Software-Managed Caches: Architectural Support for Real-Time Embedded Systems

Bruce Jacob

Electrical (and Computer) Engineering
University of Maryland, College Park

OUTLINE:

- **Motivation: the problem with caches**
- **One Solution: software-managed**

UNIVERSITY OF MARYLAND

SOFTWARE-MANAGED CACHES

Bruce Jacob

University of Maryland

Why Traditional Caches Suck

Solutions to the Problem

USE DSP-STYLE DATA CACHES

- Software explicitly manages movement
- What about instructions?

WIRE DOWN REGIONS OF MEMORY

- Usually at a page granularity (in TLB)
- Requires operating system assistance

PARTITION THE CACHES

- Solves part of the problem
- Inter-partition consistency an issue

DISABLE CACHES

SOFTWARE-MANAGED CACHES

Bruce Jacob

University of Maryland

Software-Managed Caches

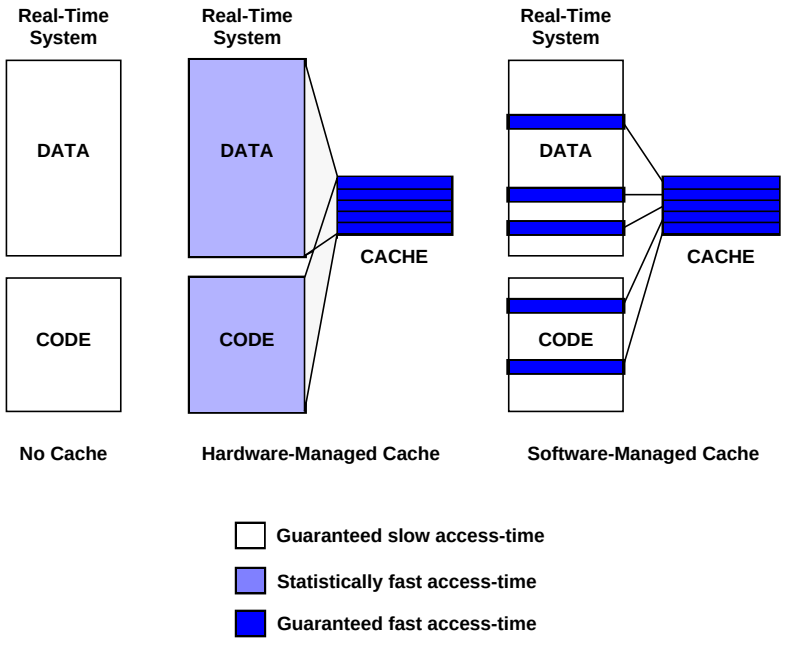

Software-Managed Caches

Top bits determine memory-access behavior

OS ensures that virtual-physical translation is the same for both subspaces

Other possibilities:

- **Physical/virtual**
- **Faulting/non-faulting**
- **Which cache or memory structure**

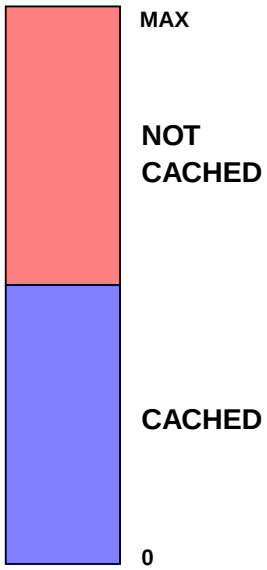

SOFTWARE-MANAGED CACHES
Bruce Jacob
University of Maryland

Application Behavior

```
int *array = malloc (N * sizeof int);
int *stream = malloc (N * sizeof int);
int *mix = malloc (N * sizeof int);

for (i=0; i<N; i++)
    x = array[i];

stream |= MIN_NEG_INT; /* 0x80000000 */
for (i=0; i<N; i++)
    x = stream[i];

for (i=0; i<N; i++)
    x = (cache_it (i)) ? mix[i]
            : (mix | MIN_NEG_INT)[i];
```

