

Blocking versus Non-Blocking Shared-Memory Multicore Synchronization: Programmability, Scalability and Performance

Shinhyung Yang¹, Seongho Jeong¹, Byunguk Min¹, Yeonsoo Kim¹,
Bernd Burgstaller¹ and Johann Blieberger²

¹Department of Computer Science, Yonsei University, Korea

²Institute of Computer Engineering, Automation Systems Group,
TU Wien, Austria

June 13, 2019

Table of Contents

- 1 Introduction
- 2 Lock Elision
- 3 Sequential Consistency
- 4 Concurrent Objects
 - Nonblocking stack: COstack
- 5 Benchmark & Measurements
 - Locks: AdaPO-lock vs. POSIX & C++ Mutexes
 - Blocking Queue: Ada_TAS_MPMC
 - Non-blocking vs. Blocking Queues
 - Progress Guarantees
 - Performance gain of AR over SC
- 6 Conclusions and Future Work

Synchronization

- mutual exclusion locks

Synchronization

- mutual exclusion locks
- Ada's protected objects (POs)

Synchronization

- mutual exclusion locks
- Ada's protected objects (POs)
- Entries and procedures of a PO execute one after another

Synchronization

- mutual exclusion locks
- Ada's protected objects (POs)
- Entries and procedures of a PO execute one after another
- makes it straight-forward for programmers to reason about updates to the shared data encapsulated by a PO

Synchronization

- mutual exclusion locks
- Ada's protected objects (POs)
- Entries and procedures of a PO execute one after another
- makes it straight-forward for programmers to reason about updates to the shared data encapsulated by a PO
- mutual-exclusion property of (highly-contended) **locks** stands in the way to **scalability of parallel programs** on many-core architectures

Synchronization

- mutual exclusion locks
- Ada's protected objects (POs)
- Entries and procedures of a PO execute one after another
- makes it straight-forward for programmers to reason about updates to the shared data encapsulated by a PO
- mutual-exclusion property of (highly-contended) **locks** stands in the way to **scalability of parallel programs** on many-core architectures
- locks do not allow **progress guarantees**, because a task may fail inside a critical section, e.g., by entering an endless loop, and thereby prevent other tasks from accessing shared data

Synchronization

- mutual exclusion locks
- Ada's protected objects (POs)
- Entries and procedures of a PO execute one after another
- makes it straight-forward for programmers to reason about updates to the shared data encapsulated by a PO
- mutual-exclusion property of (highly-contended) **locks** stands in the way to **scalability of parallel programs** on many-core architectures
- locks do not allow **progress guarantees**, because a task may fail inside a critical section, e.g., by entering an endless loop, and thereby prevent other tasks from accessing shared data
- \Rightarrow allow method calls to **overlap** in time

Lock Elision

Protecting shared-data with POs

- A coarsed-grained lock is prone to task serialization
- A fine-grained locking is error-prone and complex
- Lock elision reduces serialization with lock-based code

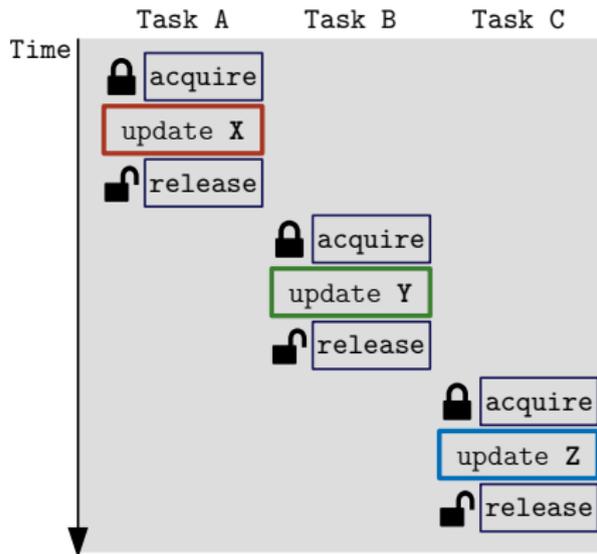


Figure : Task serialization with locks

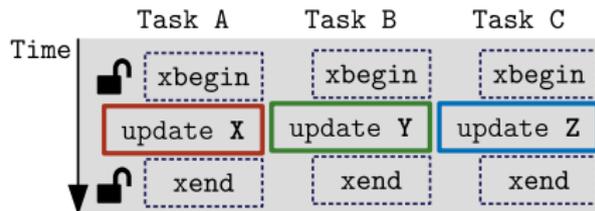


Figure : Lock elision with *Intel TSX*

```
1 function xbegin return uint32;  
2 procedure xend;  
3 function xtest return uint32;  
4 procedure xabort;
```

