

# NON-SPECULATIVE LOAD→LOAD REORDERING IN TSO<sup>1</sup>

**Alberto Ros**

Universidad de Murcia

October 17th, 2017

<sup>1</sup> A. Ros, T. E. Carlson, M. Alipour, and S. Kaxiras, "Non-Speculative Load-Load Reordering in TSO". ISCA, 2017.

# OUTLINE

1 BACKGROUND

2 WRITERSBLOCK

3 RESULTS

4 CONCLUSIONS

# PROGRAM ORDER

- Programmer intuition: instructions execute in the order they appear in the program

## THREAD 1

```
$r0 = X; // load  
$r1 = Y; // load
```

# PROGRAM ORDER

- Programmer intuition: instructions execute in the order they appear in the program

## THREAD 1

```
$r0 = X; // load  
$r1 = Y; // load
```

- What happens if the order changes?

## THREAD 1

```
$r1 = Y; // load  
$r0 = X; // load
```

# PROGRAM ORDER

- Programmer intuition: instructions execute in the order they appear in the program

## THREAD 1

```
$r0 = X; // load
$r1 = Y; // load
```

## THREAD 2

```
Y = 1; // store
X = 1; // store
```

- What happens if the order changes?

## THREAD 1

```
$r1 = Y; // load
$r0 = X; // load
```

## THREAD 2

```
Y = 1; // store
X = 1; // store
```

# TOTAL STORE ORDER (TSO)

- The memory consistency model defines the behavior of the programs
  - In particular, the behavior of the memory operations: load and store

# TOTAL STORE ORDER (TSO)

- The memory consistency model defines the behavior of the programs
  - In particular, the behavior of the memory operations: load and store
- x86 processors (Intel, AMD) implement a Total Store Order (TSO)

## TSO RULES

- load → load
- store → store
- load → store

# TOTAL STORE ORDER (TSO)

- The memory consistency model defines the behavior of the programs
  - In particular, the behavior of the memory operations: load and store
- x86 processors (Intel, AMD) implement a Total Store Order (TSO)

## TSO RULES

- load → load
- store → store<sup>2</sup>
- load → store

<sup>2</sup> A. Ros and S. Kaxiras, “Racer: TSO Consistency via Race Detection”, MICRO, 2016.



# TOTAL STORE ORDER (TSO)

- The memory consistency model defines the behavior of the programs
  - In particular, the behavior of the memory operations: load and store
- x86 processors (Intel, AMD) implement a Total Store Order (TSO)

## TSO RULES

- load→load
- store→store<sup>2</sup>
- load→store

- This talk focus on the **load→load** order

<sup>2</sup> A. Ros and S. Kaxiras, “Racer: TSO Consistency via Race Detection”, MICRO, 2016.

# POSSIBLE RESULTS UNDER LOAD → LOAD (E.G. TSO)

INITIALLY  $X=0, Y=0$

```
lx: $r0 = X;
```

```
ly: $r1 = Y;
```

```
sy: Y = 1;
```

```
sx: X = 1;
```

# POSSIBLE RESULTS UNDER LOAD→LOAD (E.G. TSO)

INITIALLY X=0, Y=0

lx: \$r0 = X;  
ly: \$r1 = Y;

sy: Y = 1;  
sx: X = 1;

## SIX POSSIBLE INTERLEAVINGS AND VALUES FOR (\$R0, \$R1)

lx ly	lx sy	lx sy	lx ly	sy lx	sy ly	sy sx
sy sx	ly sx	sy sx	sy sx	lx sx	lx ly	lx ly
(0,0)	(0,1)	(0,1)	(0,1)	(0,1)	(0,1)	(1,1)

- (1,0) is not possible if load→load & store→store

## RELAXING LOAD→LOAD

INITIALLY X=0, Y=0

lx: \$r0 = X;	sy: Y = 1;
ly: \$r1 = Y;	sx: X = 1;

## SIX POSSIBLE INTERLEAVINGS AND VALUES FOR (\$R0, \$R1)

ly lx	ly sy	ly sy	ly lx	sy lx	sy ly	sy sx
sy sx	lx sx	sx lx	sx	sx	sx	ly lx
(0,0)	(0,0)	(1,0)	(0,1)	(1,1)	(1,1)	(1,1)

- (1,0) is possible by relaxing load→load

# LOAD → LOAD REORDERING

- **Waiting** for a load to finish to start the execution of the next load is very **inefficient**

# LOAD → LOAD REORDERING

- **Waiting** for a load to finish to start the execution of the next load is very **inefficient**
- High-performance processors execute multiple load operations simultaneously
  - Memory level parallelism
- Load operations can execute out of order
- This is correct for **single-core** processors

# LOAD→LOAD REORDERING

- Executing load operations out of order can **relax** the **load→load order**

# LOAD→LOAD REORDERING

- Executing load operations out of order can **relax** the **load→load order**

Program

ld<sub>1</sub>

ld<sub>2</sub>

ld<sub>3</sub>

ld<sub>4</sub>

ld<sub>5</sub>

ld<sub>6</sub>





# LOAD→LOAD REORDERING

- Executing load operations out of order can **relax** the **load→load order**

Program

ld<sub>1</sub>

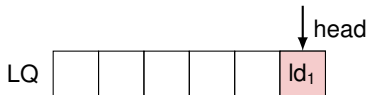
ld<sub>2</sub>

ld<sub>3</sub>

ld<sub>4</sub>

ld<sub>5</sub>

ld<sub>6</sub>



# LOAD→LOAD REORDERING

- Executing load operations out of order can **relax** the **load→load order**

Program

ld<sub>1</sub>

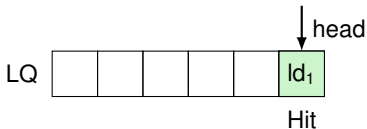
ld<sub>2</sub>

ld<sub>3</sub>

ld<sub>4</sub>

ld<sub>5</sub>

ld<sub>6</sub>

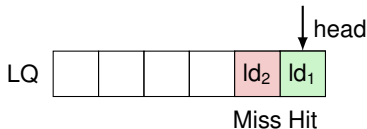


# LOAD→LOAD REORDERING

- Executing load operations out of order can **relax** the **load→load order**

Program

ld<sub>1</sub>  
ld<sub>2</sub>  
ld<sub>3</sub>  
ld<sub>4</sub>  
ld<sub>5</sub>  
ld<sub>6</sub>

