

Designing Multicores for Programmability: The Bulk Multicore Architecture

Josep Torrellas

Department of Computer Science
University of Illinois at Urbana-Champaign
<http://iacoma.cs.uiuc.edu>



SNU, Sep 2011



The Multicore Era

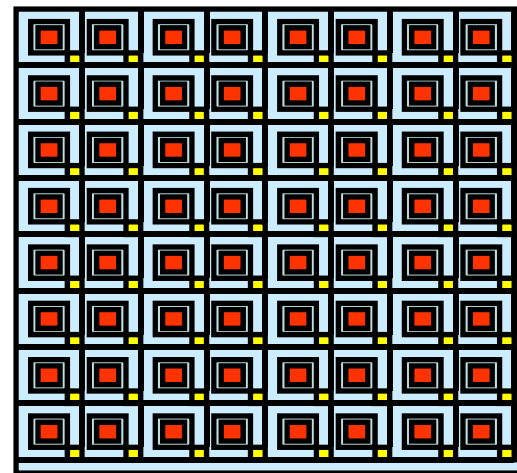
- Goals: performance, energy-efficiency & programmability
- What is a Programmable Architecture?
 - Attains **high efficiency while relieving** the programmer from low-level tasks
 - Helps **minimize the chance** of (parallel) programming errors

General-purpose shared-memory multicore

- **Novel scalable cache-coherence (signatures & chunks)**
 - Relieves programmer/runtime from managing shared data
- **High-performance sequential memory consistency**
 - Provides a SW-friendly environment
- **HW primitives for low-overhead program development & debugging**
(data-race detection, deterministic replay, address disambiguation)
 - Helps reduce the chance of parallel programming errors
 - Overhead low enough to be “on” during production runs

The Bulk Multicore

- Idea: **Eliminate the commit of individual instructions at a time**
- Mechanism:
 - Processors continuously commit **chunks** of instructions at a time (e.g. 5,000 **dynamic** instr)
 - Chunks execute **atomically** and in **isolation** (using buffering and undo)
 - Memory effects of chunks summarized in **HW signatures**
 - Chunks can be invisible to SW or generated by compiler
- Advantages over current:
 - Higher programmability
 - Higher performance
 - Simpler processor hardware



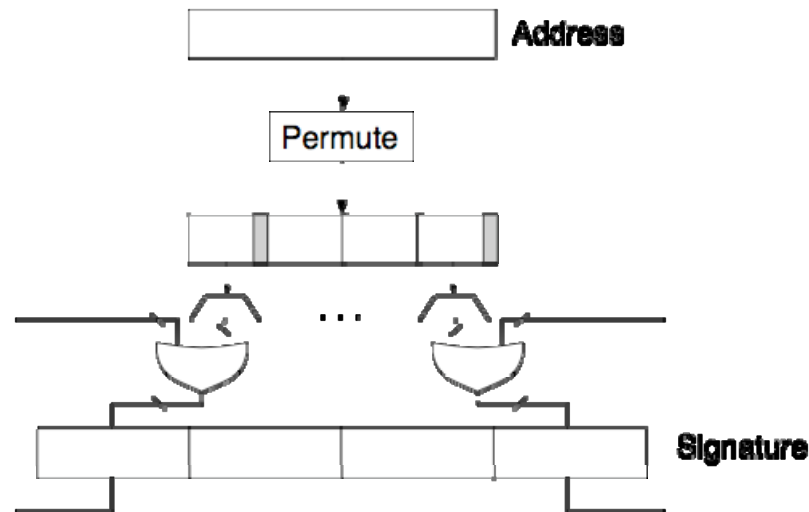
The Bulk
Multicore

Rest of the Talk

- **The Bulk Multicore**
- How it improves performance
- How it improves programmability

