CHESS
CENTER FOR HYBRID & EMBEDDED SOFTWARE SYSTEMS

College of Engineering, UC Berkeley

Precision Timed Machines (PRET)

Isaac Liu
Jan Reineke
Sungjun Kim
Hiren Patel
Edward A. Lee

http://chess.eecs.berkeley.edu/

Mission

The traditional computing abstractions only concern themselves with the "functional" aspects of a program and not its timing properties. This allows the use of techniques like speculative execution, caches, interrupts, and dynamic compilation that offer improved average-case performance at the expense of predictable execution times. The PRET project aims to improve the timing predictability at all layers of abstraction by carefully reexamining and reworking various architectural and compiler advancements with an eye toward their effects on timing behavior and worst-case bounds.

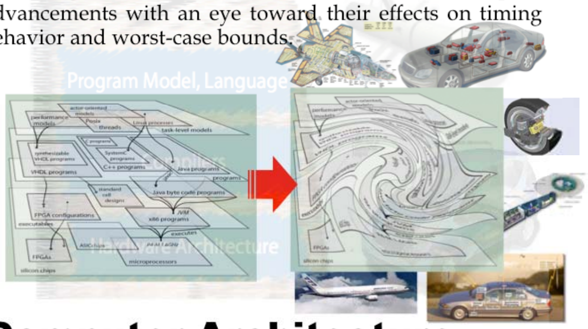

Computer Architecture

In our research, we have pursued a bottom-up approach. Starting with the underlying pipeline, we have modified and added constructs to improve the timing predictability at the architectural and instruction-set-architecture level

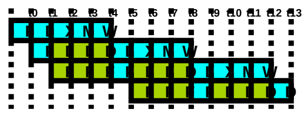
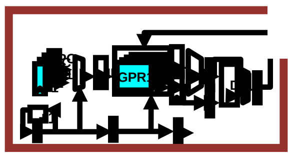

Interleaved multithreading

By using interleaving threads through our piprline, we are able to remove the data hazard and dependencies in the pipeline, creating a timing predictable pipeline.

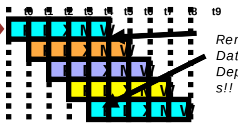

Memory Hierarchy

Conventional memory systems uses a hierarchy of memory units to bridge the latency. CPUs use registers for fast data processing operations. However, the memory hierarchy is designed as "best effort" latency and bandwidth requirements. Our goal is to look at modern memory hiearchies and design a predictable memory system.

A DRAM device consists of several banks along with controller logic to decode addresses. Concurrent accesses to banks are possible, but I/O pins are shared.

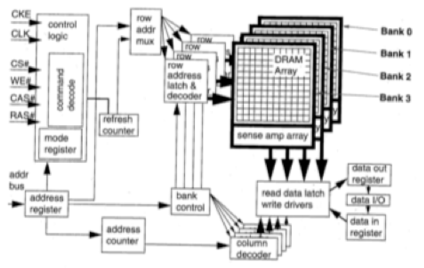

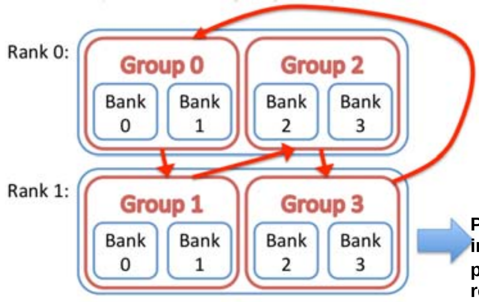

Provides four independent and predictable resources

Predictable DRAM Access

DDR2 DRAM devices utilizes bank parallelism to achieve better performance. It's difficult to predict DRAM access time because accesses to the same bank need to wait for the previous access to finish, while accessing a different bank can be done concurrently. Our DRAM controller exposes groups of banks as independent resources that are accessed sequentially which hides the latency of accessing a single bank. A TDMA scheduler provides predictable access latencies.

Scratchpad Memories

Caches can use different hardware replacement policies in attempt to prefetch the data needed from main memory. However, when the software issues a load or store instruction, it does not know the state of the cache, or what data has been prefetched in it. Scratchpad memories are another form of fast access memory which is managed in software, much like registers that use explicit load/store instructions to manipulate contents for fast data processing. With software management, better real time guarantees can be provided.

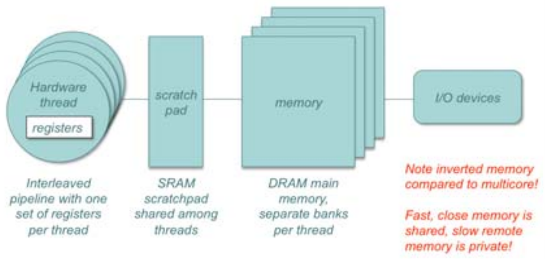
Hardware thread
registers
scratch pad
memory
I/O devices
Interleaved pipeline with one set of registers per thread
SRAM scratchpad shared among threads
DRAM main memory, separate banks per thread
Note inverted memory compared to multicore!
Fast, close memory is shared, slow remote memory is private!

Above shows one possible memory hierarchy for PRET. The memory system is a major source of headache for analyzing or predicting execution, as it causes a wide range of execution time for even the same programming running on the same system. Our goal is to design a predictable memory system that provide systems with predictable memory access times.

Acknowledgement

This work was supported in part by the Center for Hybrid and Embedded Software Systems (CHESS) at UC Berkeley, which receives support from the National Science Foundation (NSF awards #0720882 (CSR-EHS:PRET) and #0720841 (CSR-CPS)), the U. S. Army Research Office (ARO#W911NF-07-2-0019), the U. S. Air Force Office of Scientific Research (MURI #FA9550-06-0312), the Air Force Research Lab (AFRL), the State of California Micro Program, and the following companies: Agilent, Bosch, HSBC, Lockheed-Martin, National Instruments, and Toyota.