Listing 1: TSX lock elision
intrinsic

Adapting GNU Ada Run-time Library (GNARL)

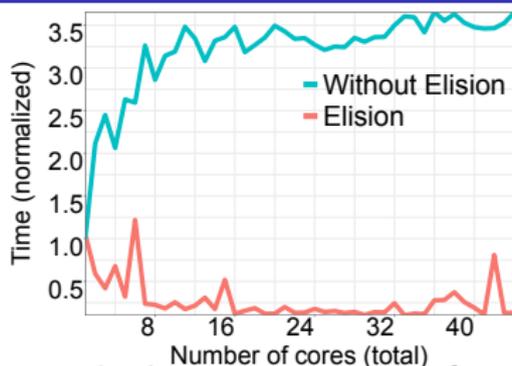
- GNARL employs one POSIX lock per PO for synchronization
- Lock elision is incorporated into Write_Lock

```
1 procedure Write_Lock           -- GNARL lock acquisition procedure
2   result := Try_Elision       -- Attempt lock elision
3   if result = fail then      -- If failed:
4     acquire PO.lock          -- fall-back to acquire POSIX lock
5   end if
6   return
7 end Write_Lock
```

Adapted Write_Lock

- 1 Invoke `Try_Elision`
- 2 If failed:
 - Fall-back to default routine to acquire POSIX lock

Summary: Lock-elision of POs



Ex: hash table

- Pros:
 - PO lock-elision shields programmers from non-blocking synchronization problem
 - high scalability in case of low probability of data-conflicts (e.g., hash-table)
- Cons:
 - requires HW support to be efficient. Not mainstream yet (e.g., not available on ARM platform)
 - Intel TSX requires fallback-path (lock)
 - Intel TSX capacity overflows with large amount of shared data (e.g., linked lists)
 - not generally applicable to all types of data structures

Sequential Consistency

- allow method calls to **overlap** in time

Sequential Consistency

- allow method calls to **overlap** in time
- synchronization on a **finer granularity** within a method's code, via atomic *read-modify-write* (RMW) operations

Sequential Consistency

- allow method calls to **overlap** in time
- synchronization on a **finer granularity** within a method's code, via atomic *read-modify-write* (RMW) operations
- atomic operations are provided either by the CPU's instruction set architecture (ISA), or the language run-time (with the help of the CPU's ISA)

Sequential Consistency

- allow method calls to **overlap** in time
- synchronization on a **finer granularity** within a method's code, via atomic *read-modify-write* (RMW) operations
- atomic operations are provided either by the CPU's instruction set architecture (ISA), or the language run-time (with the help of the CPU's ISA)
- e.g., CAS compare&swap operation

Sequential Consistency

- allow method calls to **overlap** in time
- synchronization on a **finer granularity** within a method's code, via atomic *read-modify-write* (RMW) operations
- atomic operations are provided either by the CPU's instruction set architecture (ISA), or the language run-time (with the help of the CPU's ISA)
- e.g., CAS compare&swap operation
- **sequential consistency** ensures that method calls act as if they occurred in a **sequential, total order** that is **consistent with the program order** of each participating task

Non-blocking Synchronization Techniques

- difficult to implement

Non-blocking Synchronization Techniques

- difficult to implement
- the design of non-blocking data structures is an area of active research

Non-blocking Synchronization Techniques

- difficult to implement
- the design of non-blocking data structures is an area of active research
- a programming language must provide a strict memory model

Lock-free Synchronization – Example

```
1 -- Initial values:
```

```
2 Flag := False;
```

```
3 Data := 0;
```

```
1 -- Task 1:
```

```
2 Data := 1;
```

```
3 Flag := True;
```

Lock-free Synchronization – Example

```
1 -- Initial values:  
2 Flag := False;  
3 Data := 0;
```

```
1 -- Task 1:  
2 Data := 1;  
3 Flag := True;
```

```
1 -- Task 2:  
2 loop  
3   R1 := Flag;  
4   exit when R1;  
5 end loop;  
6 R2 := Data;
```

Lock-free Synchronization – Example

```
1 -- Initial values:
```

```
2 Flag := False;
```

```
3 Data := 0;
```

```
1 -- Task 1:
```

```
2 Data := 1;
```

```
3 Flag := True;
```

```
1 -- Task 2:
```

```
2 loop
```

```
3   R1 := Flag;
```

```
4   exit when R1;
```

```
5 end loop;
```

```
6 R2 := Data;
```

store–store re-ordering of the assignments in lines 2 and 3 of Task 1

⇒ reading R2 = 0 in Line 6 of Task 2.

