

Hardware Synthesis of Weakly Consistent C Concurrency



John Wickerson
Imperial College London

a talk based on joint work with



Nadesh Ramanathan
Imperial College London



Shane T. Fleming
Imperial College London



George A. Constantinides
Imperial College London

Starting point

Concurrent languages with built-in support for strong/weak atomics



OpenCL



CPUs



GPUs

AMD



NVIDIA



FPGAs



XILINX



FPGAs

- Reconfigurable hardware

FPGAs

TECH

Intel Completes Acquisition of Altera

\$16.7 billion deal underscores Intel CEO's plan to expand chip maker's business

By **DON CLARK**

Updated Dec. 28, 2015 9:27 p.m. ET

Intel Corp. on Monday completed its biggest-ever acquisition, part of Chief Executive Brian Krzanich's plan to use new tactics to expand the chip maker's business.

The \$16.7 billion purchase of Altera Corp. makes Intel, known for microprocessors used in computers, the second-largest maker of

programmable after they leave the factory. Altera's

Xilinx Virtex 5 FPGA in Driverless Cars

Posted by Nadeesha Thewarapperuma

Find me on: [in](#) [g+](#)

Aug 4, 2015 11:00:00 AM

[Tweet](#)

[in](#) Share

19

[f](#) Like

Share

23

[G+1](#)

