

# Software-Oriented Memory-Management Design

**Bruce Jacob**

**Advanced Computer Architecture Lab  
University of Michigan**

## OUTLINE:

- **Motivation**
- **Architecture**
- **Evaluation**
- **Problems & Solutions**

## Motivation


### Trends in Microprocessor Design:

- **Faster clock speeds**
- **Larger and cheaper memories**
- **More on-chip devices**
- **Increased flexibility**
- **Increased focus on testability and reliability**

### One Conclusion:

- **Make it SIMPLE**
- **Do it in SOFTWARE**

## Implications for Memory Management?



### Translation Lookaside Buffer

- Large, fully associative, on critical path

### Finite State Machine

- Good performance, limited flexibility


## SOFTVM Design

**VIRTUAL CACHE HIERARCHY**  
replaces  
**Translation Lookaside Buffer**

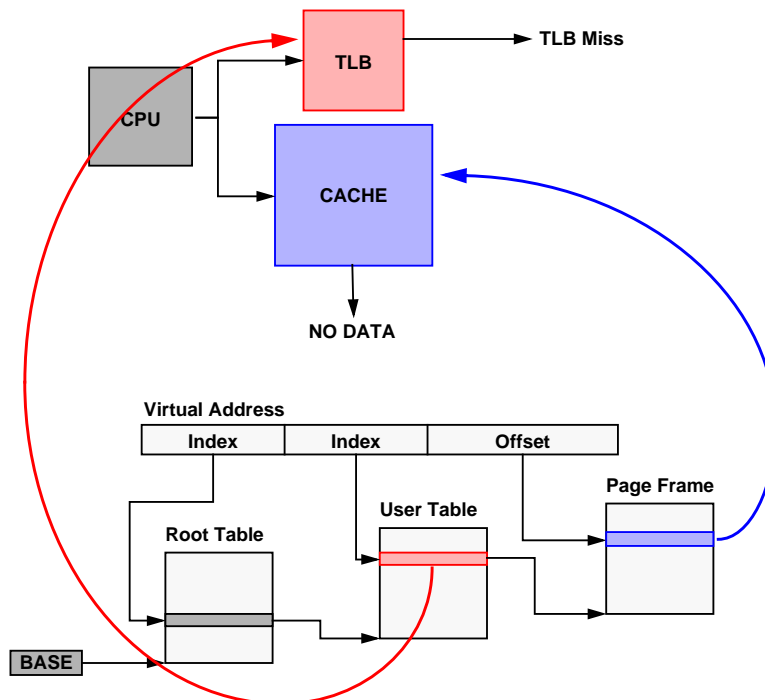
**CACHE-MISS INTERRUPT**  
replaces  
**Finite State Machine**