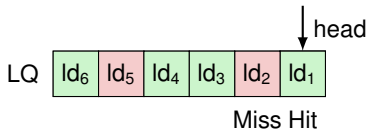


# LOAD→LOAD REORDERING

- Executing load operations out of order can **relax** the **load→load order**

Program

ld<sub>1</sub>  
ld<sub>2</sub>  
ld<sub>3</sub>  
ld<sub>4</sub>  
ld<sub>5</sub>  
ld<sub>6</sub>



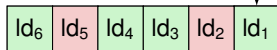
## LOAD→LOAD REORDERING

- Executing load operations out of order can **relax** the **load→load order**

Program

ld<sub>1</sub>  
ld<sub>2</sub>  
ld<sub>3</sub>  
ld<sub>4</sub>  
ld<sub>5</sub>  
ld<sub>6</sub>

LQ



Miss Hit

Order

ld<sub>1</sub>  
ld<sub>3</sub>  
ld<sub>4</sub>  
ld<sub>6</sub>  
**ld<sub>2</sub>**  
**ld<sub>5</sub>**

# LOAD→LOAD REORDERING

- In multicore processors reordering loads can affect the expected result
- It is necessary to always maintain the load→load order?

# LOAD→LOAD REORDERING

- In multicore processors reordering loads can affect the expected result
- It is necessary to always maintain the load→load order?

INITIALLY X=0, Y=0

```
$r0 = X;      |   Y = 1;  
$r1 = Y;      |   X = 1;
```

*/\* (1,0) not allowed \*/*

# LOAD→LOAD REORDERING

- In multicore processors reordering loads can affect the expected result
- It is necessary to always maintain the load→load order?

## INITIALLY X=0, Y=0

```
$r0 = X;      |   Y = 1;
$r1 = Y;      |   X = 1;
```

*/\* (1,0) not allowed \*/*

## POSSIBLE EXECUTION

```
$r0 = Y;      |
$r1 = X;      |
               |   Y = 1;
               |   X = 1;
```

*/\* (0, 0) allowed \*/*



# LOAD→LOAD REORDERING

- In multicore processors reordering loads can affect the expected result
- It is necessary to always maintain the load→load order?

## INITIALLY X=0, Y=0

```
$r0 = X;      |   Y = 1;
$r1 = Y;      |   X = 1;
```

*/\* (1,0) not allowed \*/*

## POSSIBLE EXECUTION

```
$r0 = Y;      |
               |   Y = 1;
               |   X = 1;
$r1 = X;      |
```

*/\* (1, 0) not allowed \*/*

# LOAD→LOAD REORDERING

- In multicore processors reordering loads can affect the expected result
- It is necessary to always maintain the load→load order?

## INITIALLY X=0, Y=0

```
$r0 = X;      |   Y = 1;
$r1 = Y;      |   X = 1;
```

*/\* (1,0) not allowed \*/*

## POSSIBLE EXECUTION

```
$r0 = Y;      |
               |   Y = 1;
               |   X = 1;
$r1 = X;      |
```

*/\* (1, 0) not allowed \*/*

- No, if the other cores do not see the reordering

## RELAXING LOAD→LOAD

INITIALLY X=0, Y=0

lx: \$r0 = X;	sy: Y = 1;
ly: \$r1 = Y;	sx: X = 1;

## SIX POSSIBLE INTERLEAVINGS AND VALUES FOR (\$R0, \$R1)

ly	ly	ly	sy	sy	sy
lx	sy	sy	ly	ly	lx
sy	lx	sx	lx	sx	ly
sx	sx	lx	sx	lx	lx
(0,0)	(0,0)	(1,0)	(0,1)	(1,1)	(1,1)

- (1,0) is possible by relaxing load→load

# LOAD→LOAD SPECULATION

- Solution: To allow **speculative load→load reordering**
- Some definitions: performed, ordered, source of speculation (SoS)

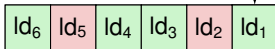
# LOAD→LOAD SPECULATION

- Solution: To allow **speculative load→load reordering**
- Some definitions: performed, ordered, source of speculation (SoS)

Program

ld<sub>1</sub>  
ld<sub>2</sub>  
ld<sub>3</sub>  
ld<sub>4</sub>  
ld<sub>5</sub>  
ld<sub>6</sub>

LQ



Order

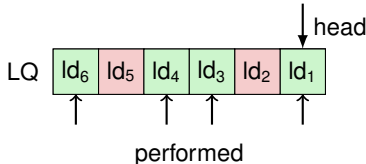
ld<sub>1</sub>  
ld<sub>3</sub>  
ld<sub>4</sub>  
ld<sub>6</sub>  
ld<sub>2</sub>  
ld<sub>5</sub>

# LOAD→LOAD SPECULATION

- Solution: To allow **speculative load→load reordering**
- Some definitions: performed, ordered, source of speculation (SoS)

Program

ld<sub>1</sub>  
ld<sub>2</sub>  
ld<sub>3</sub>  
ld<sub>4</sub>  
ld<sub>5</sub>  
ld<sub>6</sub>



Order

ld<sub>1</sub>  
ld<sub>3</sub>  
ld<sub>4</sub>  
ld<sub>6</sub>  
ld<sub>2</sub>  
ld<sub>5</sub>

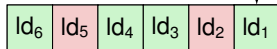
# LOAD→LOAD SPECULATION

- Solution: To allow **speculative load→load reordering**
- Some definitions: performed, ordered, source of speculation (SoS)

Program

ld<sub>1</sub>  
ld<sub>2</sub>  
ld<sub>3</sub>  
ld<sub>4</sub>  
ld<sub>5</sub>  
ld<sub>6</sub>

LQ



M-spec

Order

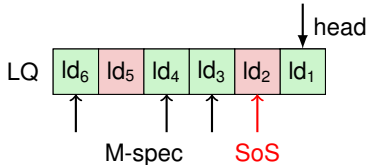
ld<sub>1</sub>  
ld<sub>3</sub>  
ld<sub>4</sub>  
ld<sub>6</sub>  
ld<sub>2</sub>  
ld<sub>5</sub>

# LOAD→LOAD SPECULATION

- Solution: To allow **speculative load→load reordering**
- Some definitions: performed, ordered, source of speculation (SoS)

Program

ld<sub>1</sub>  
ld<sub>2</sub>  
ld<sub>3</sub>  
ld<sub>4</sub>  
ld<sub>5</sub>  
ld<sub>6</sub>