Lock-free Synchronization – Example

```
1 -- Initial values:
```

```
2 Flag := False;
```

```
3 Data := 0;
```

```
1 -- Task 1:
```

```
2 Data := 1;
```

```
3 Flag := True;
```

```
1 -- Task 2:
```

```
2 loop
```

```
3   R1 := Flag;
```

```
4   exit when R1;
```

```
5 end loop;
```

```
6 R2 := Data;
```

store–store re-ordering of the assignments in lines 2 and 3 of Task 1

⇒ reading R2 = 0 in Line 6 of Task 2.

```
1 Data : Integer with Volatile; -- Ada2012
```

```
2 Flag : Boolean with Atomic; -- Ada2012
```

Ada's Volatile Variables

- guarantee that all tasks agree on the same order of updates

Ada's Volatile Variables

- guarantee that all tasks agree on the same order of updates
- \Rightarrow sequentially consistent

Ada's Volatile Variables

- guarantee that all tasks agree on the same order of updates
- \Rightarrow sequentially consistent
- however: relaxing SC for the sake of performance on contemporary CPU architectures

Concurrent Objects

J. Blieberger, B. Burgstaller, *Safe Non-blocking Synchronization in Ada 202x*,
Ada-Europe'18, 2018, p. 53–69

Concurrent Objects

J. Blieberger, B. Burgstaller, *Safe Non-blocking Synchronization in Ada 202x*,
Ada-Europe'18, 2018, p. 53–69

- support for weak memory model

Concurrent Objects

J. Blieberger, B. Burgstaller, *Safe Non-blocking Synchronization in Ada 202x*,
Ada-Europe'18, 2018, p. 53–69

- support for weak memory model
- for non-blocking synchronization

Concurrent Objects

J. Bliederger, B. Burgstaller, *Safe Non-blocking Synchronization in Ada 202x*, Ada-Europe'18, 2018, p. 53–69

- support for weak memory model
- for non-blocking synchronization
- for synchronization on a finer granularity (RMW operations)
- encapsulation of non-blocking synchronization by high-level language construct
 - what protected objects (POs) do for blocking synchronization

Example – Generic Release-Acquire Object (1/2)

```
1 generic
2   type Data is private;
3 package Generic_Release_Acquire is
4
5   concurrent RA
6   is
7     procedure Write (d: Data);
8     entry Get (D: out Data);
9   private
10    Ready: Boolean := false with Synchronized,
11      Memory_Order_Read => Acquire,
12      Memory_Order_Write => Release;
13    Da: Data;
14  end RA;
15
16 end Generic_Release_Acquire;
```

Example – Generic Release-Acquire Object (2/2)

```
1 package body Generic_Release_Acquire is
2
3   concurrent body RA is
4
5     procedure Write (D: Data) is
6     begin
7       Da := D;
8       Ready := true;
9     end Write;
10
11    entry Get (D: out Data)
12      until Ready is
13      -- spin-lock until released, i.e., Ready = true;
14      -- only sync. variables and constants allowed
15      -- in guard expression
16    begin
17      D := Da;
18    end Get;
19  end RA;
20
21 end Generic_Release_Acquire;
```

Example Lock-free Stack (1/2)

```
1  subtype Data is Integer;
2
3  type List;
4  type List_P is access List;
5  type List is
6      record
7          D: Data;
8          Next: List_P;
9      end record;
10
11 Empty: exception;
12
13 concurrent Lock_Free_Stack
14 is
15     entry Push(D: Data);
16     entry Pop(D: out Data);
17 private
18     Head: List_P with Read_Modify_Write ,
19         Memory_Order_Read => Relaxed ,
20         Memory_Order_Write_Success => Release ,
21         Memory_Order_Write_Failure => Relaxed;
22 end Lock_Free_Stack;
```

Example Lock-free Stack (2/2)

```
1 concurrent body Lock_Free_Stack is
2   entry Push (D: Data)
3     until Head = Head'OLD is
4       New_Node: List_P := new List;
5     begin
6       New_Node.all := (D => D, Next => Head);
7       Head := New_Node; -- RMW
8     end Push;
9
10    entry Pop(D: out Data)
11      until Head = Head'OLD is
12        Old_Head: List_P;
13      begin
14        Old_Head := Head;
15        if Old_Head /= null then
16          Head := Old_Head.Next; -- RMW
17          D := Old_head.D;
18        else
19          raise Empty;
20        end if;
21      end Pop;
22 end Lock_Free_Stack;
```

Benchmark Configuration

- **Platform 1:** 2 CPU Intel Xeon E5-2697 v3 system
 - 14 x86_64 cores per CPU
- **Platform 2:** 4 CPU AWS Graviton on Amazon AWS
 - 4 ARMv8 cores per CPU
- Scalability experiment policies:
 - One Ada task assigned per core
 - Cores of a CPU populated consecutively
 - Once all cores of a CPU are populated, the next CPU receives tasks
- Tasks run synchronization-constructs in **tight-loop**
 - Incurs high contention
 - Brings out the best in each synchronization-construct :)