**Thesis tests validity  
of software-oriented design  
on memory management**

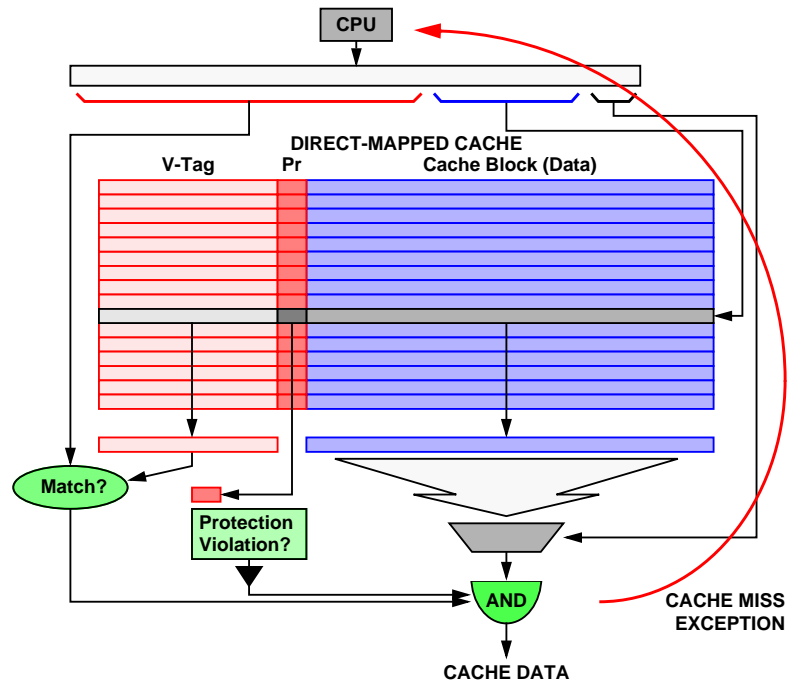
# TLB Architecture



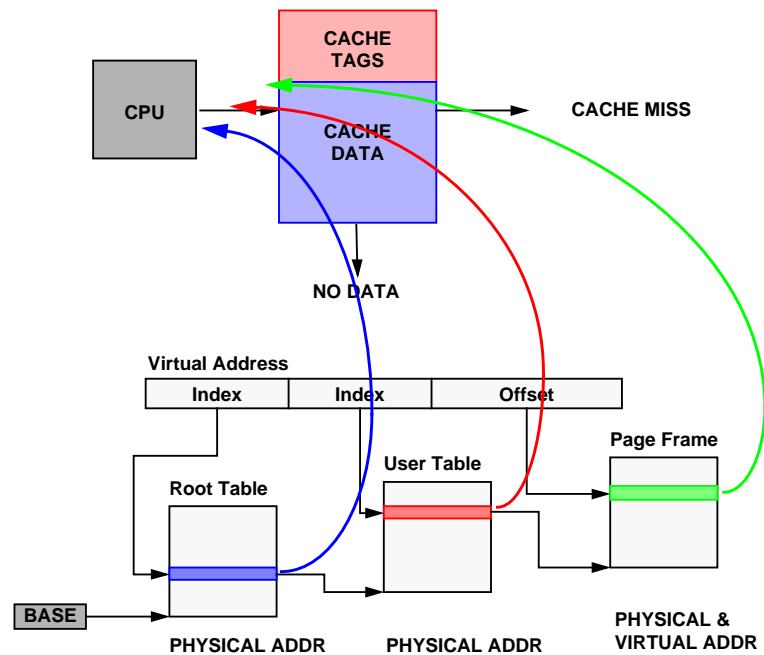
# TLB Miss



# SOFTVM Architecture



# SOFTVM Cache Miss



## Expectations

- ✓ **REDUCTION in DIE AREA**
- ✗ **SENSITIVE to CACHE ORGANIZATION**
- ✓ **MISSES INFREQUENT rel. to TLB**
- ✗ **MISSES EXPENSIVE rel. to TLB**
- ✓ **INCREASED FLEXIBILITY**

## Evaluation


### Simulated 5 Virtual Memory Designs:

- **SOFTVM** - No TLB
- **Ultrix/MIPS** - SW-mgd TLB, part.
- **Mach/MIPS** - SW-mgd TLB, part.
- **BSD/Intel** - HW-mgd TLB, no part.
- **PA-RISC** - SW-mgd TLB, no part.

**TLBs: 128-entry I-TLB, 128-entry D-TLB**  
**Fully associative, Random replacement,**  
**16 entries in Protected Partition**

