Software-Managed Address Translation

Bruce Jacob and Trevor Mudge

Advanced Computer Architecture Lab
University of Michigan

OUTLINE:

- **Background & Motivation**
- ***Softvm* Design**
- **Experiments & Results**

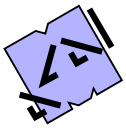

ADVANCED COMPUTER ARCHITECTURE LAB

SOFTWARE-MANAGED ADDRESS TRANSLATION

Bruce Jacob

University of Michigan

Historical Perspective

MIPS:

- **HW page-table-walking unnecessary**

SPUR:

- **HW storage for PTEs unnecessary**

VMP:

- **Software-controlled caches work**

CONCLUSION:

Can get rid of TLB altogether

SOFTWARE-MANAGED ADDRESS TRANSLATION

Bruce Jacob

University of Michigan

TLB Architecture

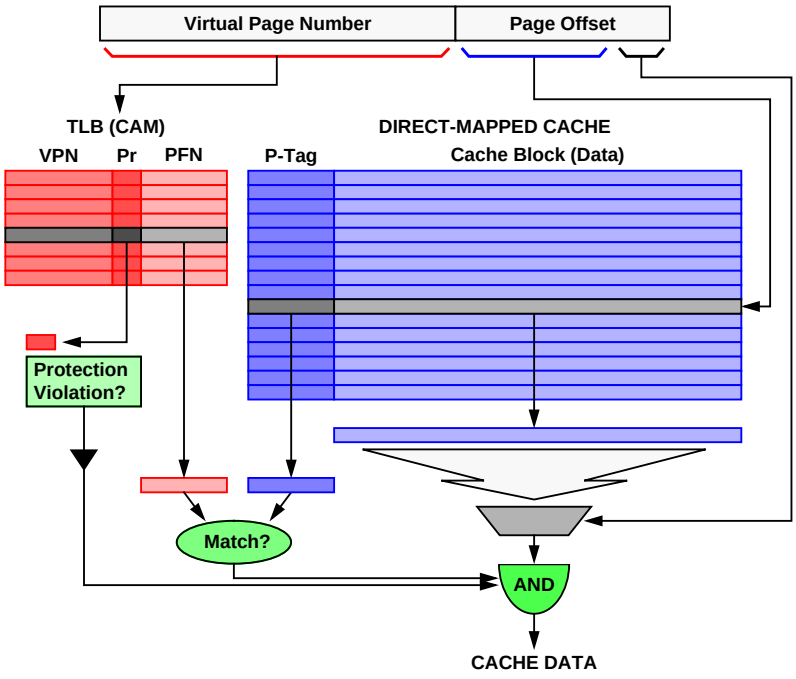