COstack (1/2)

- **COstack**: Non-blocking stack using C++11 atomic library with weaker memory model
- Function `bool node.compare_exchange_strong(type *expected, type *desired)`
 - **Success**: If the node has not been changed by other threads, then the node is atomically changed to `desired`, and returns `true`
 - **Fail**: If the value of node has changed by other threads, then the `expected` is atomically changed to `node`, and returns `false`
- **Push**: RMW spins in the `while` loop until `head` is changed to `new_node` and returns `true`

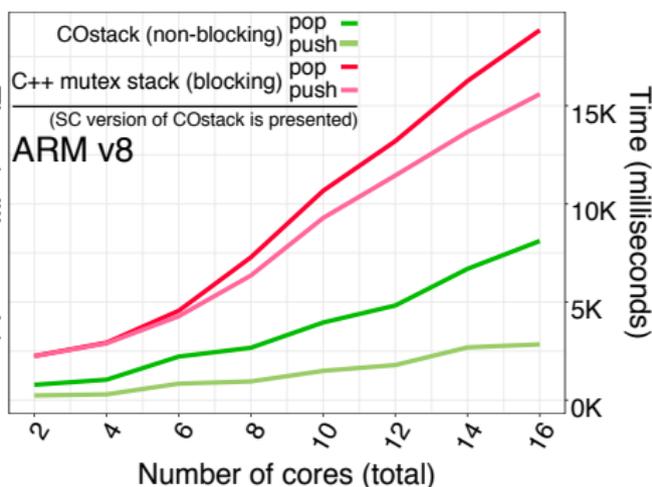
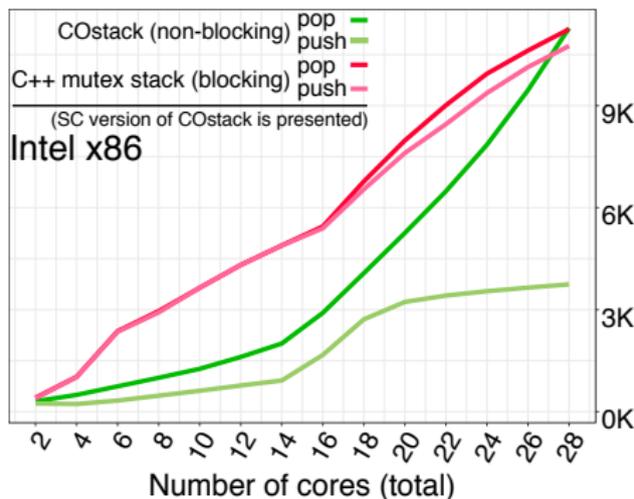
```
1 void push(T const &data) {
2     node *const new_node = new node(data);
3     new_node->next = head.load(std::memory_order_acquire);
4     while (!head.compare_exchange_strong(new_node->next,
5         new_node)) ; // RMW
6 }
```

CStack (2/2)

- **Pop:** `old_head` is in while loop until `old_head` is updated to the latest head variable.

```
1  void pop() {
2      node *old_head = head.load();
3      do {
4          node *temp;
5          do {
6              temp = old_head;
7              old_head = head.load();
8          } while (old_head != temp);
9      } while (old_head &&
10             !head.compare_exchange_strong(old_head, old_head->next))
11             ;
11 }
```

Non-blocking vs. Blocking Stacks



- non-blocking COstack mimics Concurrent Objects proposed for Ada202x
- actually implemented in C++
- COstack performs better than blocking C++ mutex stack
- more performance gains observed on ARMv8

Locks: AdaPO-lock vs. POSIX & C++ Mutexes

- PO-lock using PO's monitor-style synchronization:

```
1 protected body Lock is  
2   procedure CriticalSection is  
3     begin  
4       null; -- Critical section code here...  
5     end CriticalSection;  
6 end Lock;
```