**Virtual cache hierarchies, Perfect memory**


## SOFTVM VM-simulation



**Root-level handler:** 20 inst, 1 PTE load

**User-level handler:** 10 inst, 1 PTE load

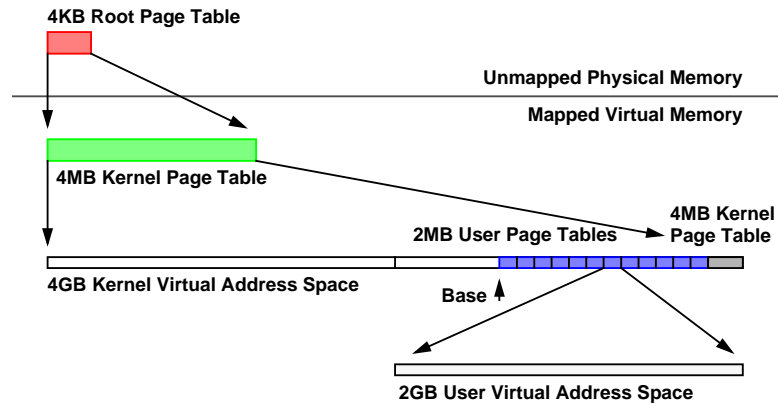
## ULTRIX VM-simulation



**Root-level handler:** 20 inst, 1 PTE load

**User-level handler:** 10 inst, 1 PTE load

# MACH VM-simulation




**Root-level handler:** 500 inst, 1 PTE load

**Kernel-level handler:** 20 inst, 1 PTE load

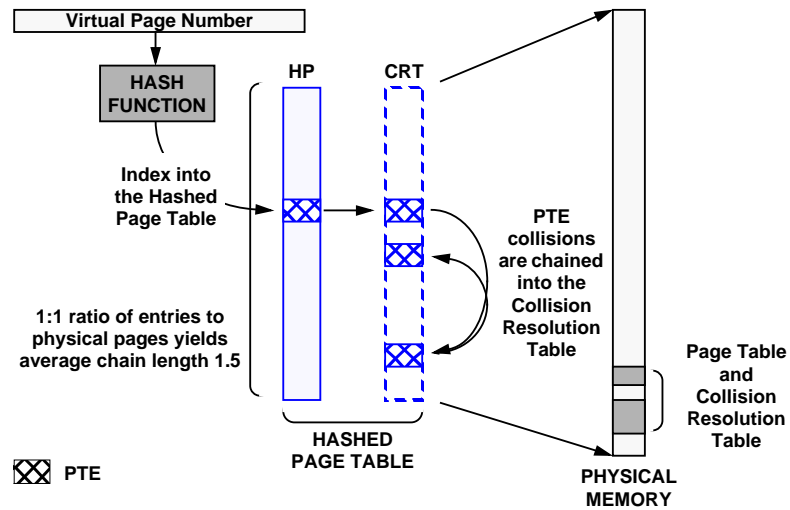
**User-level handler:** 10 inst, 1 PTE load