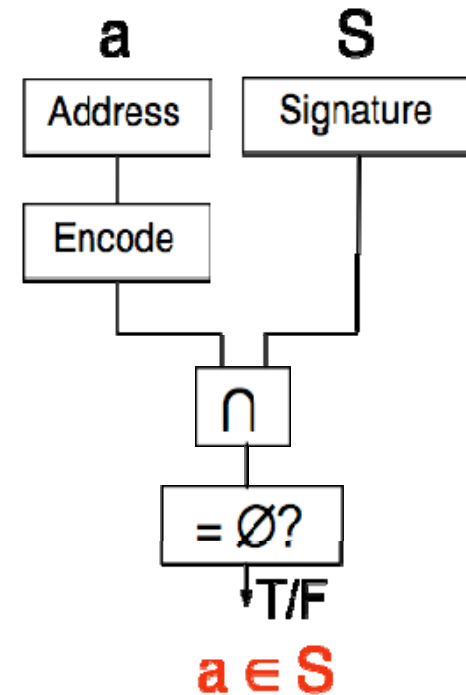


Hardware Mechanism: Signatures [ISCA06]

- Hardware accumulates the addresses read/written in signatures



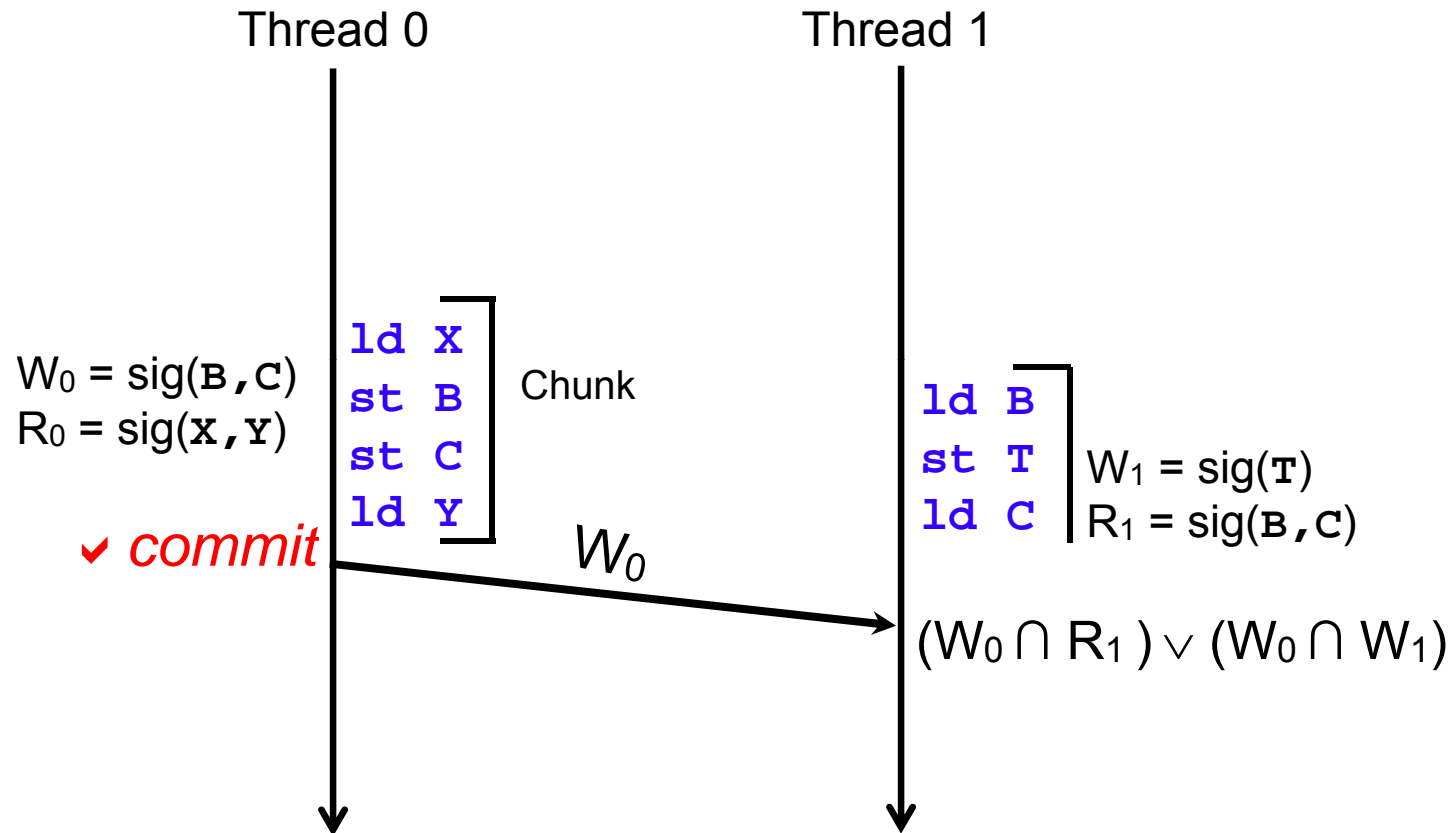
- Read and Write signatures
- Summarize the footprint of a **Chunk** of code

Signature Operations In Hardware



Inexpensive Operations on Groups of Addresses


Executing Chunks Atomically & In Isolation: Simple!



