2000.
- [MP12] M. Zahran and M. Franklin. Hierarchical multithreading for exploiting parallelism at multiple granularities. In *Proceedings of 5th Workshop on Multithreaded Execution Architecture and Compilation*, pages 35–42, 2001.
- [MP13] T. Ungerer, B. Robič, and J. Šilc. A survey of processors with explicit multithreading. *ACM Computing Surveys*, 35(1):29–63, 2003.
- [MP14] D. Collins, , and D. M. Tullsen. Clustered multithreaded architectures pursuing both IPC and Cycle time. In *Proceedings of the 18th International Parallel and Distributed Processing Symposium*, 2004.
- [MP15] T. Ohsawa, M. Takagi, S. Kawahara, and S. Matsushita. Pinot: Speculative multi-threading processor architecture exploiting parallelism over a wide range of granularities. In *Proceedings of the 38th Annual ACM/IEEE International Symposium on Microarchitecture*, pages 81–92, Barcelona, Spain, November 2005.

Micro-architectural Support for Multithreading

- [MSM1] D. E. Culler, A. Sah, K. E. Schauser, T. von Eicken, and J. Wawrzynek. Fine-grain parallelism with minimal hardware support: A compiler-controlled threaded abstract machine. In *Proceedings of the Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 164–175, Santa Clara, CA, 1991.
- [MSM2] B. Boothe and A. Ranade. Improved multithreading techniques for hiding communication latency in multiprocessors. In *Proceedings of the 19th Annual International Symposium on Computer Architecture*, pages 214–223, Queensland, Australia, 1992.

- [MSM3] D. M. Tullsen, S. J. Eggers, J. S. Emer, H. M. Levy, J. L. Lo, and R. L. Stamm. Exploiting choice: Instruction fetch and issue on an implementable simultaneous multithreading processor. In *Proceedings of the 23rd annual International Symposium on Computer Architecture*, pages 191–202, Philadelphia, PA, 1996.
- [MSM4] D. M. Tullsen, J. L. Lo, S. J. Eggers, and H. M. Levy. Supporting fine-grain synchronization on a simultaneous multithreaded processor. In *Proceedings of the 5th International Symposium on High Performance Computer Architecture*, January 1999.
- [MSM5] R. Rajwar and J. R. Goodman. Speculative lock elision: Enabling highly concurrent multithreaded execution. In *Proceedings of the 34th Annual ACM/IEEE International Symposium on Microarchitecture*, pages 294–305, Austin, Texas, 2001.
- [MSM6] R. S. Chappell, F. Tseng, A. Yoaz, and Y. N. Patt. Microarchitectural support for precomputation microthreads. In *Proceedings of the 35th Annual ACM/IEEE International Symposium on Microarchitecture*, pages 74–84, 2002.

Performance Study

- [PS1] W.-D. Weber and A. Gupta. Exploring the benefits of multiple hardware contexts in a multiprocessor architecture: Preliminary results. In *Proceedings of the 16th annual international symposium on Computer Architecture*, pages 273–280, Jerusalem, Israel, 1989.
- [PS2] A. Gupta, J. Hennessy, K. Gharachorloo, T. Mowry, and W.-D. Weber. Comparative evaluation of latency reducing and tolerating techniques. In *Proceedings of the 18th Annual International Symposium on Computer Architecture*, pages 254–263, Toronto, Ontario, Canada, 1991.
- [PS3] R. Thekkath and S. J. Eggers. The effectiveness of multiple hardware contexts. In *Proceedings of the sixth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 328–337, San Jose, CA, 1994.
- [PS4] R. Thekkath and S. J. Eggers. Impact of sharing-based thread placement on multithreaded architectures. In *Proceedings of the 21st Annual International Symposium on Computer Architecture*, pages 176–186, Chicago, IL, 1994.
- [PS5] X.-M. Tian, S. Nemawarkar, G. R. Gao, and H. Hum. Data locality sensitivity of multithreaded computations on a distributed-memory multiprocessor. In *Proceedings of the 1996 conference of the Centre for Advanced Studies on Collaborative Research*, page 37, Toronto, Ontario, Canada, 1996.
- [PS6] A. Sohn, M. Sato, N. Yoo, and J.-L. Gaudiot. Effects of multithreading on data and workload distribution for distributed-memory multiprocessors. In *Proceedings of the 10th International Parallel Processing Symposium*, pages 116–122, 1996.
- [PS7] J.-Y. Tsai, Z. Jiang, E. Ness, and P.-C. Yew. Performance study of a concurrent multithreaded processor. In *Proceedings of the the Fourth International Symposium on High-Performance Computer Architecture*, page 24, 1998.
- [PS8] J. L. Lo, L. A. Barroso, S. J. Eggers, K. Gharachorloo, H. M. Levy, and S. S. Parekh. An analysis of database workload performance on simultaneous multithreaded processors. In *Proceedings of the 25th Annual International Symposium on Computer Architecture*, pages 39–50, Barcelona, Spain, 1998.

- [PS9] H. Kwak, B. Lee, A. R. Hurson, S.-H. Yoon, and W.-J. Hahn. Effects of multithreading on cache performance. *IEEE Transactions on Computers*, 48(2):176–184, 1999.
- [PS10] W. M. Zuberek. Analysis of performance bottlenecks in multithreaded multiprocessor systems. *Fundamenta Informaticae*, 50(2):223–241, 2002.
- [PS11] G. E. Blelloch and P. B. Gibbons. Effectively sharing a cache among threads. In *Proceedings of the 16th Annual ACM Symposium on Parallelism in Algorithms and Architectures*, pages 235–244, Barcelona, Spain, 2004.