# INTEL VM-simulation



**User-level handler:** 7 cycles, 2 PTE loads

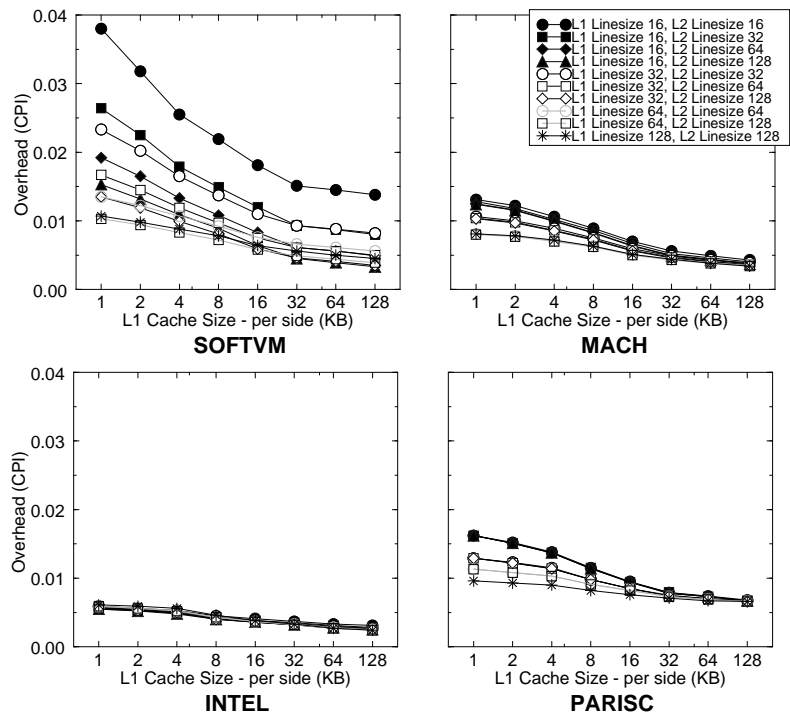
# PARISC VM-simulation



**User-level handler: 20 inst,  $\geq 1$  PTE loads**

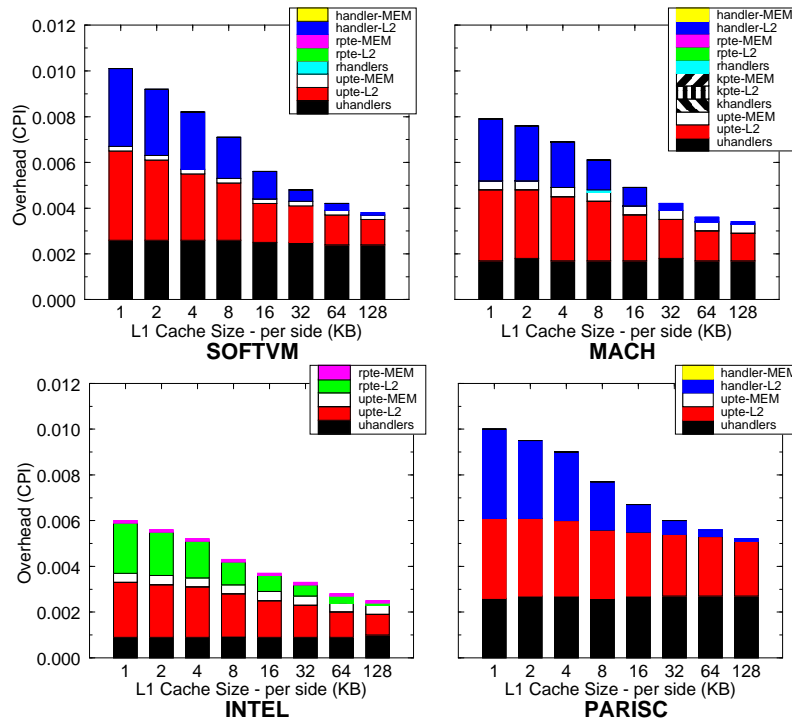
**PTEs are 2x size of other tables**

# VM Performance: GCC

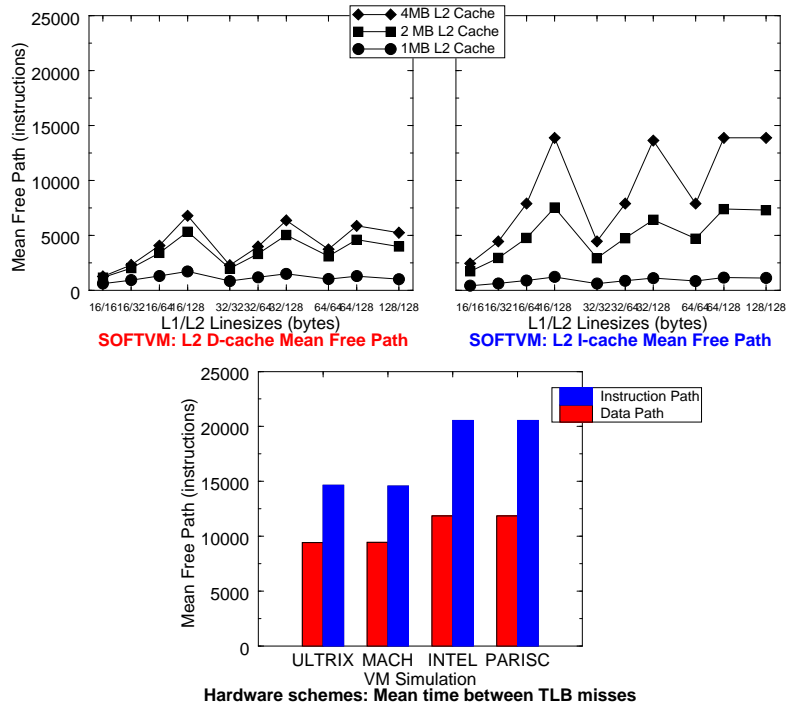




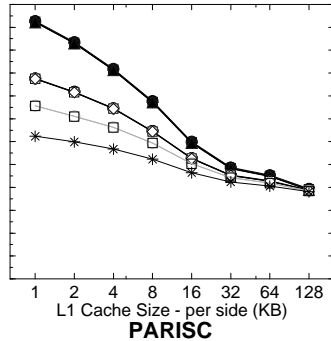
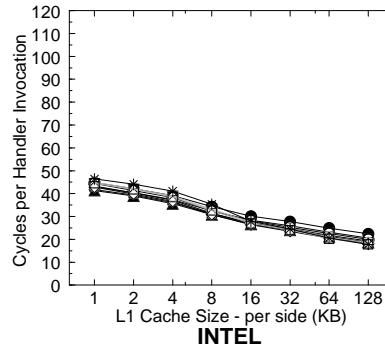
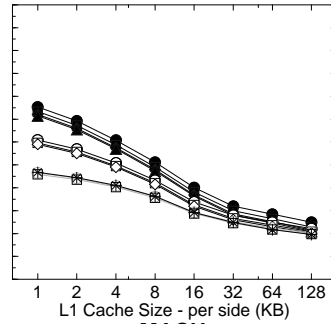
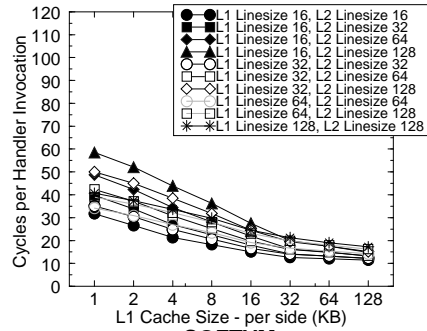
# VM Break-downs: GCC



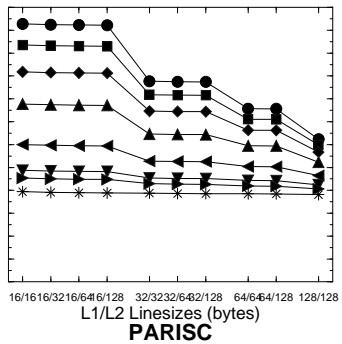
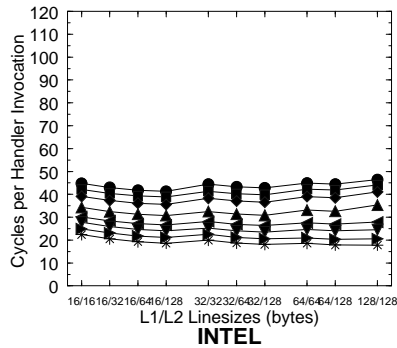
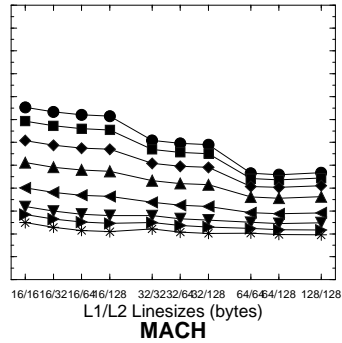
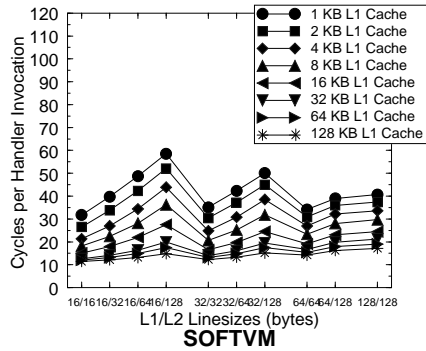
# How Often: GCC