SOFTWARE-MANAGED ADDRESS TRANSLATION

Bruce Jacob

University of Michigan

TLB Miss

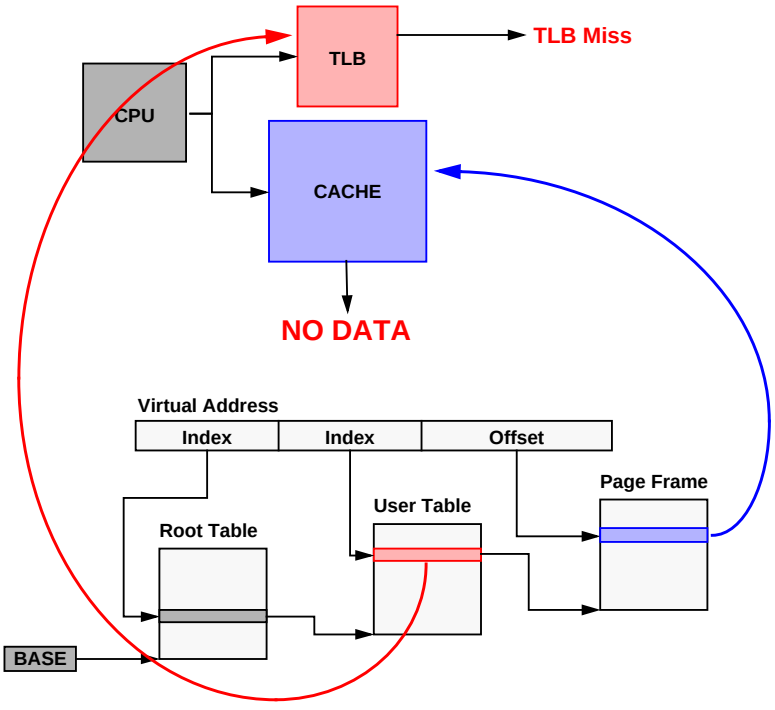

SOFTWARE-MANAGED ADDRESS TRANSLATION

Bruce Jacob

University of Michigan

SOFTVM Architecture

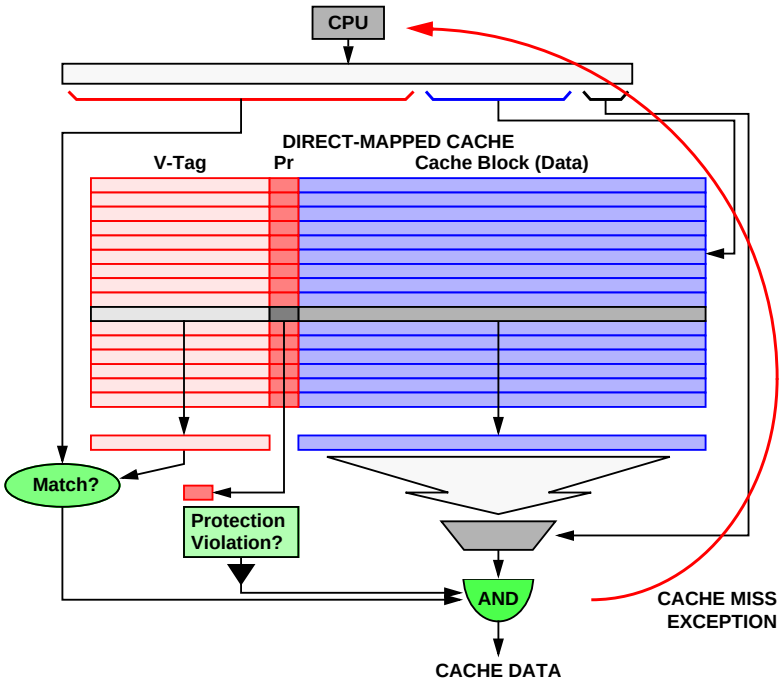

SOFTWARE-MANAGED ADDRESS TRANSLATION

Bruce Jacob

University of Michigan

SOFTVM Miss

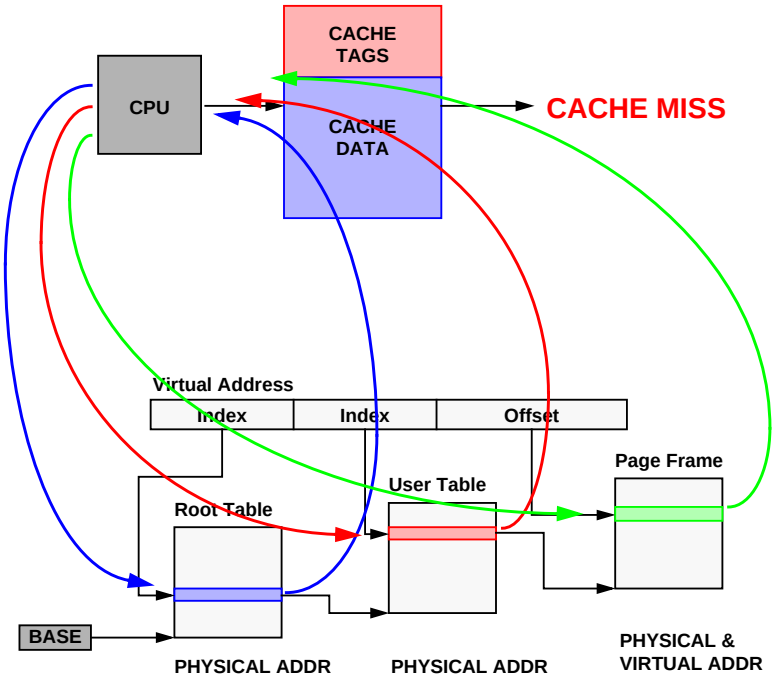

Cost of Software Solution

Cost to handle TLB Miss: ~12 cycles

- **Build PA & Load RPTE**
  Build PA & Load UPTE
  Insert UPTE into TLB
  Retry Load

Cost to handle Cache Miss: ~15 cycles

- **Build PA & Load RPTE**
  Build PA & Load UPTE
  Prepare caches for USER-DATA
  Build PA & Load USER-DATA
  Retry Load