5

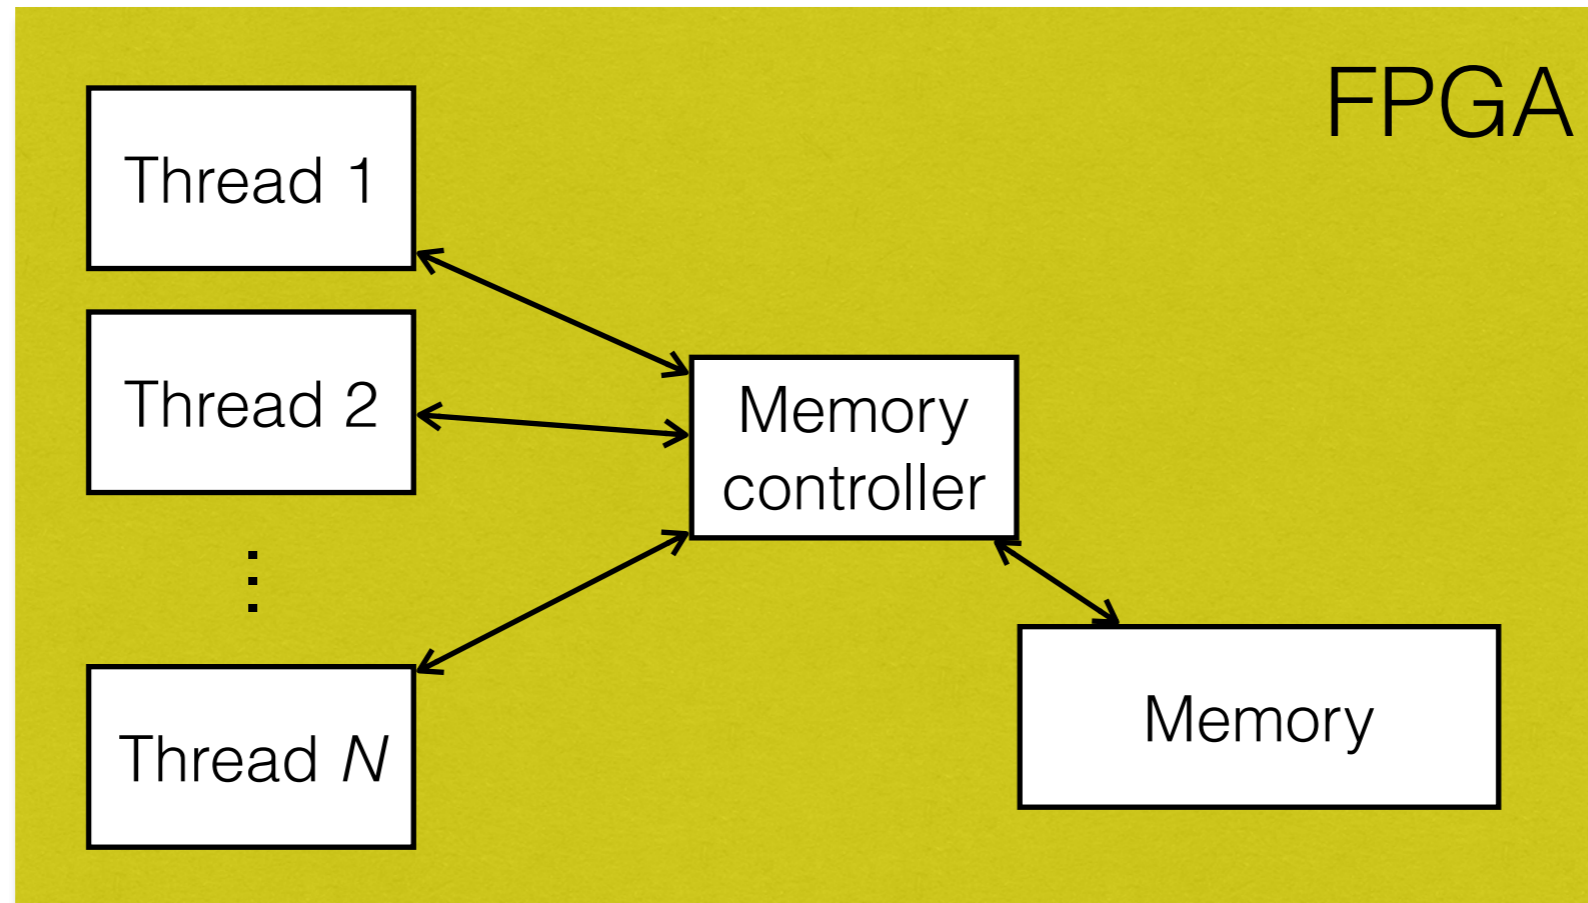


art

for
r of
tera's

LegUp

- Open-source hardware synthesis tool developed at the University of Toronto since 2009
- Supports pthreads, and OpenMP (using locks)



Synthesis example

- Can we implement atomic stores/loads using just ordinary stores/loads?

```
int x=0; int y=0;
-----
T1() {
1.1  x=1;
1.2  y=1;
}
-----
T2() {
2.1  if(y==1)
2.2  r0=x;
}
-----
assert(r0 != 0)
```

Synthesis example

- Can we implement atomic stores/loads using just ordinary stores/loads?

```
int x=0; int y=0;
-----
T1(int a) {
1.1 x=a/3;
1.2 y=1;
}
T2() {
2.1 if(y==1)
2.2 r0=x;
}
-----
assert(r0 != 0)
```



```
int x=0; int y=0;
```

| | |
|--|--|
| <pre>T1(int a) { 1.1 x=a/3; 1.2 y=1; }</pre> | <pre>T2() { 2.1 if(y==1) 2.2 r0=x; }</pre> |
|--|--|

```
assert(r0 ≠ 0)
```

| | | | | | | | | |
|--------|---|---|---|---|---|-----|----|----|
| Cycle: | 1 | 2 | 3 | 4 | 5 | ... | 35 | 36 |
|--------|---|---|---|---|---|-----|----|----|

| | | | | | | | | |
|-----|------|--|--------|--|--|--|------|--|
| 1.1 | ld a | | | | | | | |
| 1.1 | | | divide | | | | | |
| 1.1 | | | | | | | st x | |
| 1.2 | st y | | | | | | | |

| | | | | | | | |
|-----|--|--|------|--|---------------------|--|--|
| 2.1 | | | ld y | | | | |
| 2.2 | | | ld x | | | | |
| 2.2 | | | | | slt y==1? x:null | | |

Implementing atomics

- **Unsound:** only respects RW/WR/WW dependencies

```
r0=w;
```

```
r1=x;
```

```
r2=y.ld(ACQ);
```

```
r3=z;
```

Implementing atomics

- **Unsound:** only respects RW/WR/WW dependencies

| Cycle: | 1 | 2 |
|----------------------------|---------------------------------|---|
| <code>r0=w;</code> | <code>ld_{na} w</code> | |
| <code>r1=x;</code> | <code>ld_{na} x</code> | |
| <code>r2=y.ld(ACQ);</code> | <code>ld_{ACQ} y</code> | |
| <code>r3=z;</code> | <code>ld_{na} z</code> | |