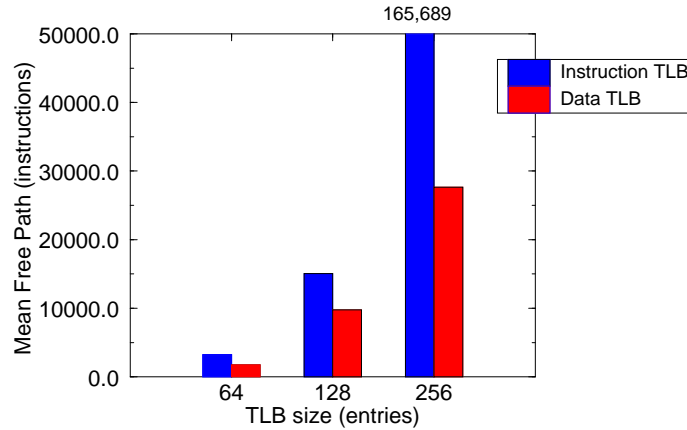
# How Much: GCC



# How Much: GCC



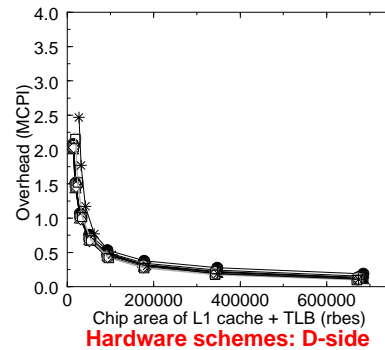
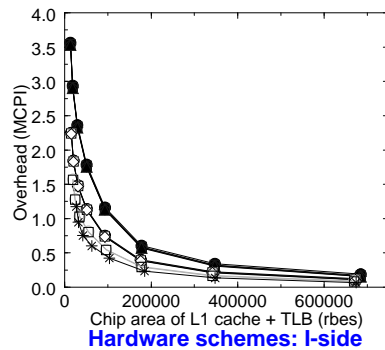
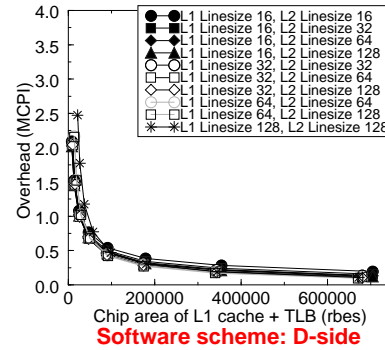

# TLB Sensitivity: GCC



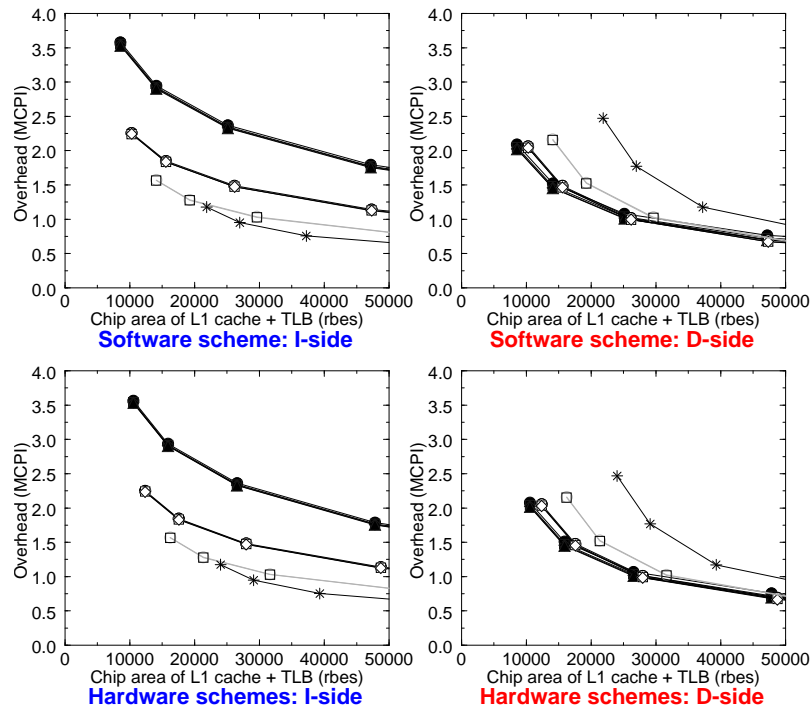
Mean free path of GCC — ULTRIX page table, 4MB caches