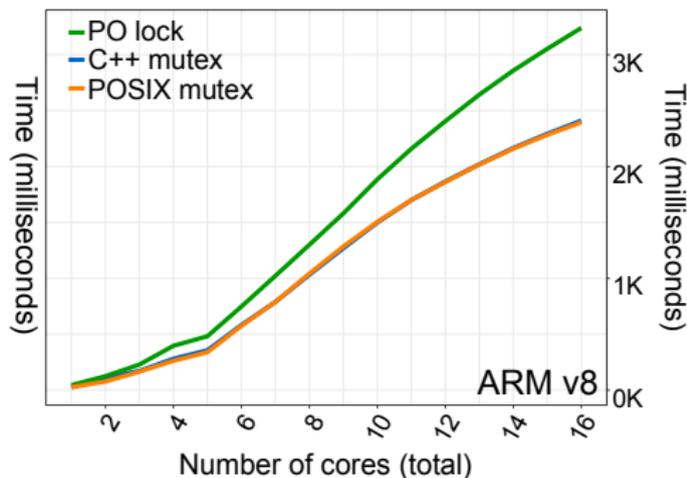
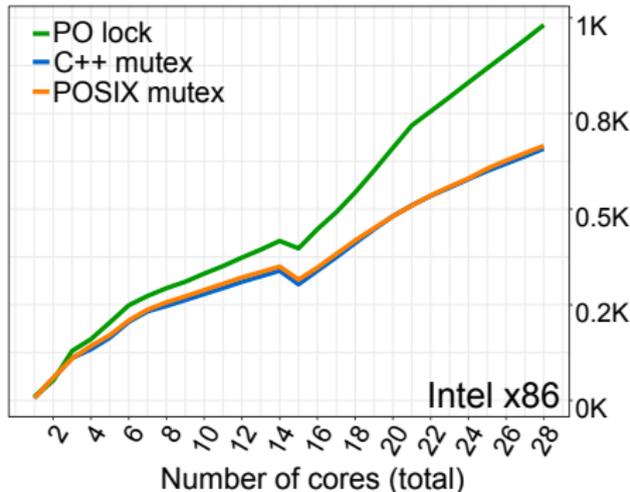
- POSIX mutex consisting of lock-unlock primitives:

```
1 pthread_mutex_t mutex;  
2  
3 pthread_mutex_lock(&mutex);  
4   // Critical section code here...  
5 pthread_mutex_unlock(&mutex);
```

- C++ mutex:

```
1 std::mutex mtx;  
2  
3 mtx.lock();  
4   // Critical section code here...  
5 mtx.unlock();
```

Performance: PO-lock, POSIX & C++ Mutexes



- All three locks provide mutual exclusion but no fairness
 - (The egg-shell model applies only to entries)
- PO adds noticeable performance-overhead compared to “plain” mutexes
 - Programmer-supplied PO configuration mechanism could help to avoid this overhead

Ada_TAS_MPMC (Multi-Producer-Multi-Consumer Queue) (1/2)

- GCC atomic intrinsics are used for synchronization instead of protected objects
- type `__sync_lock_test_and_set_4` (type *ptr, type value)
⇒ It atomically writes value into *ptr, and return the previous contents of *ptr. Size of *ptr and value is 4 byte each.

```
1  function TAS_4(  
2      current: access Unsigned;  
3      Newval: Unsigned) return Unsigned;  
4  pragma Import (Intrinsic, TAS_4,  
5      "__sync_lock_test_and_set_4");
```

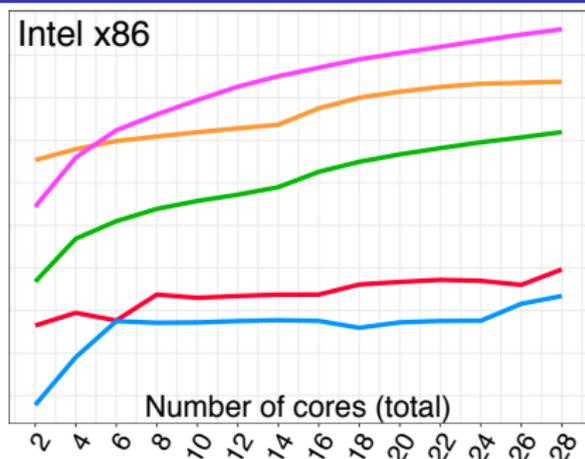
- `__sync_lock_release_4` (type *ptr)
⇒ It releases the lock acquired by `__sync_lock_test_and_set`.
It atomically writes constant 0 to *ptr

Ada_TAS_MPMC (2/2)

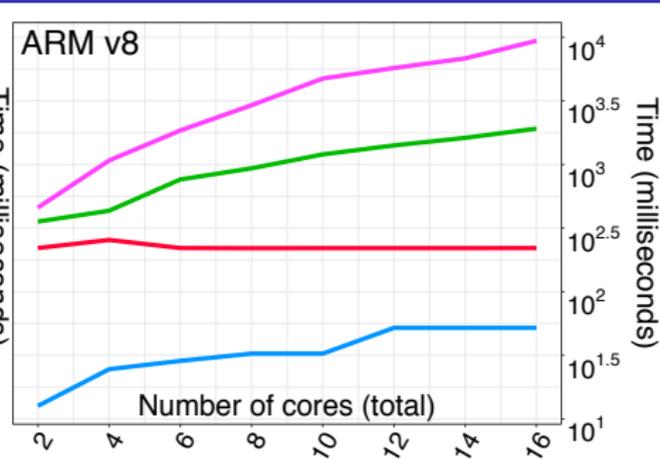
- **Lock & Unlock:** Q.Locked is initialized as 0. By TAS_4 on Q.Locked, lock is acquired. \Rightarrow By writing 0 to Q.Locked, lock is released.
- **Producer:** It enqueues to a queue Q for iter times.
- **Consumer:** It dequeues element from Q to Data for iter times. If Q is empty, release the lock, and redo inner loop.

```
1  Q: My_USQ.Queue;
2  task body Producer is
3  begin
4      for I in 1..iter loop
5          while TAS_4(Q.Locked' Access, 1) = 1 loop -- Acquire
6              null;
7          end loop;
8
9          Q.Enqueue(New_Item => i);
10
11         Lock_Release_4(Q.Locked' Access);          -- Release
12     end loop;
```