Chunk Operation + Signatures: Bulk

[ISCA07]

- Execute each chunk **atomically** and in **isolation**
- (Distributed) arbiter ensures a **total order** of chunk commits



- **Supports Sequential Consistency** [Lamport79]:
 - **High performance:** Instructions are fully reordered by HW
Loads and stores make it in any order to the sig
Fences are NOOPS
 - **Low HW complexity:** Need not snoop ld buffer for consistency

Summary: Benefits of Bulk Multicore

- Gains in HW simplicity, performance, and programmability
- Hardware simplicity:
 - Memory consistency support moved away from core
 - Toward **commodity cores**

Rest of the Talk

- The Bulk Multicore
- How it improves performance
- How it improves programmability

High Performance

- HW reorders accesses heavily (**intra-** and **inter-**chunk)
- If chunks driven by compiler: Novel compiler optimizations

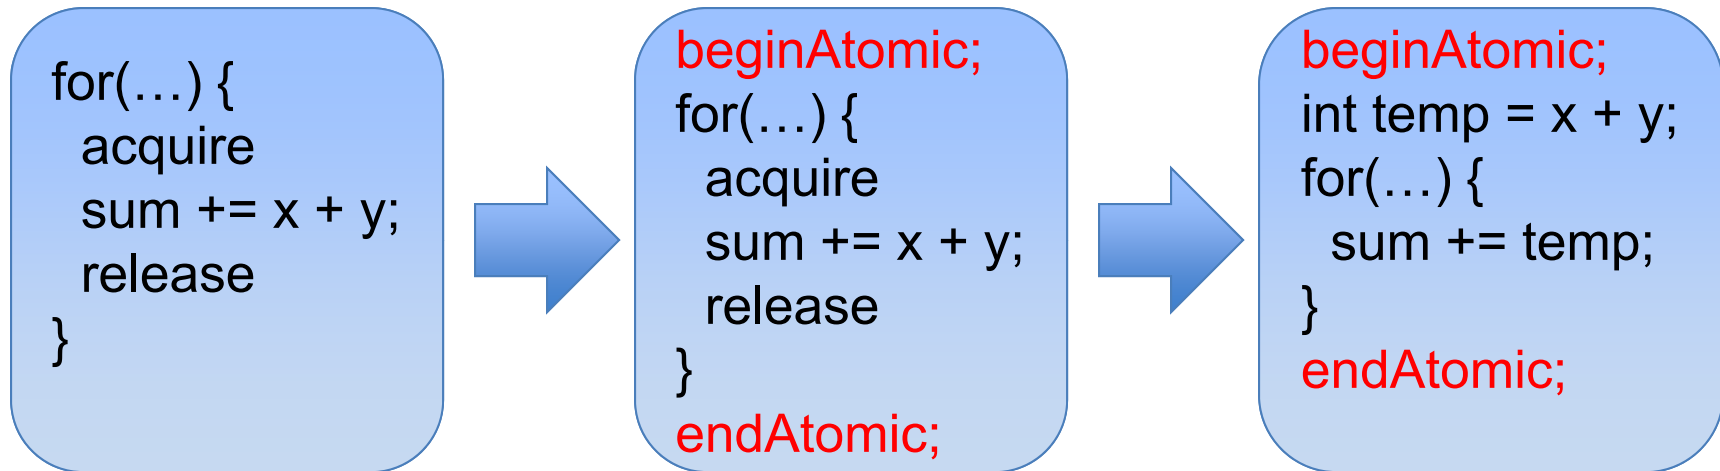
BulkCompiler: Compiler for Bulk Multicore [MICRO-09]