Order

ld<sub>1</sub>  
ld<sub>3</sub>  
ld<sub>4</sub>  
ld<sub>6</sub>  
ld<sub>2</sub>  
ld<sub>5</sub>

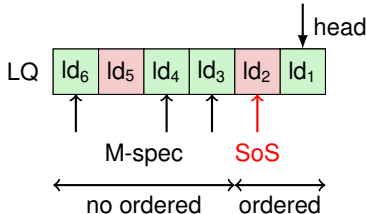


# LOAD→LOAD SPECULATION

- Solution: To allow **speculative load→load reordering**
- Some definitions: performed, ordered, source of speculation (SoS)

Program

ld<sub>1</sub>  
ld<sub>2</sub>  
ld<sub>3</sub>  
ld<sub>4</sub>  
ld<sub>5</sub>  
ld<sub>6</sub>



Order

ld<sub>1</sub>  
ld<sub>3</sub>  
ld<sub>4</sub>  
ld<sub>6</sub>  
ld<sub>2</sub>  
ld<sub>5</sub>

# SQUASH AND RE-EXECUTE UPON INVALIDATION

- Current multicore avoid incorrect results
  - With the help of the cache coherence protocol

ly

sy

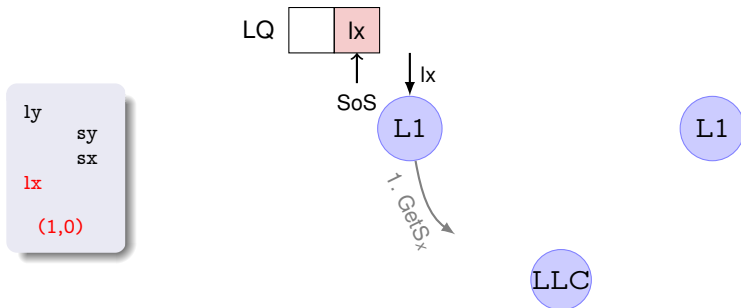
sx

lx

(1,0)

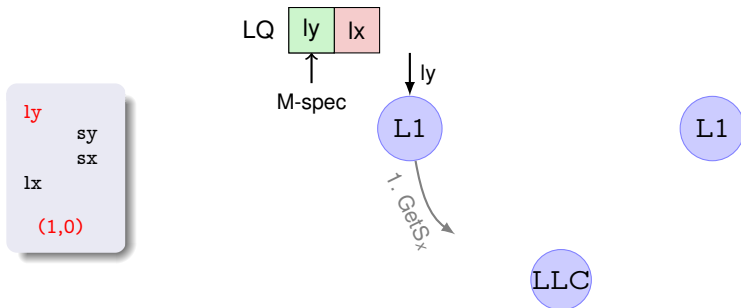
# SQUASH AND RE-EXECUTE UPON INVALIDATION

- Current multicore avoid incorrect results
  - With the help of the cache coherence protocol



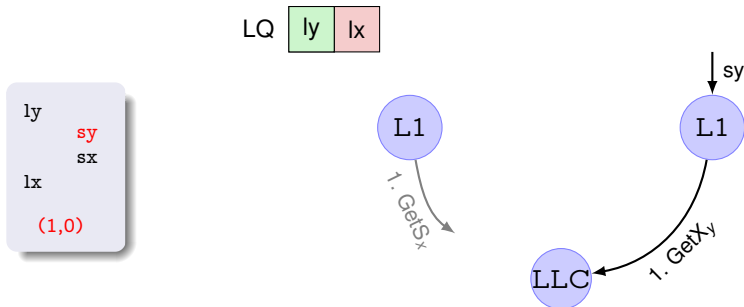
# SQUASH AND RE-EXECUTE UPON INVALIDATION

- Current multicore avoid incorrect results
  - With the help of the cache coherence protocol



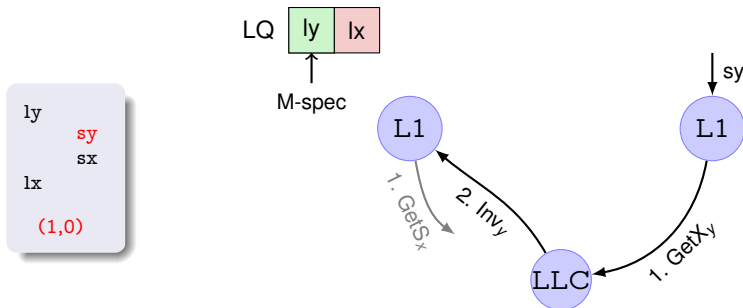
# SQUASH AND RE-EXECUTE UPON INVALIDATION

- Current multicore avoid incorrect results
  - With the help of the cache coherence protocol



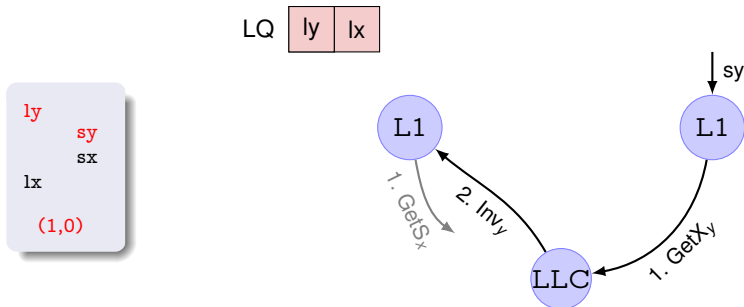
# SQUASH AND RE-EXECUTE UPON INVALIDATION

- Current multicore avoid incorrect results
  - With the help of the cache coherence protocol



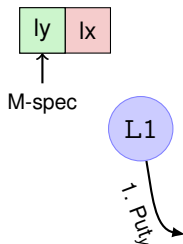
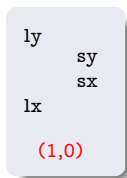
# SQUASH AND RE-EXECUTE UPON INVALIDATION

- Current multicore avoid incorrect results
  - With the help of the cache coherence protocol
  - Squashing and re-executing on remote writes



# SQUASH AND RE-EXECUTE UPON EVICTIONS

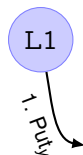
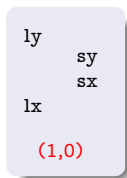
- What happens when a cache block loaded by an M-spec load is evicted?
  - If the directory stops tracking the block, the M-spec load will not receive an invalidation





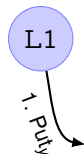
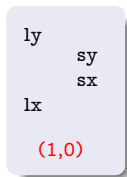
# SQUASH AND RE-EXECUTE UPON EVICTIONS

- What happens when a cache block loaded by an M-spec load is evicted?
  - If the directory stops tracking the block, the M-spec load will not receive an invalidation



# SQUASH AND RE-EXECUTE UPON EVICTIONS

- What happens when a cache block loaded by an M-spec load is evicted?
  - If the directory stops tracking the block, the M-spec load will not receive an invalidation
- Solution: Squash and re-execute upon evictions
  - This impacts the performance of **sequential** applications!!!