Implementing atomics

- **SC:** all memory accesses strictly ordered

| Cycle: | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
|----------------------------|--------------------|---|--------------------|---|---------------------|---|--------------------|---|
| <code>r0=w;</code> | ld _{na} w | | | | | | | |
| <code>r1=x;</code> | | | ld _{na} x | | | | | |
| <code>r2=y.ld(ACQ);</code> | | | | | ld _{ACQ} y | | | |
| <code>r3=z;</code> | | | | | | | ld _{na} z | |

Implementing atomics

- **SC-atomics:** all atomics are strictly ordered

| Cycle: | 1 | 2 | 3 | 4 | 5 | 6 |
|----------------------------|--------------------------------|---|---------------------------------|---|--------------------------------|---|
| <code>r0=w;</code> | <code>ld_{na} w</code> | | | | | |
| <code>r1=x;</code> | <code>ld_{na} x</code> | | | | | |
| <code>r2=y.ld(ACQ);</code> | | | <code>ld_{ACQ} y</code> | | | |
| <code>r3=z;</code> | | | | | <code>ld_{na} z</code> | |

Implementing atomics

- **Weak-atomics:** acquires cannot move down; releases cannot move up; SCs cannot move at all

| Cycle: | 1 | 2 | 3 | 4 |
|----------------------------|---------------------------------|---|--------------------------------|---|
| <code>r0=w;</code> | <code>ld_{na} w</code> | | | |
| <code>r1=x;</code> | <code>ld_{na} x</code> | | | |
| <code>r2=y.ld(ACQ);</code> | <code>ld_{ACQ} y</code> | | | |
| <code>r3=z;</code> | | | <code>ld_{na} z</code> | |

Correctness

- We used the Alloy Analyzer to check that there is no execution (with ≤ 9 memory events) that:
 - is allowed by our scheduling constraints, but
 - is inconsistent according to the C memory model.

Automatically Comparing Memory Consistency Models

John Wickerson
Imperial College London, UK
j.wickerson@imperial.ac.uk

Tyler Sorensen
Imperial College London, UK
t.sorensen15@imperial.ac.uk

Mark Batty
University of Kent, UK
m.j.batty@kent.ac.uk

George A. Constantinides
Imperial College London, UK
g.constantinides@imperial.ac.uk



Abstract


A *memory consistency model* (MCM) is the part of a programming language specification that defines which memory locations are visible to a processor. Because

