

- General Chair's Message p. vii
- Program Chair's Message p. viii
- Conference Organization p. ix
- Reviewers p. x
- Plenary Session
- Opening Remarks
- Keynote Speech: Greg Papadopoulos, CTO, Sun Microsystems Inc., USA
- Session 1 Multithreading and Speculation
  - Execution-based Prediction Using Speculative Slices C. Zilles and G. Sohi p. 2
  - Speculative Precomputation: Long-range Prefetching of Delinquent Loads J. Collins and H. Wang and D. Tullsen and C. Hughes and Y. Lee and D. Lavery and J. Shen p. 14
  - Dynamically Allocating Processor Resources between Nearby and Distant ILP R. Balasubramonian and S. Dwarkadas and D. Albonesi p. 26
- Session 2 Memory System Issues
  - Tolerating Memory Latency through Software-Controlled Pre-Execution in Simultaneous Multithreading Processors C. Luk p. 40
  - Data Prefetching by Dependence Graph Precomputation M. Annavaram and J. Patel and E. Davidson p. 52
  - Concurrency, Latency, or System Overhead: Which Has the Largest Impact on Uniprocessor DRAM-System Performance? V. Cuppu and B. Jacob p. 62
- Session 3 Processor Architecture
  - Focusing Processor Policies via Critical-Path Prediction B. Fields and S. Rubin and R. Bodik p. 74
  - Automated Design of Finite State Machine Predictors for Customized Processors T. Sherwood and B. Calder p. 86
  - Better Exploration of Region-Level Value Locality with Integrated Computation Reuse and Value Prediction Y. Wu and D. Chen and J. Fang p. 98
- Session 4 Communication Support
  - CryptoManiac: A Fast Flexible Architecture for Secure Communication L. Wu and C. Weaver and T. Austin p. 110
  - QoS Provisioning in Clusters: An Investigation of Router and NIC Design K. Yum and E. Kim and C. Das p. 120
- Session 5 Cache Management
  - Locality vs. Criticality S. Srinivasan and R. Ju and A. Lebeck and C. Wilkerson p. 132
  - Dead-Block Prediction and Dead-Block Correlating Prefetchers A. Lai and C. Fide and B. Falsafi p. 144
  - Code Layout Optimizations for Transaction Processing Workloads A. Ramirez and L. Barroso and K. Gharachorloo and R. Cohn and J. Larriba-Pey and P. Lowney and M. Valero p. 155
- Session 6A Architectural Impact of Emerging Technologies
  - Exploring and Exploiting Wire-Level Pipelining in Emerging Technologies M. Niemier and P. Kogge p. 166
  - NanoFabrics: Spatial Computing Using Molecular Electronics S. Goldstein and M. Budiu p. 178

- Session 6B Shared-Memory Multiprocessors
- A Simple Method for Extracting Models from Protocol Code D. Lie and A. Chou and D. Engler and D. Dill p. 192
- Removing Architectural Bottlenecks to the Scalability of Speculative Parallelization M. Prvulovic and M. Garzaran and L. Rauchwerger and J. Torrellas p. 204
- Session 7 Energy-Effective Designs
- Power and Energy Reduction Via Pipeline Balancing R. Bahar and S. Manne p. 218
- Energy-Effective Issue Logic D. Folegnani and A. Gonzalez p. 230
- Cache Decay: Exploiting Generational Behavior to Reduce Cache Leakage Power S. Kaxiras and Z. Hu and M. Martonosi p. 240
- Session 8 Performance Tools and Evaluations
- Variability in the Execution of Multimedia Applications and Implications for Architecture C. Hughes and P. Kaul and S. Adve and R. Jain and C. Park and J. Srinivasan p. 254
- Measuring Experimental Error in Microprocessor Simulation R. Desikan and D. Burger and S. Keckler p. 266
- Rapid Profiling via Stratified Sampling S. Sastry and R. Bodik and J. Smith p. 278
- Author Index p. 291