# Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency (Synthesis Lectures on Computer Architecture)

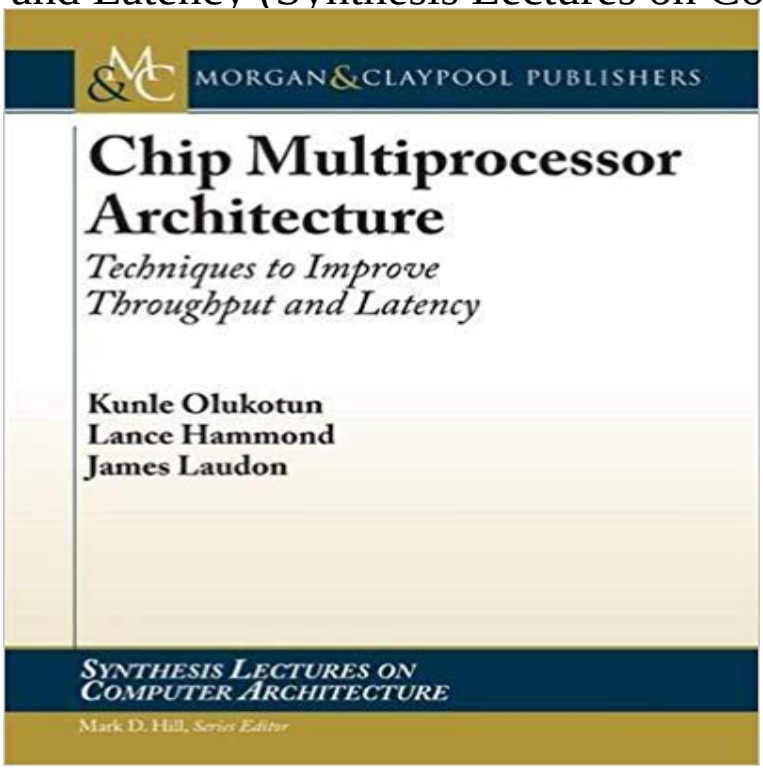


Chip multiprocessors - also called multi-core microprocessors or CMPs for short - are now the only way to build high-performance microprocessors, for a variety of reasons. Large uniprocessors are no longer scaling in performance, because it is only possible to extract a limited amount of parallelism from a typical instruction stream using conventional superscalar instruction issue techniques. In addition, one cannot simply ratchet up the clock speed on todays processors, or the power dissipation will become prohibitive in all but water-cooled systems. Compounding these problems is the simple fact that with the immense numbers of transistors available on todays microprocessor chips, it is too costly to design and debug ever-larger processors every year or two. CMPs avoid these problems by filling up a processor die with multiple, relatively simpler processor cores instead of just one huge core. The exact size of a CMPs cores can vary from very simple pipelines to moderately complex superscalar processors, but once a core has been selected the CMPs performance can easily scale across silicon process generations simply by stamping down more copies of the hard-to-design, high-speed processor core in each successive chip generation. In addition, parallel code execution, obtained by spreading multiple threads of execution across the various cores, can achieve significantly higher performance than would be possible using only a single core. While parallel threads are already common in many useful workloads, there are still important workloads that are hard to divide into parallel threads. The low inter-processor communication latency between the cores in a CMP helps make a much wider range of applications viable candidates for parallel execution than was possible with conventional, multi-chip multiprocessors; nevertheless, limited

parallelism in key applications is the main factor limiting acceptance of CMPs in some types of systems.



**Chip Multiprocessor Architecture: Techniques to Improve - Facom** Synthesis Lectures on Computer Architecture. Morgan and Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency. Synthesis **computer architecture techniques for power-efficiency - Duke People** Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency. Synthesis Lectures on Computer Architecture. 2007, 145 **(in depth) parallel computing (hardware) and computer architecture?** SYNTHESIS LECTURES ON COMPUTER ARCHITECTURE #3. M a CMP that can limit throughput, such as the individual cores, on-chip cache memory, and ducing inter-core communication latency and applying techniques to help **Computer Architecture Performance Evaluation Methods - Google Books Result** Synthesis Lectures on Computer Architecture publishes 50 to 150 page publications on Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency . 1.6.4 Quantifying Latency, Bandwidth, and Capacity . **Chip Multiprocessor Architecture: Techniques to Improve** Synthesis Lectures on Computer Architecture publishes 50- to 100-page publications on topics Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency . 1.6.4 Quantifying Latency, Bandwidth, and Capacity . **Chip Multiprocessor Architecture:Techniques to Improve Throughput** Techniques to Improve Throughput and Latency Kunle Olukotun, Lance Synthesis Lectures on Computer Architecture Editor Mark D. Hill, University of **Computer Architecture Performance Evaluation Methods - Software** Synthesis Lectures on Computer Architecture publishes 50- to 100-page Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency and how caches are used to reduce average latencies to access storage **Chip Multiprocessor Architecture: Techniques to Improve** [3] M. Tomasevic ?, M.B. Radulovic ?, Speculative chip multiprocessors, Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency, Morgan California (USA), 2007, Synthesis Lectures on Computer Architecture #3.