1. Introduction

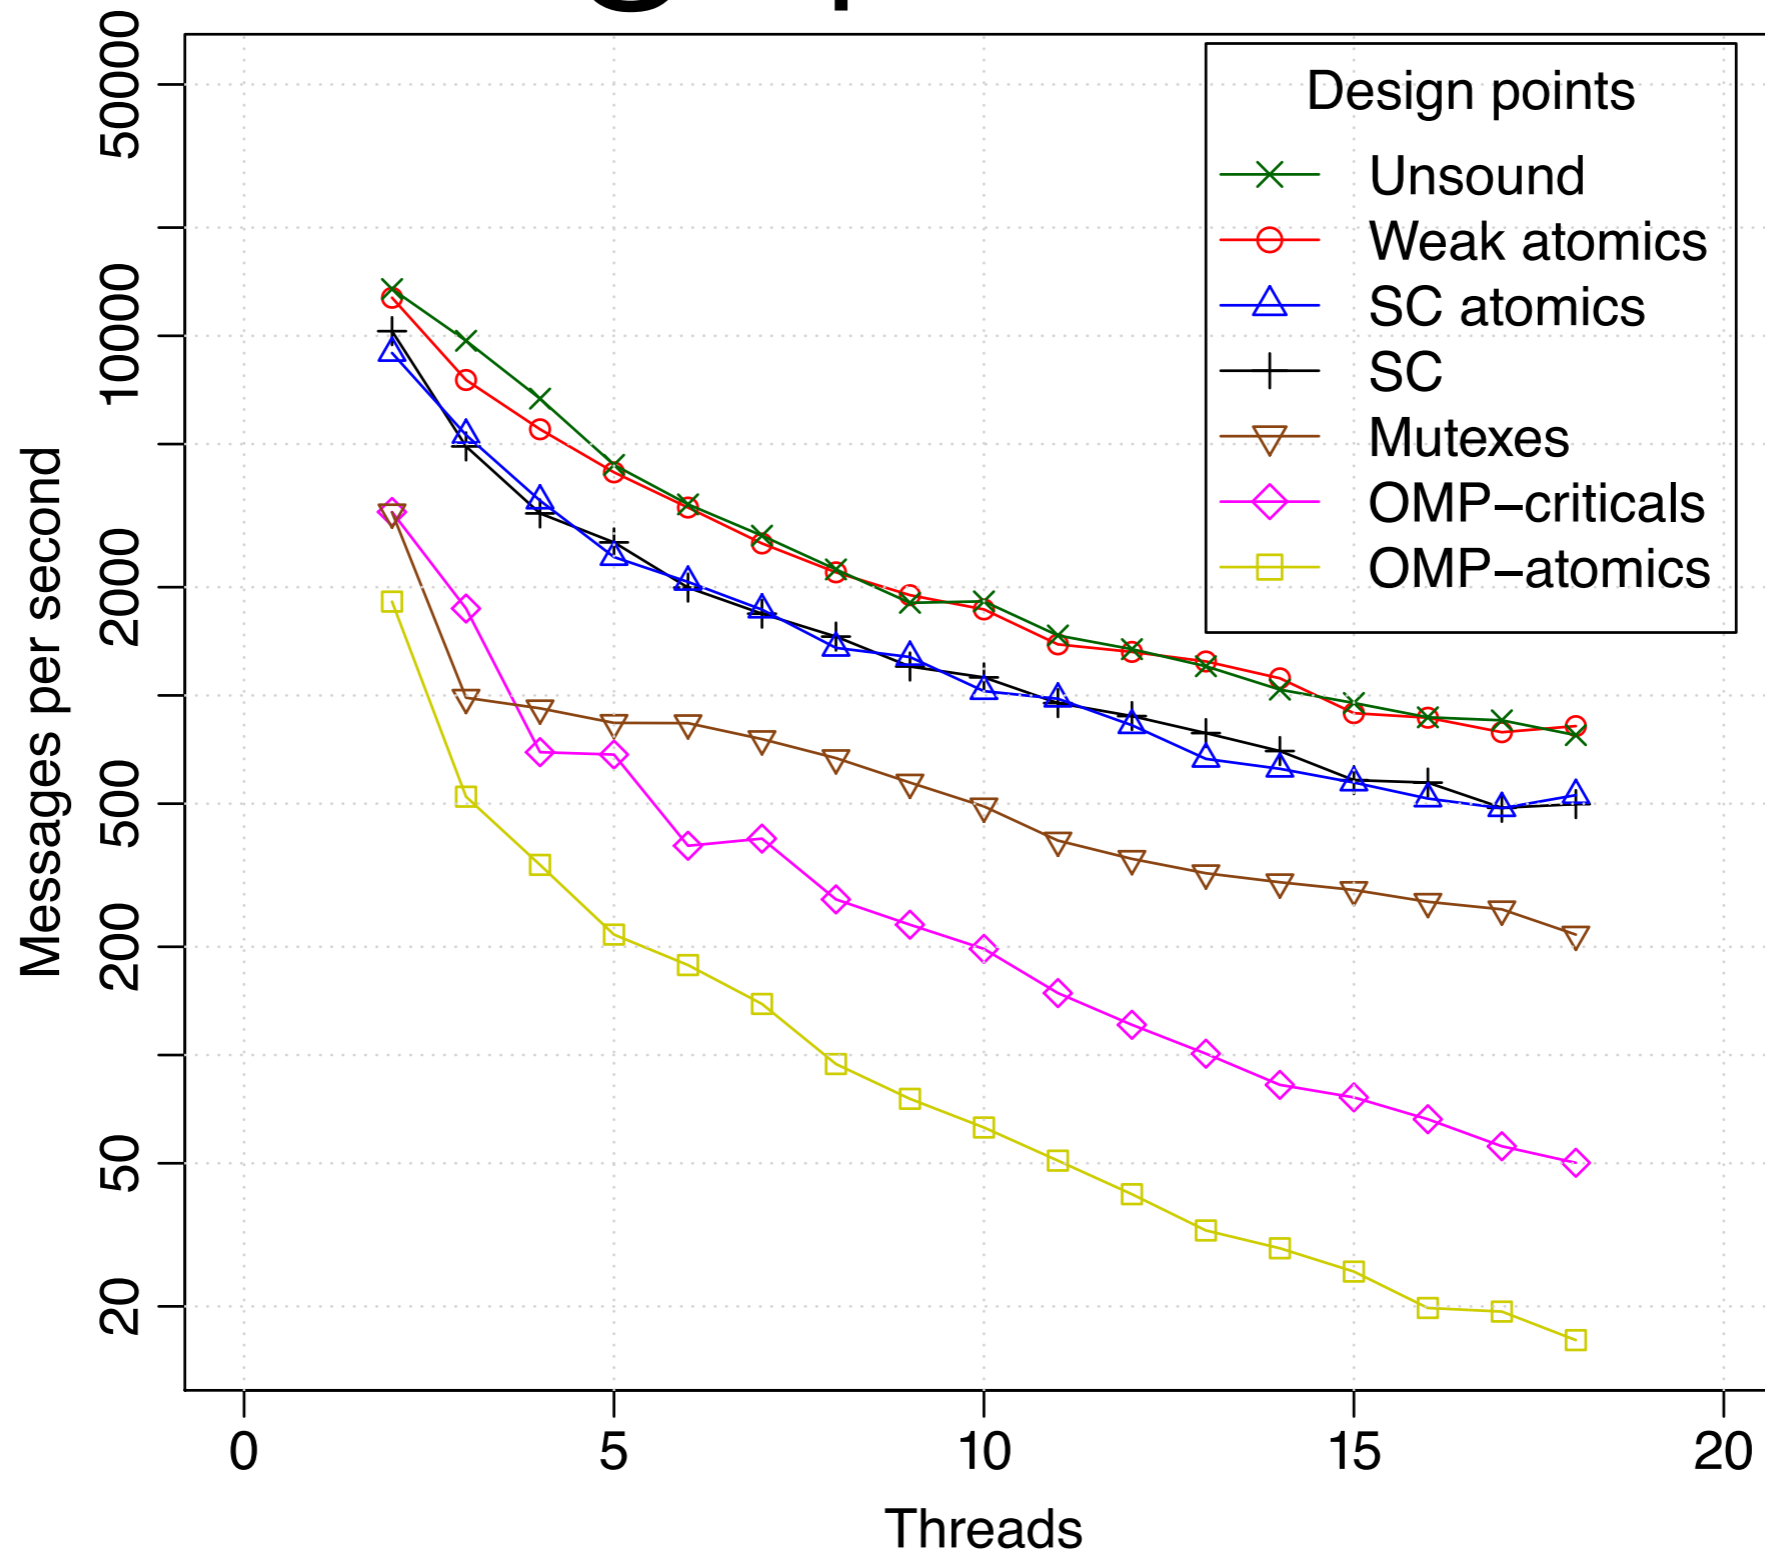
In the specification of a concurrent programming language or a parallel architecture, the *memory consistency model* (MCM) defines which values can legally be read from shared memory locations [4]. The MCM is chosen through to enable portability, but specific