# PROBLEMS OF SPECULATION

- Memory-related speculation is the **current solution** to issue loads out of order while keeping the load→load order

# PROBLEMS OF SPECULATION

- Memory-related speculation is the **current solution** to issue loads out of order while keeping the load→load order
- Why is **good**?
  - Because squashing is not frequent!

# PROBLEMS OF SPECULATION

- Memory-related speculation is the **current solution** to issue loads out of order while keeping the load→load order
- Why is **good**?
  - Because squashing is not frequent!
- Why is **bad**?
  - Because speculative loads hold critical resources (LQ, RoB)
  - Because the processor needs to keep continuously the rollback path

# PROBLEMS OF SPECULATION

- Memory-related speculation is the **current solution** to issue loads out of order while keeping the load→load order
- Why is **good**?
  - Because squashing is not frequent!
- Why is **bad**?
  - Because speculative loads hold critical resources (LQ, RoB)
  - Because the processor needs to keep continuously the rollback path

## QUESTION

Can we execute loads out of order without speculation and guaranteeing load→load?

# OUTLINE

1 BACKGROUND

2 WRITERSBLOCK

3 RESULTS

4 CONCLUSIONS

# WRITERSBLOCK ACHIEVEMENT

- Current multicore processors speculatively execute loads out of order
- If a conflict happens, loads are squashed and re-executed



# WRITERSBLOCK ACHIEVEMENT

- Current multicore processors speculatively execute loads out of order
- If a conflict happens, loads are squashed and re-executed

## WRITERSBLOCK

Makes possible **removing this speculation**, executing loads out of order, and making that an executed load is **never squashed** because of the consistency model

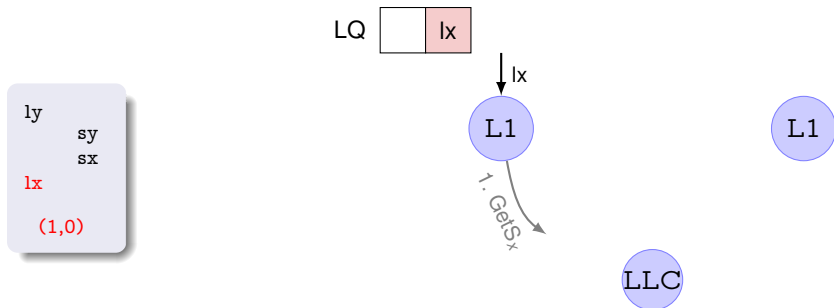
# How?

- With the help of the cache coherence protocol

ly  
sy  
sx  
lx  
  
(1,0)

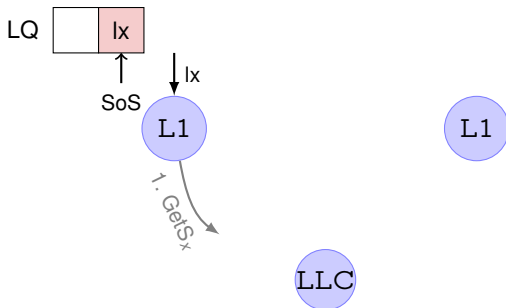
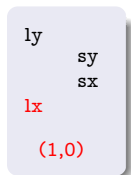
# How?

- With the help of the cache coherence protocol



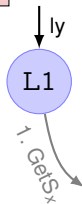
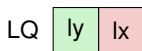
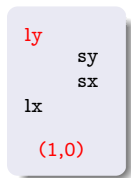
# How?

- With the help of the cache coherence protocol



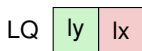
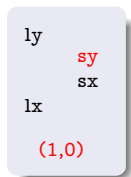
# How?

- With the help of the cache coherence protocol

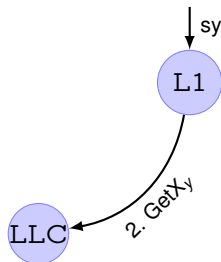
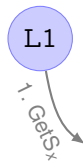


# How?

- With the help of the cache coherence protocol

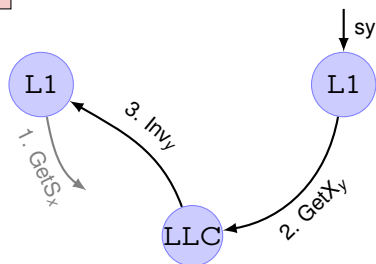
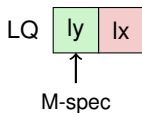
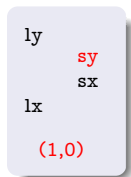


M-spec



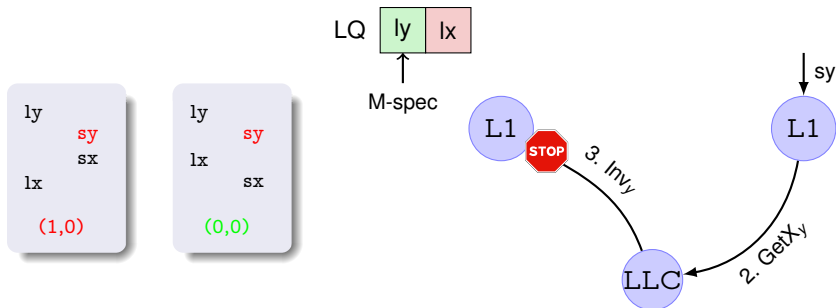
# How?