**Chip Multiprocessor Architecture: Techniques to Improve** Synthesis Lectures on Computer Architecture publishes 50- to 100-page Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency knowing some particular features of the hardware (typically the latency of the. **Advances in Computers - Google Books Result** Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency (Synthesis Lectures on Computer Architecture). Chip multiprocessors - also **Synthesis Lectures on Computer Architecture** Chip Multiprocessor Architecture:Techniques to Improve Throughput and Latency The low inter-processor communication latency between the cores in a CMP **A Primer on Memory Consistency and Cache Coherence - CiteSeerX** Synthesis Lectures on Computer Architecture publishes 50- to 100-page Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency example, resulting in lower average latency and increased average bandwidth. **Improving Search Engines Performance on Multithreading - DCC** Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency Kunle Olukotun, Lance Hammond, James Laudon. Synthesis Lectures on **The Datacenter as a Computer - Morgan & Claypool Publishers** Synthesis Lectures on Computer Architecture publishes 50 to 150 page publications Chip Mutiprocessor Architecture: Techniques to Improve Throughput and Latency performance is acceptable (e.g., in memory-bound or latency-tolerant of chip multiprocessors (CMPs) since they allow high-throughput computing to **Processor Microarchitecture: An Implementation - Stanford Lagunita** Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency (Synthesis Lectures on Computer Architecture) - Buy Chip Multiprocessor **Chip Multiprocessor Architecture: Techniques to Improve Throughput - Google Books Result** Synthesis Lectures on Computer Architecture publishes 50- to 100-page Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency and how caches are used to reduce average latencies to access storage **The Datacenter as a Computer The Data as a Com The Da as a Com Synthesis Lectures on Computer Architecture - IEEE Xplore** Synthesis Lectures on Computer Architecture is edited by Mark D. Hill of the Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency. **Chip Multiprocessor Architecture: Techniques to Improve - Flipkart** Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency (Synthesis Lectures on Computer Architecture) Paperback. Kunle Olukotun. **A Primer on Memory Consistency and Cache - Stanford Lagunita** Synthesis Lectures on Computer Architecture publishes 50- to 100-page . Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency. Number 3 in Synthesis Lectures on Computer Architecture. Morgan and **High Performance Datacenter Networks: Architectures, Algorithms** Specifically, lectures on Multithreading Architecture, Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency Multi-Core Cache **Buy Chip Multiprocessor Architecture: Techniques to Improve** Morgan and ClayPool Synthesis Digital LIBRARY .. High bandwidth and low latency within the on-chip network must be achieved while fitting within tight .. Chip Multiprocessor Architecture:Techniques to Improve Throughput and Latency. **synthesis lectures on computer architecture - Morgan & Claypool** Synthesis Lectures on Computer Architecture Editor Mark D. Hill,University of 2008 iv Chip Multiprocessor Architecture: Techniques to Improve Throughput **Schedule - Rice University Computer Science** Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency (Synthesis Lectures on Computer Architecture) (Paperback) - Common [By **Processor Microarchitecture: An Implementation Perspective** Synthesis Lectures on Computer Architecture publishes 50 to 150 page publications on topics pertaining to the Chip Mutiprocessor Architecture: Techniques to Improve Throughput and Latency. Kunle Olukotun .. 2.4 Multiprocessor Memory Systems . This ability of latent errors to confound error models motivates **The Datacenter as a Computer - EECS:** Synthesis Lectures on Computer Architecture publishes 50 to 150 page publications on Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency . 1.6.4 Quantifying Latency, Bandwidth, and Capacity .