- Takes code with **synchronization** operations (locks, barriers..)
- Adds instructions to drive chunking
 - *beginAtomic PC*
 - Starts new chunk
 - Takes as argument the PC of the **Safe-Version** of code
 - *endAtomic*
 - Finishes the current chunk

→ Atomicity allows compiler to:
→ Optimize the code within chunks
→ Ignore memory model restrictions

Example of BulkCompiler Optimization

- sum, x, y are shared variables



- HW guarantees atomic execution (no synchs needed)
- Compiler allowed to perform arbitrary optimizations inside
- If another thread accesses sum, x, y
 - HW detects failed speculation, squashes, and retries chunk

More Complete Transformation

- **Low-contention** critical sections:
 - Group many of them in same Atomic Region (AR)

beginAtomic

i₁

acquire M1

i₂

release M1

i₃

acquire M2

i₄

release M2

i₅

endAtomic

beginAtomic

i₁

while (M1 == taken) {}

i₂

i₃

while (M2 == taken) {}

i₄

i₅

endAtomic



- Remove acquire / release
- Insert **plain** spins on lock variables
 - Lock may be owned
 - Owner will squash you on release
- **Optimize** and reorder the code

More Complete Transformation

- **Low-contention** critical sections:
 - Group many of them in same Atomic Region (AR)

beginAtomic

i₁

acquire M1

i₂

release M1

i₃

acquire M2

i₄

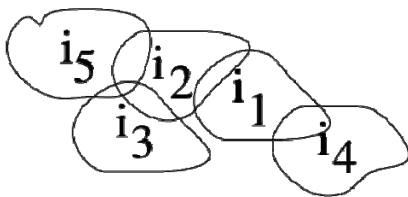
release M2

i₅

endAtomic



beginAtomic
while (M1 == taken) {}
while (M2 == taken) {}
endAtomic



- Remove acquire / release
- Insert **plain** spins on lock variables
 - Lock may be owned
 - Owner will squash you on release
- **Optimize** and reorder the code

More Complete Transformation

- **Low-contention** critical sections:
 - Group many of them in same Atomic Region (AR)

beginAtomic

i₁

acquire M1

i₂

release M1

i₃

acquire M2

i₄

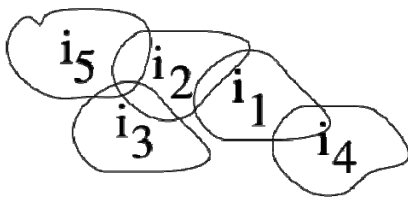
release M2

i₅

endAtomic



beginAtomic
while (M1 == taken) {}
while (M2 == taken) {}
endAtomic



- Remove acquire / release
- Insert **plain** spins on lock variables
 - Lock may be owned
 - Owner will squash you on release
- **Optimize** and reorder the code



No conventional compiler can do this

Rest of the Talk

- The Bulk Multicore
- How it improves performance
- How it improves programmability


High programmability

- Invisible to the programming model/language
- Supports **Sequential Consistency (SC)**
 - * **Software correctness tools** assume SC
- Enables **novel always-on debugging** techniques
 - * Only keep **per-chunk** state, not per-load/store state
 - * Deterministic replay of parallel programs **with no log**
 - * Data race detection at **production-run speed**

Concept: Deterministic Replay of MP Execution

- During **Execution**: HW records into a log the order of dependences between threads
- The log has captured the “interleaving” of threads
- During **Replay**: Re-run the program
 - Enforcing the dependence orders in the log

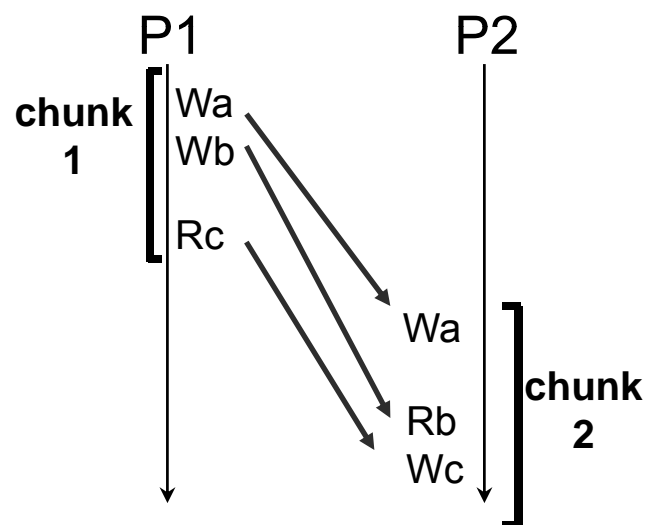
Conventional Schemes



- Potentially large logs

Bulk: Log Necessary is Minuscule [ISCA08]

