

MICRO toppicks 2010

そういえば読んでないということに気づいた .

iCFP: Tolerating All-Level Cache Misses in In-Order Processors

- in-order processor の性能向上の話 .
- キャッシュミスがおこっても , 投機的にプログラムを実行して , ミスに関する命令だけをケアする
- Miss Level Parallelism(MLP)

Near-Optimal Cache Block Placement with Reactive Nonuniform Cache Architectures

- Reactive NUCA

A Task-Centric Memory Model for Scalable Accelerator Architectures

- Rigel におけるソフトウェアでのキャッシュコヒレンシのメカニズムなど

DMP: Deterministic Shared-Memory Multiprocessing

- fully deterministic な shared-memory multiprocessing <= デバッグ , テスト , 開発が容易に
 - executes multiple threads that communicate via shared memory
 - produces the same program output if given the same program input
- all communication between threads is precisely the same for every execution
- the simplest way
 - メモリ操作のたびにトークンの取得 / 開放を行う
 - performance degradation(メモリアクセスのたびにトークンの開放を待つ / 並列実行の benefit を失う)
- 並列性を recover するために
 - invalidation-based cache coherence protocol をスレッド間通信の発生時の識別に使用
 - 投機実行
- DNP-ShTab(cache coherence protocol)
- DMP-TM, DMP-TMFwd(support for transactional memory)

Making Address-Related Prefetching Practical

- address-related prefetching をするには MB クラスのメタデータが必要だった
- two innovations for managing metadata
 - hash-based lookup: maintains an index of previously recorded miss-address seq.
 - probabilistic update: applies only a randomly selected subset of updates to the hashed index table
 - reducing metadata memory traffic by a factor of 3.4

Accelerating Critical Section Execution with Asymmetric Multicore Architectures

- is based on the asymmetric chip multiprocessor
 - consists of at least one large, high-performance core and many small, power-efficient cores
- small core は , クリティカルセクションに到達したら large core にリクエストを投げて stall する . large core はクリティカルセクションの実行を終えたら small core に通知 , small core

の処理が再開 .

Per-Thread Cycle Accounting

- tracking per-thread progress rates for each coexecuting thread (during SMT exec.)
 - to improve QoS, SLA, Performance predictability, service differentiation, proportional-share performance
- to determine the alone exec. time
 - base cycle component
 - miss event cycle component(cache, TLB, branch)
 - waiting cycle component

AnySP: Anytime Anywhere Anyway Signal Processing

- AnySP = fully programmable architecture
 - targets multiple application domain
 - addresses these challenges for next-generation mobile signal processing
- FFT, space-time block coding(STBC), low-density parity-check(LDPC) を解析
 - SIMD 幅 , レジスタの値のライフタイム , 命令ペアの頻度 , Algorithm data-reordering patterns
- AnySP processing element design
 - Configurable multi-SIMD width support
 - Temporary buffer and bypass network
 - Flexible functional units
 - SRAM-based swizzle network(post-bfabrication でのネットワークの変更が可能)
 - Multiple output adder tree support

Gate-Level Information-Flow Tracking for Secure Architectures

Predicting Voltage Droops Using Recurring Program and Microarchitectural Event Activity

- voltage fluctuations can lead to timing violations or transistor lifetime issues
 - dynamically learns to predict dangerous voltage fluctuations based on program and microarchitectural events
- to reduce the gap between nominal and worst-case operating voltages
 - voltage emergency predictor to identify imminent emergencies

Architectural Implications of Nanoscale-Integrated Sensing and Computing

- a nanoscale sensor processor(nSP)
 - addresses the aforementioned challenges
 - by integrating molecular probe sensors and molecular-scale digital logic
- a sensor array for environment monitoring/a simple processor core/a small memory for state and programs/a communication device

Gordon: An Improved Architecture for Data-Intensive Applications

- Gordon is a system architecture for data-centric applications combining low-power processors, flash memory, and data-centric programming systems
 - for data-centric applications.

- for large-scale data processing, current technology must overcome three challenges
 - uniprocessor での性能低下と CMP のプログラミングの難しさ
 - ハードドライブ容量は増える一方で、レイテンシとバンド幅はそうではない
 - 冷却, お金, エコからくる電力の制約
- data-centric allipcation(Map-Reduce と aDryad とかでプログラムされている)
 - solid-state sotorage device は, 向上したバンド幅とレイテンシ削減に寄与
- Gordon: a flash-based system architecture for massively parallel, data-centric computing
 - solid-state disks
 - low-power processors
 - data-centric programming paradigms
 - offers a flash management software layer that allows a highly parallel operation of large arrays of flash devices

Phase-Change Technology and the Future of Main Memory

- architecting a DRAM alternative
 - array architecture
 - buffer organization(area neutrality, design space, delay and energy opt.)
 - scaling and implications
- mitigating wear and energy
 - improving PCM lifetime(eliminating redundant bit writes, row shifting, segment swapping)
 - analyzing energy implications