- With the help of the cache coherence protocol



# How?

- With the help of the cache coherence protocol
  - Blocking and **delaying** the remote write (WritersBlock)



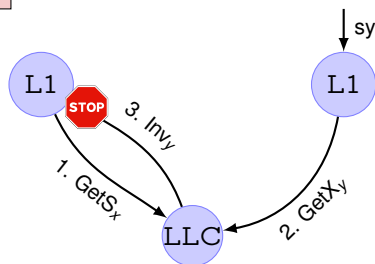
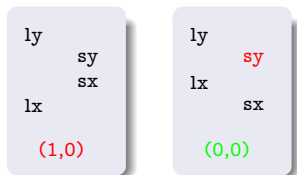


# How?

- With the help of the cache coherence protocol
  - Blocking and **delaying** the remote write (WritersBlock)
  - Until when?** Until the load stop being M-spec

LQ 

ly	lx
----	----

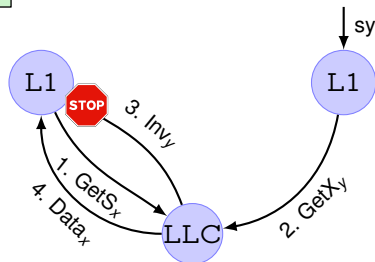
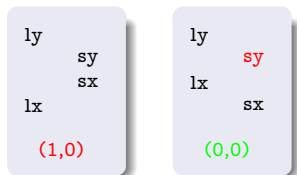


# How?

- With the help of the cache coherence protocol
  - Blocking and **delaying** the remote write (WritersBlock)
  - Until when?** Until the load stop being M-spec

LQ 

ly	lx
----	----

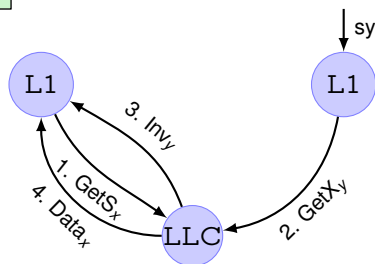
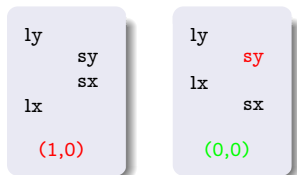


# How?

- With the help of the cache coherence protocol
  - Blocking and **delaying** the remote write (WritersBlock)
  - Until when?** Until the load stop being M-spec

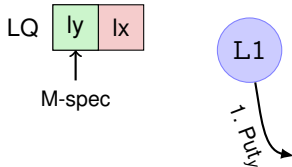
LQ 

ly	lx
----	----



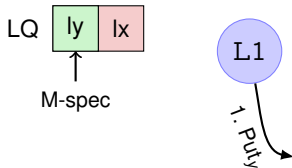
# EVICIONS

- What happens upon an eviction? Do we squash loads?



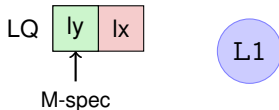
# EVICCTIONS

- What happens upon an eviction? Do we squash loads?
  - No, just need to guarantee that the invalidation will arrive upon a remote write



# EVICCTIONS

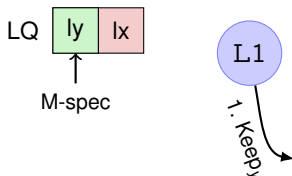
- What happens upon an eviction? Do we squash loads?
  - No, just need to guarantee that the invalidation will arrive upon a remote write
- Solution:
  - Clean blocks implement **silent** evictions<sup>3</sup>



<sup>3</sup> R. Fernandez-Pascual, A. Ros, and M. E. Acacio, “To Be Silent or Not: On the Impact of Evictions of Clean Data in Cache-Coherent Multicores”, Journal of Supercomputing, 2017.

# EVICCTIONS

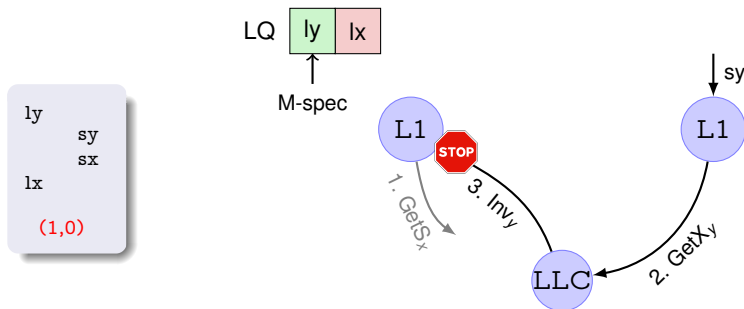
- What happens upon an eviction? Do we squash loads?
  - No, just need to guarantee that the invalidation will arrive upon a remote write
- Solution:
  - Clean blocks implement **silent** evictions<sup>3</sup>
  - Dirty blocks write back the data but **the directory still keeps track**



<sup>3</sup> R. Fernandez-Pascual, A. Ros, and M. E. Acacio, "To Be Silent or Not: On the Impact of Evictions of Clean Data in Cache-Coherent Multicores", Journal of Supercomputing, 2017.

# DEADLOCK

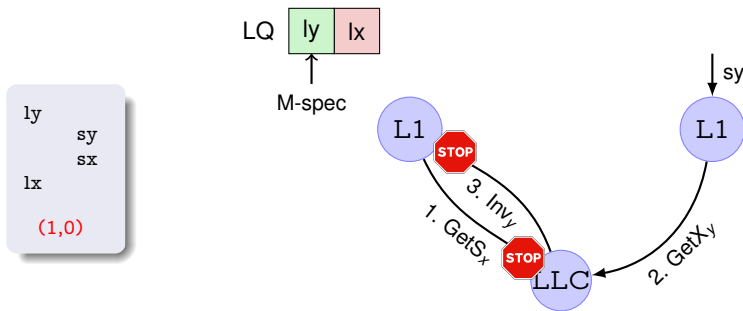
- Blocking writes can cause deadlocks
  - If  $x$  and  $y$  are two words within the same cache line





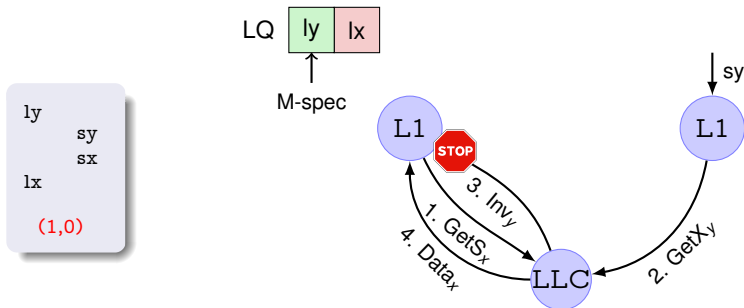
# DEADLOCK

- Blocking writes can cause deadlocks
  - If  $x$  and  $y$  are two words within the same cache line