- During **Execution**:
 - Commit the instructions in chunks, not individually



Combined Log
of all Procs:

P1

P2


Pi

If we **fix** the chunk commit interleaving:

Combined Log = NIL

Data Race Detection at Production-Run Speed [ISCA03]

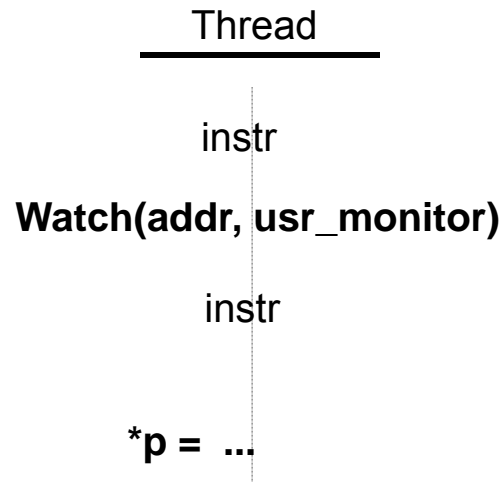
Data race: Two threads access same data without synch



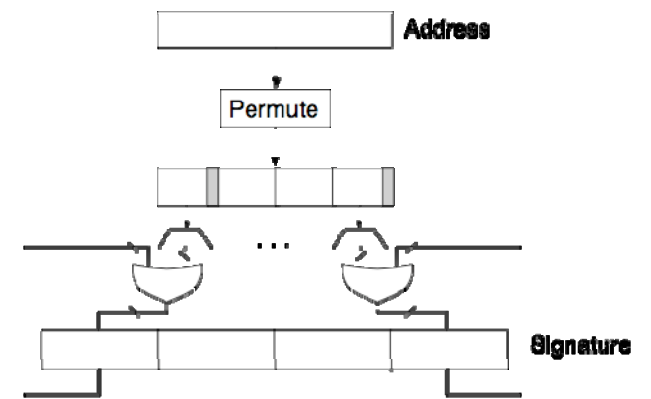
- If we detect communication between...
 - Ordered chunks: not a data race
 - Unordered chunks: data race

Extension: Signatures Visible to SW through ISA

- Enables **pervasive monitoring** [ISCA04]
 - Support numerous watchpoints for free



```
usr_monitor(Addr){  
    .....  
}
```



Extension: Signatures Visible to SW through ISA

- Enables **pervasive monitoring** [ISCA04]
 - Support numerous watchpoints for free
- Enables **novel compiler optimizations** [ASPLOS08]
 - Function memoization
 - Loop-invariant code motion
- Enables **debugging** data races & concurrency bugs [MICRO 09]

Many novel programming/compiler/tool opportunities

Summary: The Bulk Multicore

- 128+ cores/chip, coherent shared-memory (perhaps in groups)
- Simple HW with commodity cores
 - Memory consistency checks moved away from the core
- High performance shared-memory programming model
 - Execution in chunks, possibly driven by the compiler
 - Signatures for disambiguation, cache coherence, and compiler opts
- High programmability:
 - Sequential consistency
 - Sophisticated always-on development support
 - Deterministic replay of parallel programs with no log (*DeLorean*)
 - Data race detection for production runs (*ReEnact*)
 - Pervasive program monitoring (*iWatcher*)
 - Using signatures/hashes to detect races (*SigRace*, *Light64*)

Acknowledgments

Key contributors:

- Luis Ceze
- Calin Cascaval
- James Tuck
- Pablo Montesinos
- Wonsun Ahn
- Milos Prvulovic
- Pin Zhou
- YY Zhou
- Jose Martinez

The Bulk Multicore Architecture for Programmability

Josep Torrellas

Department of Computer Science
University of Illinois at Urbana-Champaign
<http://iacoma.cs.uiuc.edu>