Non-blocking vs. Blocking Queues



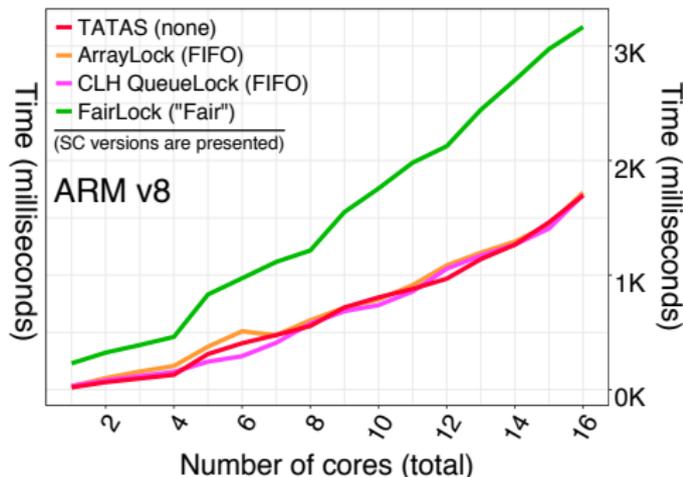
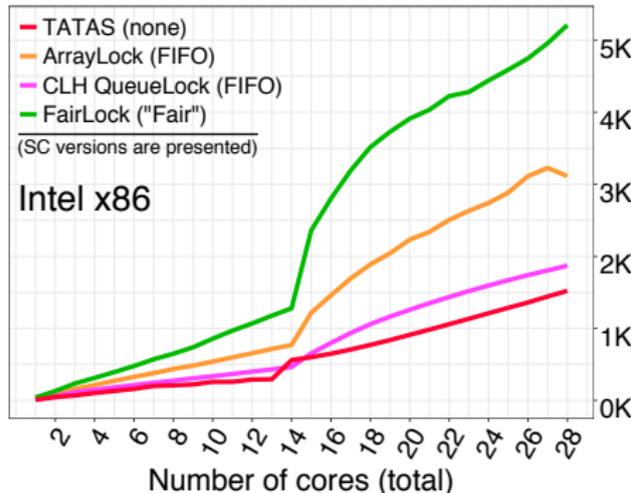
— Ada LockFree MPMC — Boost MPMC
— Ada Synchronized MPMC — Boost SPSC
— B-Queue SPSC



— Ada LockFree MPMC — Boost MPMC
— Ada Synchronized MPMC — Boost SPSC
— B-Queue SPSC