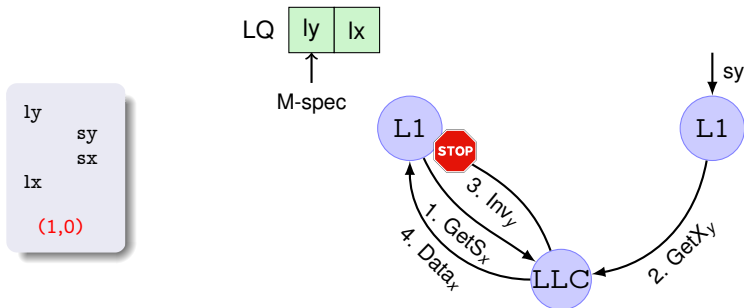
# DEADLOCK

- Blocking writes can cause deadlocks
  - If  $x$  and  $y$  are two words within the same cache line
  - Solution:** Blocked writes allow reads to be resolved



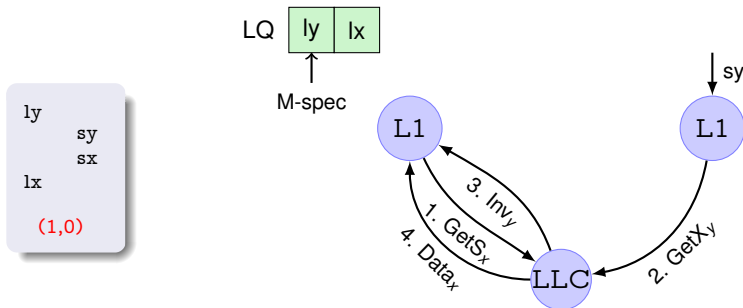
# DEADLOCK

- Blocking writes can cause deadlocks
  - If  $x$  and  $y$  are two words within the same cache line
  - Solution:** Blocked writes allow reads to be resolved



# DEADLOCK

- Blocking writes can cause deadlocks
  - If  $x$  and  $y$  are two words within the same cache line
  - Solution:** Blocked writes allow reads to be resolved



# LIVELOCK

- Resolving reads while blocking writes can cause livelock
  - Resolving a read once the data is invalidated will cause a second invalidation

# LIVELOCK

- Resolving reads while blocking writes can cause livelock
  - Resolving a read once the data is invalidated will cause a second invalidation
- **Solution**
  - Reads resolved through WritersBlock must be non-cacheable
  - and cannot resolve M-spec loads (no invalidation will be received)

# DEADLOCK AVOIDANCE

- **WRITERSBLOCK** cause writes to be blocked
  - Until a load stop being M-speculative
- Deadlocks will not happen if loads cannot be stopped by a pending write miss
- Other blocking causes:
  - MSHR address occupied by write miss  $\Rightarrow$  Duplicate read-write MSHR allocation
  - Full directory/LLC  $\Rightarrow$  Non-cacheable loads
  - Atomic Read-Modify-Write  $\Rightarrow$  Non-speculative

## CASE OF USE: OUT-OF-ORDER COMMIT

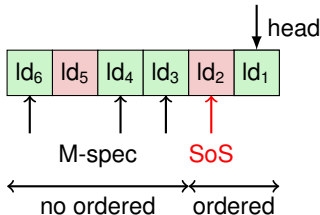
- **Out-of-order commit**<sup>4</sup> allows the processor to retire instructions from the reorder buffer (RoB) even if they are not at the head
- It cannot retire instructions that may squash

<sup>4</sup> G. B. Bell and M. H. Lipasti, "Deconstructing Commit", ISPASS, 2004.



# CASE OF USE: OUT-OF-ORDER COMMIT

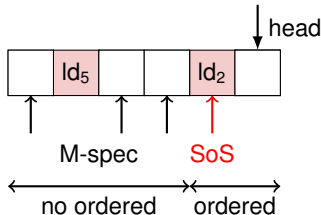
- **Out-of-order commit**<sup>4</sup> allows the processor to retire instructions from the reorder buffer (RoB) even if they are not at the head
- It cannot retire instructions that may squash
- **WRITERSBLOCK** allows the retirement of out-of-order loads
- Better RoB/LQ usage



<sup>4</sup> G. B. Bell and M. H. Lipasti, "Deconstructing Commit", ISPASS, 2004.

# CASE OF USE: OUT-OF-ORDER COMMIT

- **Out-of-order commit**<sup>4</sup> allows the processor to retire instructions from the reorder buffer (RoB) even if they are not at the head
- It cannot retire instructions that may squash
- **WRITERSBLOCK** allows the retirement of out-of-order loads
- Better RoB/LQ usage



<sup>4</sup> G. B. Bell and M. H. Lipasti, "Deconstructing Commit", ISPASS, 2004.

# OUTLINE

1 BACKGROUND

2 WRITERSBLOCK

**3 RESULTS**

4 CONCLUSIONS

# SIMULATION ENVIRONMENT

- Simulator: GEMS + OoO processor (TSO)

# SIMULATION ENVIRONMENT

- Simulator: GEMS + OoO processor (TSO)
- 16-core multicore
- Silvermont (32-entry RoB), Nehalem (128-entry RoB), and Haswell (192-entry RoB)

# SIMULATION ENVIRONMENT

- Simulator: GEMS + OoO processor (TSO)
- 16-core multicore
- Silvermont (32-entry RoB), Nehalem (128-entry RoB), and Haswell (192-entry RoB)
- Benchmarks: Splash-3<sup>5</sup> and Parsec-3.0

<sup>5</sup> C. Sakalis, C. Leonardsson, S. Kaxiras, and A. Ros, “Splash-3: A Properly Synchronized Benchmark Suite for Contemporary Research”, ISPASS, 2016.

# SIMULATION ENVIRONMENT

- Simulator: GEMS + OoO processor (TSO)
- 16-core multicore
- Silvermont (32-entry RoB), Nehalem (128-entry RoB), and Haswell (192-entry RoB)
- Benchmarks: Splash-3<sup>5</sup> and Parsec-3.0
- Protocols
  - **DIRECTORY**: Directory-based MESI protocol
  - **WRITERSBLOCK**: Extensions to **DIRECTORY**

<sup>5</sup> C. Sakalis, C. Leonardsson, S. Kaxiras, and A. Ros, “Splash-3: A Properly Synchronized Benchmark Suite for Contemporary Research”, ISPASS, 2016.