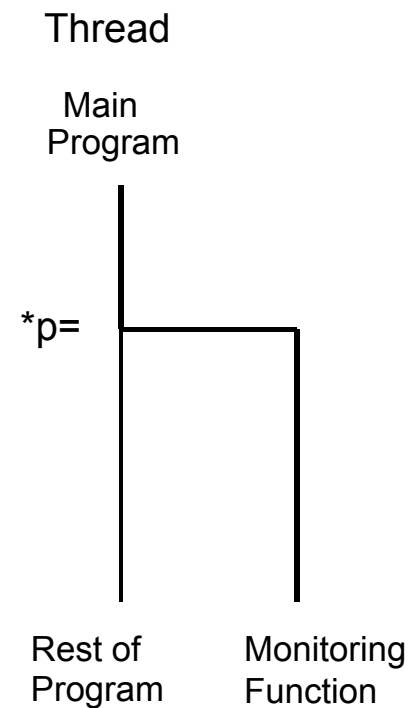
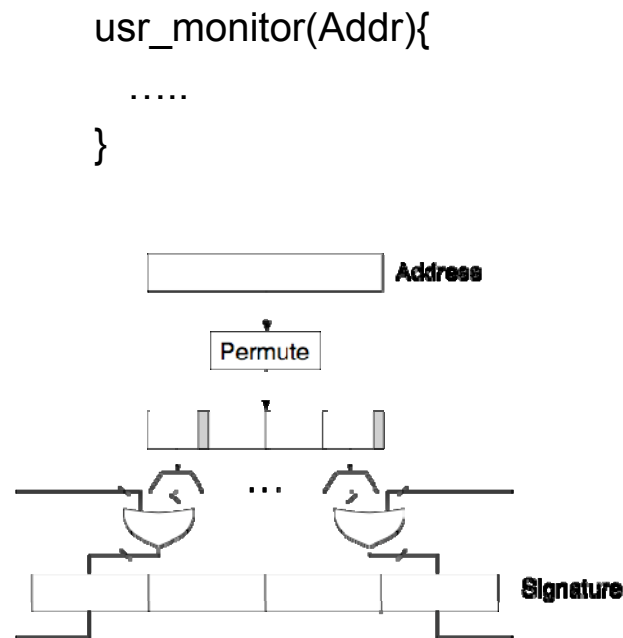
Pervasive Monitoring: Attaching a Monitor Function to Address

- Watch memory location
- Trigger monitoring function when it is accessed

```

instr
Watch(addr, usr_monitor)
instr
instr
instr
*p = ...
instr
instr
instr


```



Chunk Operation + Signatures: Bulk

[ISCA07]

- Execute each chunk **atomically** and in **isolation**
- (Distributed) arbiter ensures a **total order** of chunk commits



- Supports **Sequential Consistency** [Lamport79]:
 - **Low hardware complexity**: Need not snoop ld buffer for consistency
 - **High performance**: Instructions are fully reordered by HWloads and stores make it in any order to the sig
Fences are NOOPS

Rest of the Talk

- The Bulk Multicore
- How it improves performance
- How it improves programmability
- **Extension: Signatures visible to SW through ISA**


Signatures & Hashes Visible to SW through ISA

- Enables **pervasive monitoring** [ISCA04]
 - Support numerous watchpoints for free
- Enables **novel compiler optimizations** [ASPLOS08]
 - Function memoization
 - Loop-invariant code motion
- Enables **debugging** data races & concurrency bugs [MICRO 09]