Compiler Behavior

Group Data & Instructions to MINIMIZE CACHE CONFLICTS

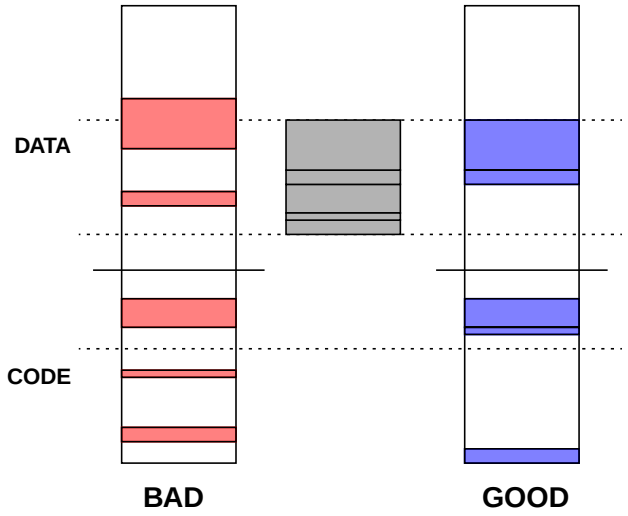

SOFTWARE-MANAGED CACHES

Bruce Jacob

University of Maryland

Issue: Discontinuities

DATA SPACE

- **Relatively easy to rearrange items**

CODE SPACE

- **Can move FUNCTIONS around easily**
- **PORTIONS of code is another matter ...**

FINE-GRAINED PLACEMENT:

- **Virtual addresses vs. physical addresses**
- **Segmented addresses potentially better**

PROBLEM: fine-grained relocation at granularity of TLB page

Performance: GCC

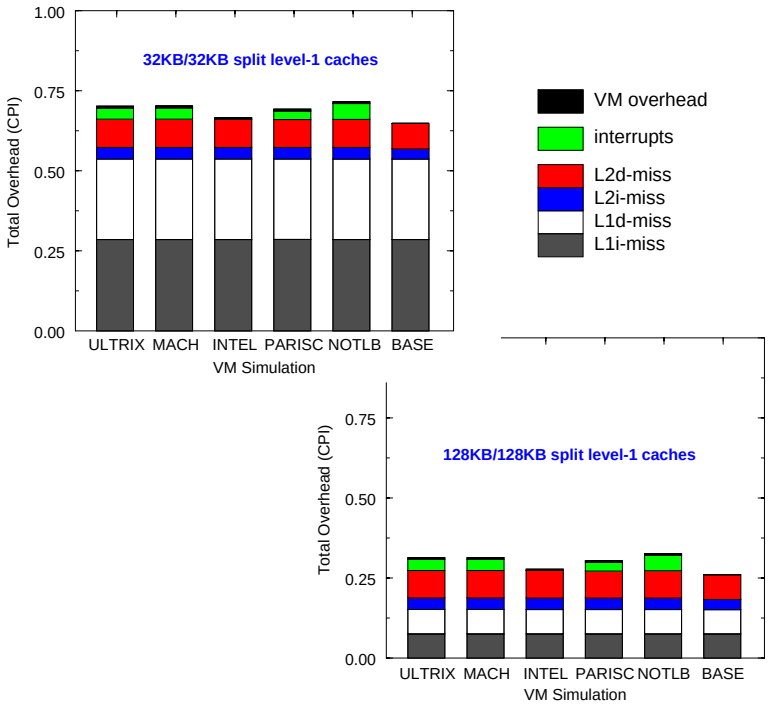

SOFTWARE-MANAGED CACHES

Bruce Jacob

University of Maryland

Summary

Hardware support has been explored:

Bruce L Jacob and Trevor N Mudge. “Software-managed address translation.” *Proc. Third International Symposium on High Performance Computer Architecture (HPCA-3)*, pp. 156-167. San Antonio TX, February 1997.

Bruce L Jacob and Trevor N Mudge. “A look at several memory management units, TLB-refill mechanisms, and page table organizations.” *Proc. Eighth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-8)*, pp. 295-306. San Jose CA, October 1998.

Hardware/software issues being explored as part of F-ZONE project

http://www.ee.umd.edu/~blj/