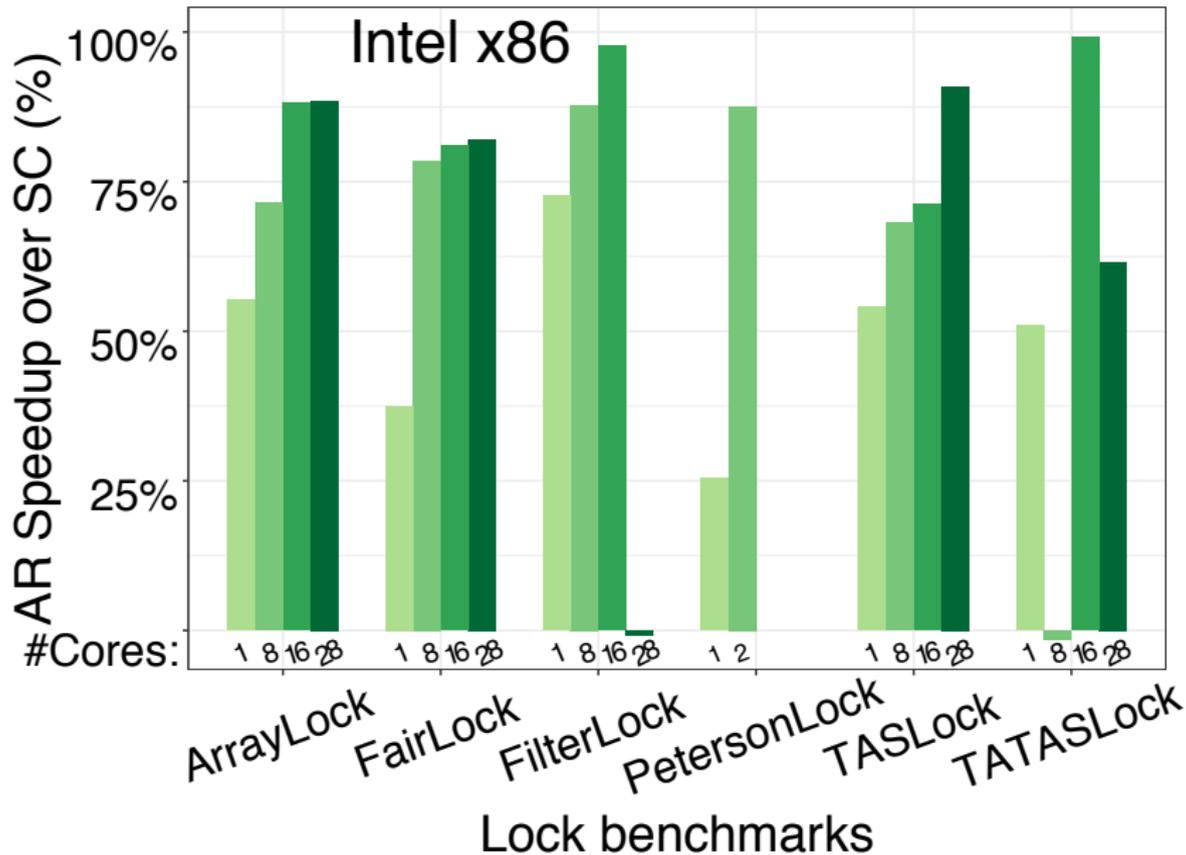
- SPSC-queues: lower sync. overhead, higher performance
- Boost MPMC and Ada LockFree MPMC constitute identical concurrent data-structures (Algorithm of Michael & Scott)
- Blocking MPMCs (Ada Synchronized, Ada TAS): clear performance disadvantage