**Hardware schemes VERY sensitive to TLB miss-rate ... Software scheme would look 5x better against 64-entry TLBs**

# Bottom Line: GCC



## Bottom Line (closeup): GCC



## Problems & Solutions ...

### USE OF EXCEPTION MECHANISM

**Exception Problem**

### CACHE-BOUND PERFORMANCE

**Multimedia Problem**

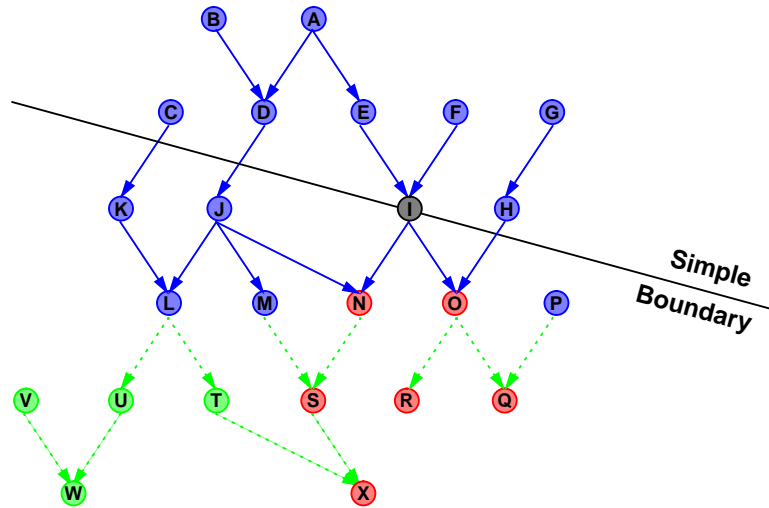
### LARGE VIRTUAL CACHES

**Synonym Problem**

### LEVEL-2 CACHES

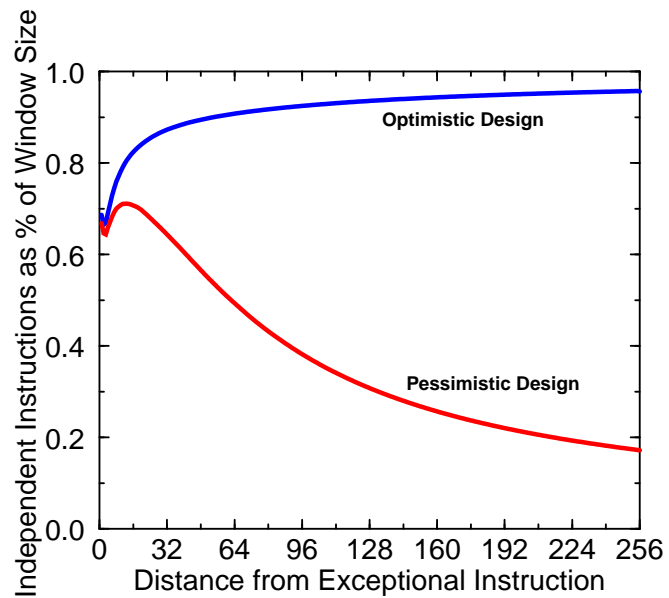
**Cost Problem**

## Exception Problem



**Cost of Interrupt =  
Cost of flushing Reorder Buffer**

## Potential Gains



# Re-Execute Buffer

ROB Entry contains **DEPENDENCY INFO**

On Interrupt, **DO NOT** Flush Reorder Buffer

At Instruction Commit:

- If **DONE**, Commit
- If waiting for **EXECUTE**, Wait
- If waiting for **OPERAND**, Re-Execute

# ROB/REB Illustration

Reorder Buffer		Re-Ex Buffer
Inst	Dep's	Buffer
A	-	-
B	-	-
C	A, B	-
D	A, D	-
E*	C	-
F	D, E	-
G	D	-
H	G	-

Detection of Exception

Reorder Buffer		Re-Ex Buffer
Inst	Dep's	Buffer
E*	-	E: -
F	D, E	-
G	D	-
H	G	-
x1	-	-
x2	x1	-
x3	x2	-
x4	x3	-