SOFTWARE-MANAGED ADDRESS TRANSLATION

Bruce Jacob

University of Michigan

Cost vs. Cache Size

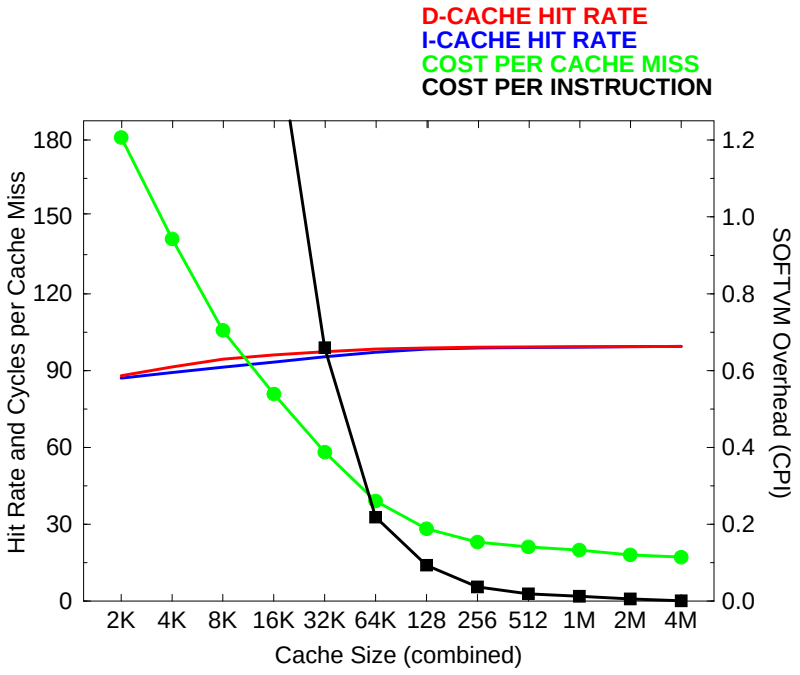

Experiments

Modified 32-BIT PowerPC Architecture, MIPS-like Page Table

Trace-Driven Simulations:

- **L1 Cache (20):** **2—256 KBytes**
- **L2 Cache (100):** **1, 2, 4 MBytes**
- **Linesizes:** **16—128 Bytes**

RESULTS:

- **Softvm:** **0.1 to 5% Overhead**
- **Mach+MIPS:** **5 to 10% Overhead**
- **Ultrix+MIPS:** **2% Overhead**

SOFTWARE-MANAGED ADDRESS TRANSLATION

Bruce Jacob

University of Michigan

Results

L1 Caches: 16/16KB D/I, 32-byte linesize
L2 Caches: 64-byte linesize

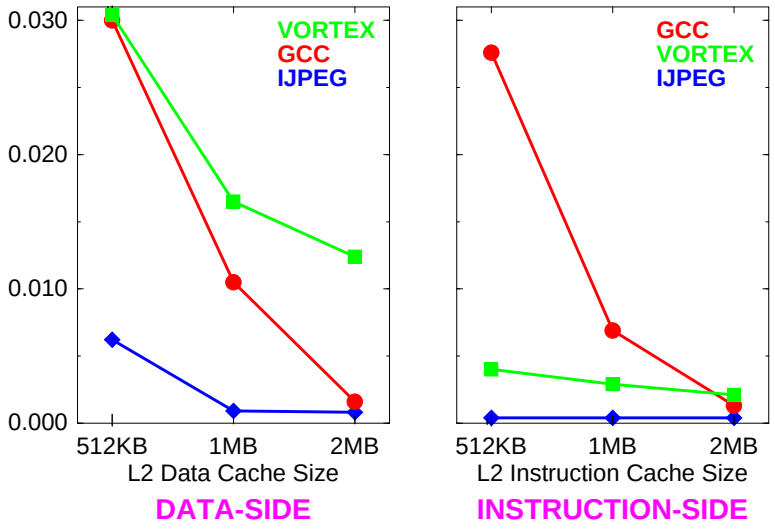

Potential Problem: STREAM

MULTIMEDIA HAS
NO TEMPORAL LOCALITY

WORST-CASE SCENARIO:

Take an exception for every cache line

SOLUTIONS:

- **Prefetch buffers**
- **Prefetch into L2 cache**
- **Provided unmapped regions to user**

SOFTWARE-MANAGED ADDRESS TRANSLATION

Bruce Jacob

University of Michigan

Design Considerations

LARGE VIRTUAL CACHES:

Synonym Problem

SOLUTION:

Segmentation
or
Large ASIDs w/ Flat Address Space

DRAWBACK:

Increases size of cache tags

Bruce Jacob

University of Michigan

Conclusions

TLB Elimination is Possible

Cycle Time can DECREASE

Performance can INCREASE

Support for Multimedia Possible

Software-Managed: FLEXIBILITY