Non-blocking Locks and Progress Guarantees

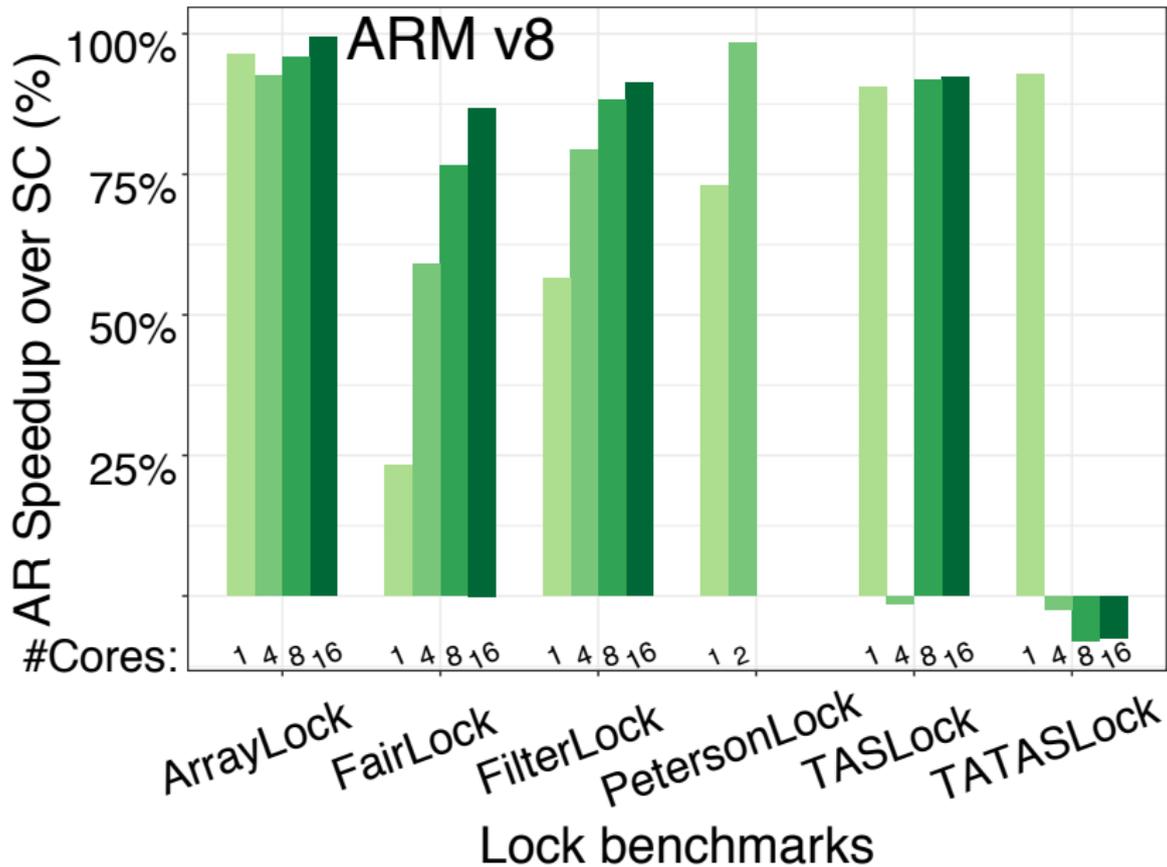


- TATAS-lock permits **starvation**
 - Optimizes throughput of task-ensemble as a whole
- CLH QueueLock and ArrayLock: **FIFO progress-guarantee**
- FairLock: no task can enter critical section twice while another task is waiting
 - Weaker guarantee than FIFO

Performance gain ratio: AR over SC (1/3)



Performance gain ratio: AR over SC (2/3)



Performance gain ratio: AR over SC (3/3)

- took considerable time and a lot of discussions to convert from SC to AR version

Performance gain ratio: AR over SC (3/3)

- took considerable time and a lot of discussions to convert from SC to AR version
- the same or even more to relax acquire and release operations (here we got real performance gains)

Performance gain ratio: AR over SC (3/3)

- took considerable time and a lot of discussions to convert from SC to AR version
- the same or even more to relax acquire and release operations (here we got real performance gains)
- some C++ memory fences turned out to be insufficient

Conclusions and Future Work (1/2)

- Ada's protected objects have advantages:

Conclusions and Future Work (1/2)

- Ada's protected objects have advantages: (1) safe

Conclusions and Future Work (1/2)

- Ada's protected objects have advantages: (1) safe (2) easy to use

Conclusions and Future Work (1/2)

- Ada's protected objects have advantages: (1) safe (2) easy to use (3) easy to comprehend.

Conclusions and Future Work (1/2)

- Ada's protected objects have advantages: (1) safe (2) easy to use (3) easy to comprehend.
- If performance is a major concern, lock-free implementations outperform the blocking approach.

Conclusions and Future Work (1/2)

- Ada's protected objects have advantages: (1) safe (2) easy to use (3) easy to comprehend.
- If performance is a major concern, lock-free implementations outperform the blocking approach.
- To gain higher performance for Ada 2012 \Rightarrow intrinsics or machine code insertions.

Conclusions and Future Work (2/2)

- Simply replacing PO locks with spin loops is not enough.

Conclusions and Future Work (2/2)

- Simply replacing PO locks with spin loops is not enough.
- Concurrent execution inside the methods is better.

Conclusions and Future Work (2/2)

- Simply replacing PO locks with spin loops is not enough.
- Concurrent execution inside the methods is better.
- For Ada 2012 programs \Rightarrow calls to intrinsics.

Conclusions and Future Work (2/2)

- Simply replacing PO locks with spin loops is not enough.
- Concurrent execution inside the methods is better.
- For Ada 2012 programs \Rightarrow calls to intrinsics.
- \Rightarrow adding *concurrent objects* (COs) and a strict memory consistency model to Ada 202x.

Conclusions and Future Work (2/2)

- Simply replacing PO locks with spin loops is not enough.
- Concurrent execution inside the methods is better.
- For Ada 2012 programs \Rightarrow calls to intrinsics.
- \Rightarrow adding *concurrent objects* (COs) and a strict memory consistency model to Ada 202x.
- Future container libraries for Ada 202x should also include lock-free and wait-free data-structures.

Conclusions and Future Work (2/2)

- Simply replacing PO locks with spin loops is not enough.
- Concurrent execution inside the methods is better.
- For Ada 2012 programs \Rightarrow calls to intrinsics.
- \Rightarrow adding *concurrent objects (COs)* and a strict memory consistency model to Ada 202x.
- Future container libraries for Ada 202x should also include lock-free and wait-free data-structures.
- Set up benchmarks for reuse.

Conclusions and Future Work (2/2)

- Simply replacing PO locks with spin loops is not enough.
- Concurrent execution inside the methods is better.
- For Ada 2012 programs \Rightarrow calls to intrinsics.
- \Rightarrow adding *concurrent objects (COs)* and a strict memory consistency model to Ada 202x.
- Future container libraries for Ada 202x should also include lock-free and wait-free data-structures.
- Set up benchmarks for reuse.
- Prototype implementation for COs.