Inst E at Head of ROB

Reorder Buffer		Re-Ex Buffer
Inst	Dep's	Buffer
x1	-	E: -
x2	x1	F: D,E
x3	x2	next: I
x4	x3	-
x5	x4	-
x6	x5	-
x7	x6	-
x8	x7	-

Handler at Head of ROB

Reorder Buffer		Re-Ex Buffer
Inst	Dep's	Buffer
xW	xV	-
xX	xW	-
xY	xX	-
xZ	xY	-
E	-	-
F	D, E	-
I	-	-
J	I	-

Re-Execution of E & F

# Multimedia Problem STREAM

**MULTIMEDIA HAS  
NO TEMPORAL LOCALITY**

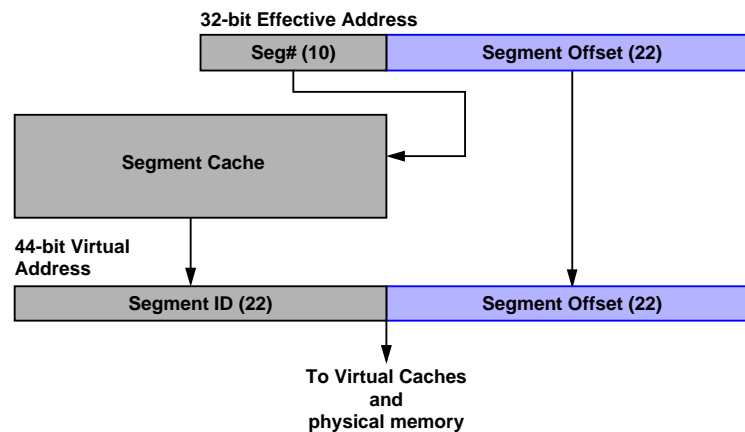
**WORST-CASE SCENARIO:**

**Take an exception for every cache line**

**SOLUTIONS:**

- Prefetch buffers
- Prefetch into L2 cache
- Provide unmapped regions to user

# Unmapped 4MB Superpages



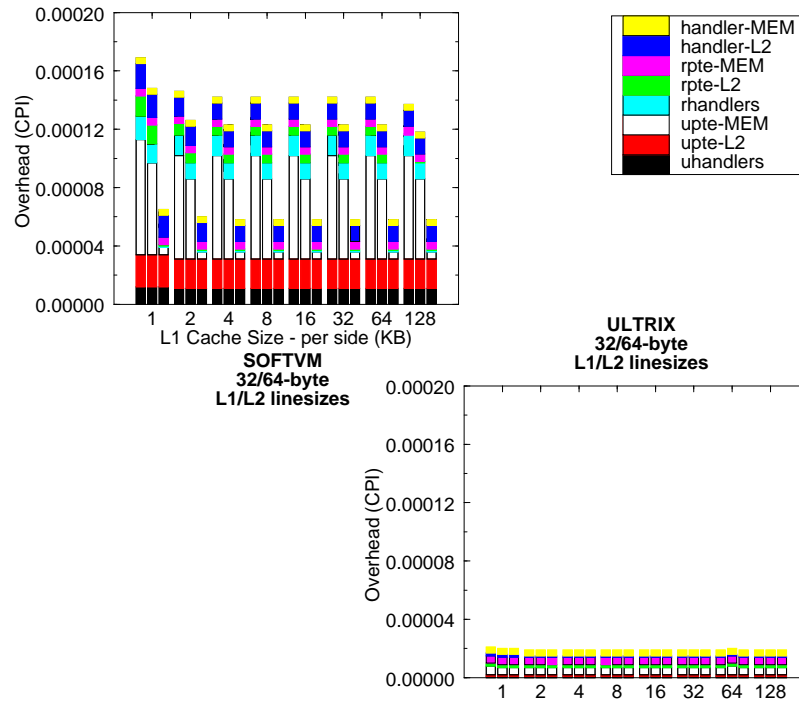
**Alternative (Non-Segmented) Solution:**