# SIMULATION ENVIRONMENT

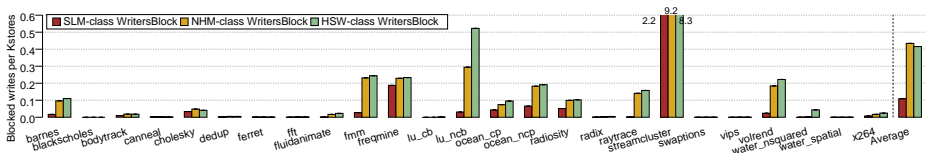
- Simulator: GEMS + OoO processor (TSO)
- 16-core multicore
- Silvermont (32-entry RoB), Nehalem (128-entry RoB), and Haswell (192-entry RoB)
- Benchmarks: Splash-3<sup>5</sup> and Parsec-3.0
- Protocols
  - **DIRECTORY**: Directory-based MESI protocol
  - **WRITERSBLOCK**: Extensions to **DIRECTORY**
- Commit technique
  - **INORDERCOMMIT**
  - **OoOCOMMIT**

<sup>5</sup> C. Sakalis, C. Leonardsson, S. Kaxiras, and A. Ros, “Splash-3: A Properly Synchronized Benchmark Suite for Contemporary Research”, ISPASS, 2016.



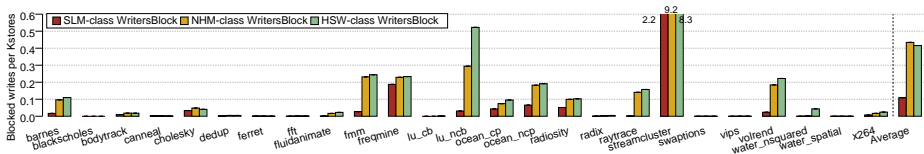
# WRITERSBLOCK: BLOCKED WRITES

- Results for **INORDERCOMMIT**
- Normalized to **DIRECTORY**



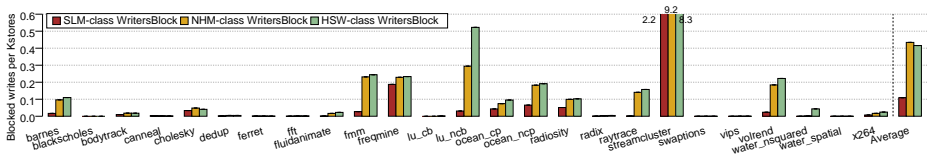
# WRITERSBLOCK: BLOCKED WRITES

- Results for **INORDERCOMMIT**
- Normalized to **DIRECTORY**
- The larger the RoB, the more misspeculations, and the more blocked writes



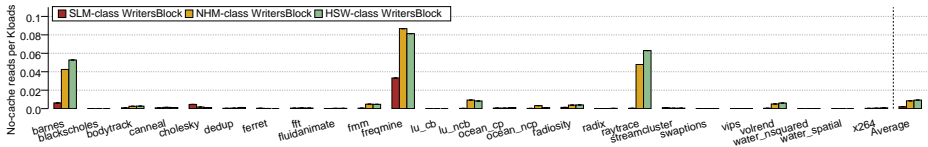
# WRITERSBLOCK: BLOCKED WRITES

- Results for **INORDERCOMMIT**
- Normalized to **DIRECTORY**
- The larger the RoB, the more misspeculations, and the more blocked writes
- Less than 5 blocks per 10,000 stores, on average



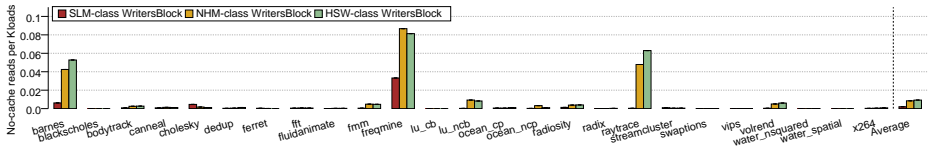
# WRITERSBLOCK: NON-CACHEABLE DATA

- Results for **INORDERCOMMIT**
- Normalized to **DIRECTORY**



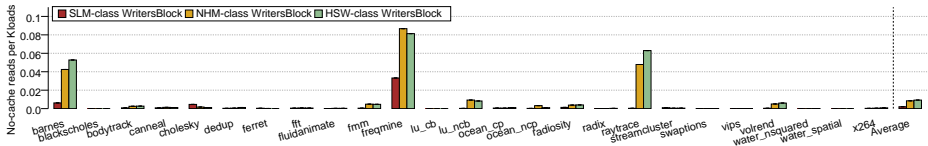
# WRITERSBLOCK: NON-CACHEABLE DATA

- Results for **INORDERCOMMIT**
- Normalized to **DIRECTORY**
- The larger the RoB, the more misspeculations, and the more non-cacheable data



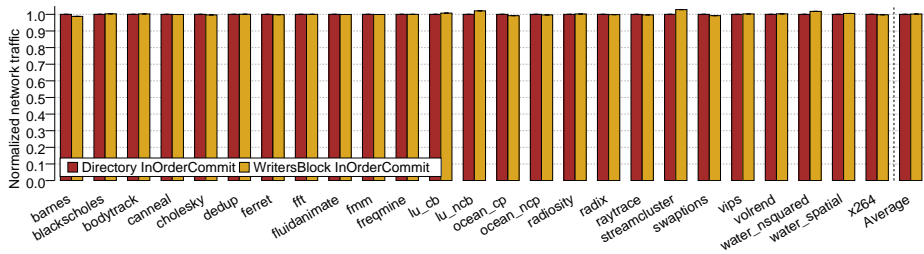
# WRITERSBLOCK: NON-CACHEABLE DATA

- Results for **INORDERCOMMIT**
- Normalized to **DIRECTORY**
- The larger the RoB, the more misspeculations, and the more non-cacheable data
- $\approx 1$  non-cacheable data per 100,000 loads, on average