Many novel programming/compiler/tool opportunities

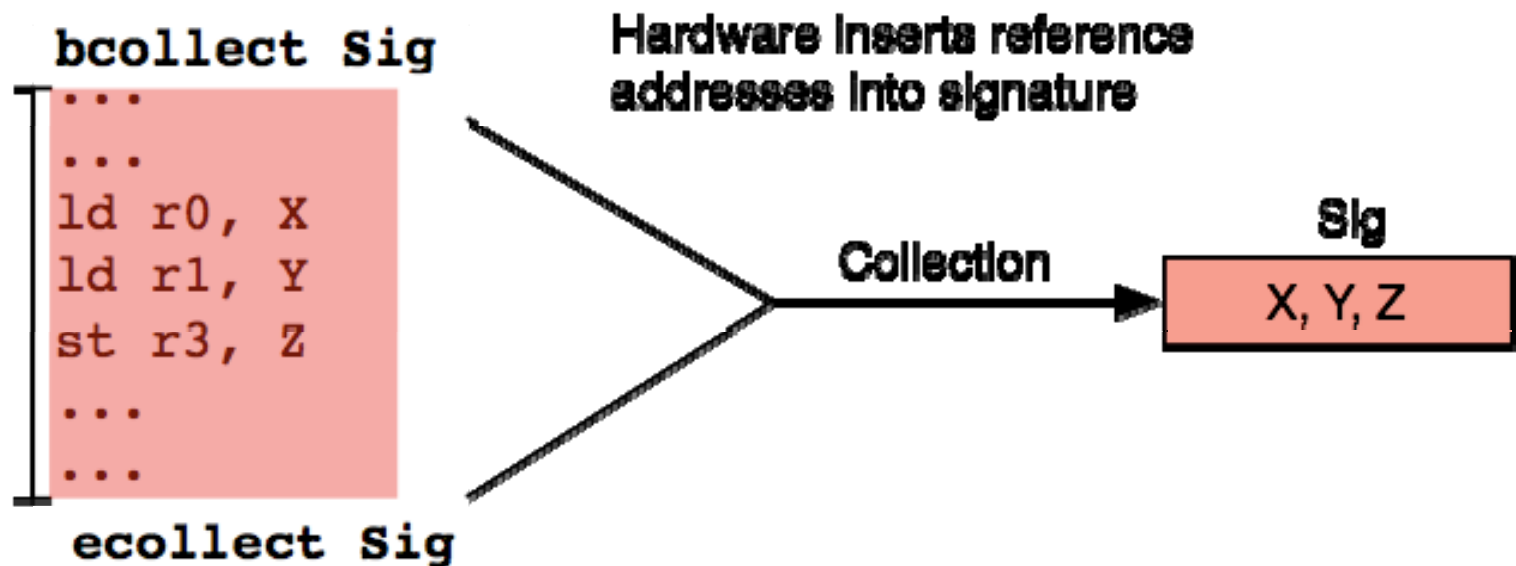
Enabling Novel Compiler Optimizations

New instruction: Begin/End collecting addresses into sig




Enabling Novel Compiler Optimizations


New instruction: Begin/End collecting addresses into sig



Instruction: Begin/End Disambiguation Against Sig



Instruction: Begin/End Disambiguation Against Sig



Optimization: Function Memoization

- Goal: skip the execution of functions


```
foo(x);
```

```
... = y  
z = ...
```

```
foo(x);
```

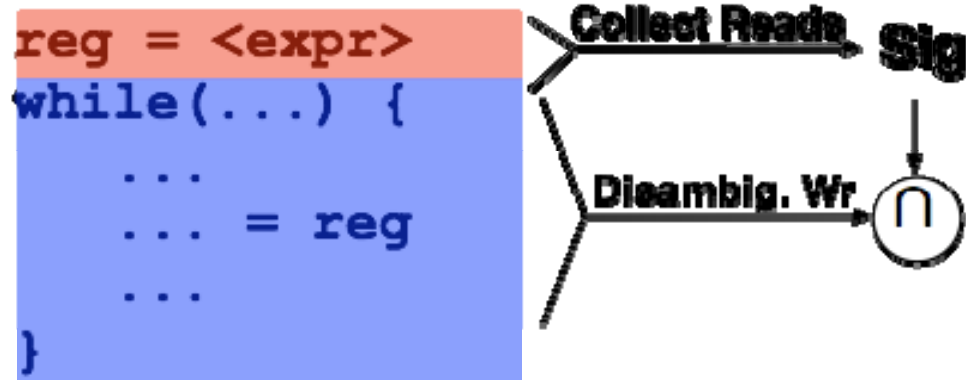
Example Opt: Function Memoization

- Goal: skip the execution of functions whose outputs are known



Example Opt: Loop-Invariant Code Motion

```
while(...) {  
    ...  
    ... = <expr>  
    ...  
}
```



Example Opt: Loop-Invariant Code Motion

```
while(...) {  
    ...  
    ... = <expr>  
    ...  
}
```

checkpoint()

reg = <expr>

while(...) {

...

... = reg


...

}

if (conflict)

rollback()

<original loop>



Different Synchronization Ops