- **Small Superpage TLB (BAT Registers)**

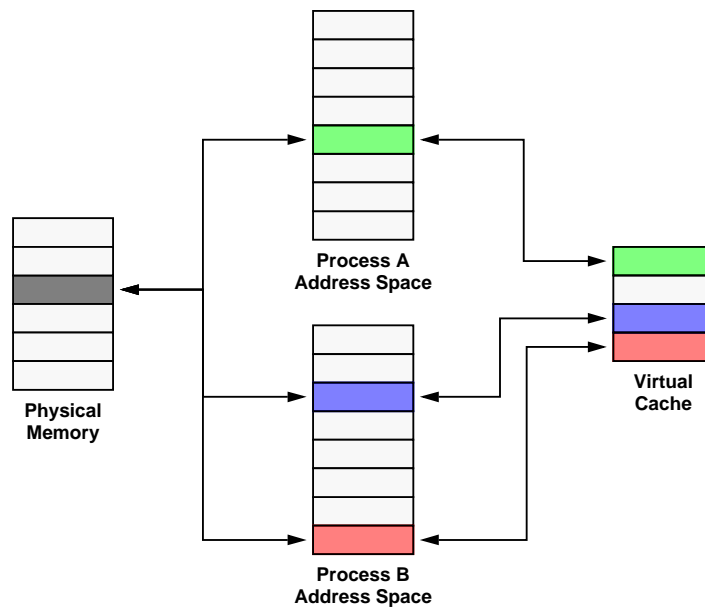




## ... with Superpages



## Virtual Cache Synonym Problem



# Solutions to Synonym Problem

## HARDWARE

- **Backpointers**

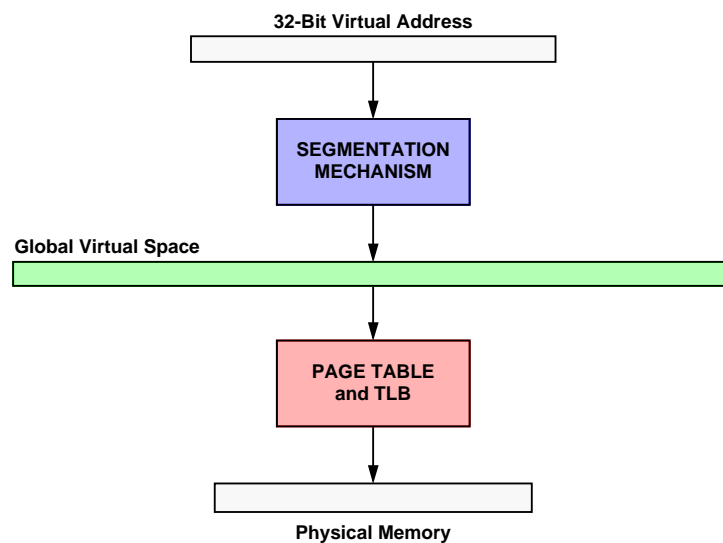
## SOFTWARE

- **OS/2: Eliminate Aliasing**
- **SunOS: Aliases map to same cache line**
- **SASOS: Eliminate Aliasing**

## HARDWARE/SOFTWARE


- **Segmentation ....**

# Segmentation



## Segmented Solution

**ALIASING** without **SYNONYM PROBLEMS**



## Cost Problem

**Solution Requires (LARGE) L2 Caches**

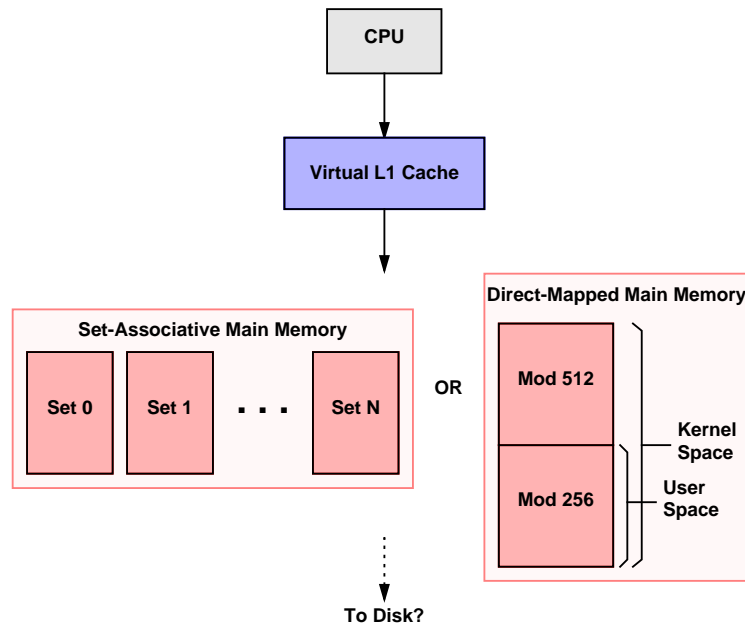
**What about embedded market?**

**What about cost-cutter market?**

**Implement L2 Caches in DRAM,  
Call it MAIN MEMORY**

**Requirement: ability to  
pin down critical regions ...**

# Virtually-Addressed Main Memory



## Conclusions

Elimination of MMU is Possible

Cycle Time can DECREASE

Performance can INCREASE

Software-Managed: FLEXIBILITY

Examples of Usefulness:

- Shared-Memory Multiprocessor
- Real-Time Processing
- Architecture Emulation
- Restricted Reconfigurable Computing