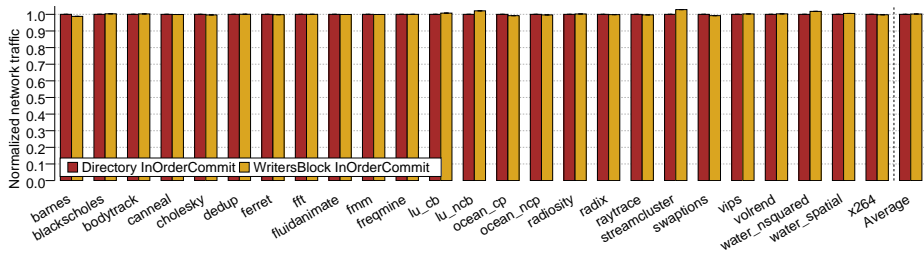
# WRITERSBLOCK: NETWORK TRAFFIC

- Results for **INORDERCOMMIT**
- Normalized to **DIRECTORY**



# WRITERSBLOCK: NETWORK TRAFFIC

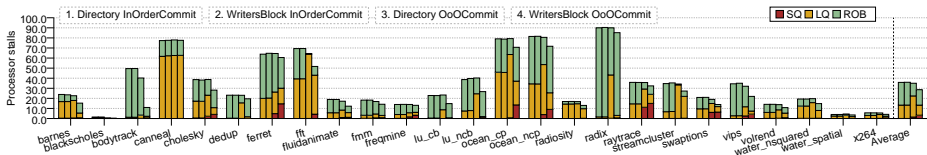
- Results for **INORDERCOMMIT**
- Normalized to **DIRECTORY**
- Network traffic on par





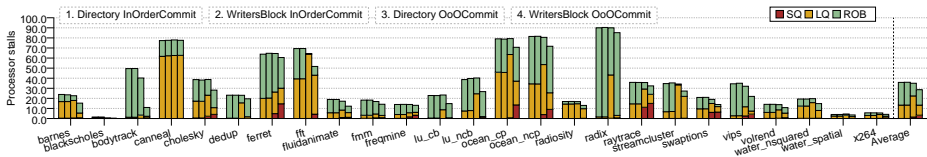
# OUT-OF-ORDER COMMIT: PROCESSOR STALLS

- Normalized to **DIRECTORY + INORDERCOMMIT**



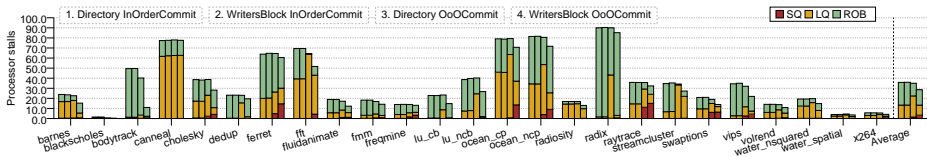
# OUT-OF-ORDER COMMIT: PROCESSOR STALLS

- Normalized to **DIRECTORY + INORDERCOMMIT**
- INORDERCOMMIT**
  - WRITERSBLOCK** does not increase SQ stalls



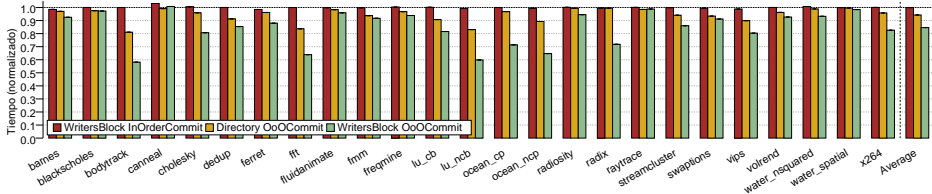
# OUT-OF-ORDER COMMIT: PROCESSOR STALLS

- Normalized to **DIRECTORY + INORDERCOMMIT**
- **INORDERCOMMIT**
  - **WRITERSBLOCK** does not increase SQ stalls
- **OoOCOMMIT**
  - **WRITERSBLOCK** reduces RoB and LQ stalls on average respect to **DIRECTORY**



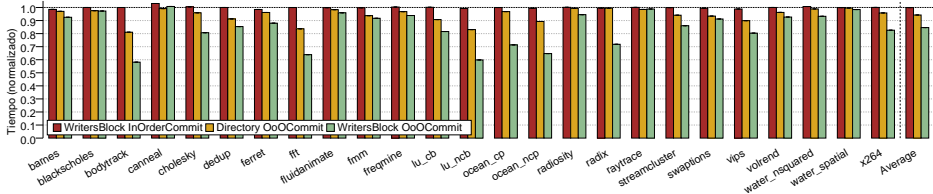
# OUT-OF-ORDER COMMIT: EXECUTION TIME

- Normalized to DIRECTORY + INORDERCOMMIT



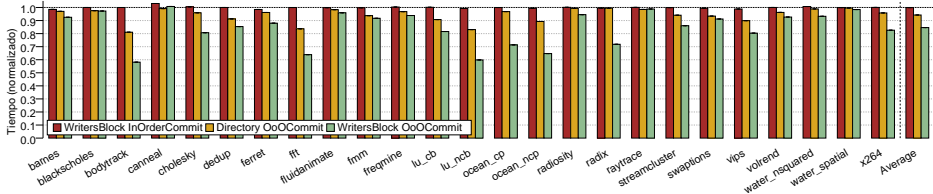
# OUT-OF-ORDER COMMIT: EXECUTION TIME

- Normalized to DIRECTORY + INORDERCOMMIT
- INORDERCOMMIT
  - WRITERSBLOCK does not harm performance on average respect to DIRECTORY



# OUT-OF-ORDER COMMIT: EXECUTION TIME

- Normalized to **DIRECTORY + INORDERCOMMIT**
- **INORDERCOMMIT**
  - **WRITERSBLOCK** does not harm performance on average respect to **DIRECTORY**
- **OoOCOMMIT**
  - **WRITERSBLOCK** improves performance by **11%** on average respect to **DIRECTORY**



# OUTLINE

1 BACKGROUND

2 WRITERSBLOCK

3 RESULTS

4 CONCLUSIONS

# CONCLUSIONS

With the help of the cache coherence protocol,  
and without harming performance,  
we can **execute** loads **out of order** and **without speculation**,  
and obtaining results as if the loads were executed in order  
(**LOAD**→**LOAD**)



# CONCLUSIONS

With the help of the cache coherence protocol,  
and without harming performance,  
we can **execute** loads **out of order** and **without speculation**,  
and obtaining results as if the loads were executed in order  
(**LOAD**→**LOAD**)

Non-speculative loads can increase performance of  
out-of-order commit by **11%**

# NON-SPECULATIVE LOAD→LOAD REORDERING IN TSO<sup>1</sup>

**Alberto Ros**

Universidad de Murcia

October 17th, 2017

<sup>1</sup> A. Ros, T. E. Carlson, M. Alipour, and S. Kaxiras, "Non-Speculative Load-Load Reordering in TSO". ISCA, 2017.