Case study: circular FIFO

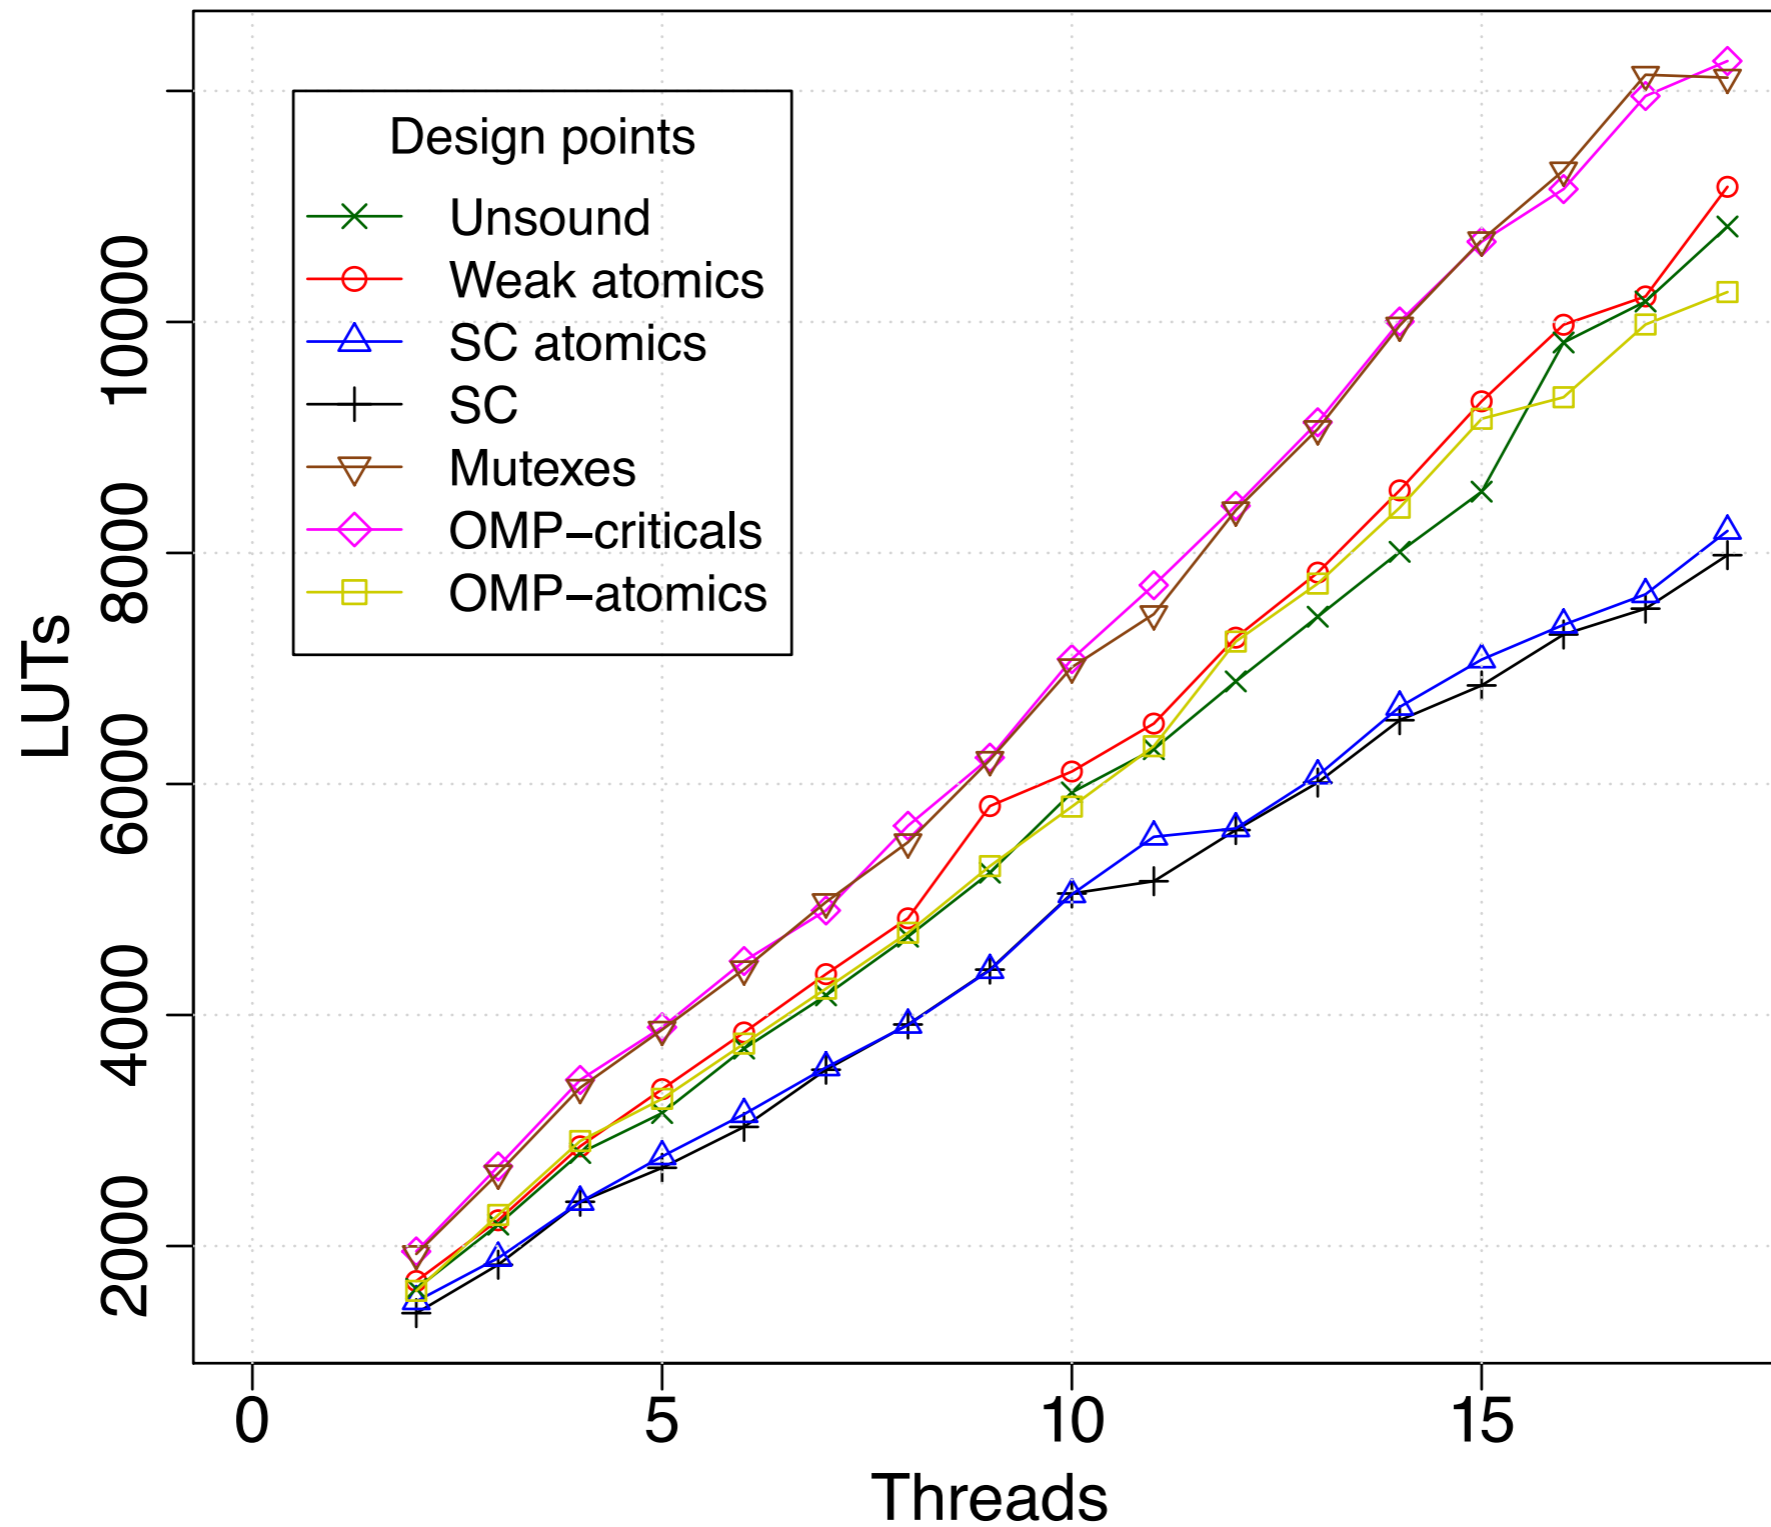
| | <code>atomic_int tail=0; head=0;</code> <code>int arr[SIZE]; res[MSGs];</code> | | |
|------|---|--|------|
| 1.1 | <code>while(prod<MSGs) {</code> | <code>while(cons<MSGs) {</code> | 2.1 |
| 1.2 | <code> thead = head.ld(ACQ);</code> | <code> ctail = tail.ld(ACQ);</code> | 2.2 |
| 1.3 | <code> ctail = tail.ld(RLX);</code> | <code> thead = head.ld(RLX);</code> | 2.3 |
| 1.4 | <code> ntail = (ctail+1)%SIZE;</code> | <code> nthead = (thead+1)%SIZE;</code> | 2.4 |
| 1.5 | <code> if(ntail != thead){</code> | <code> if(ctail != thead){</code> | 2.5 |
| 1.6 | <code> arr[ctail] = prod</code> | <code> res[cons] = arr[thead];</code> | 2.6 |
| 1.7 | <code> tail.st(ntail,REL);</code> | <code> thead.st(nthead,REL);</code> | 2.7 |
| 1.8 | <code> prod++;</code> | <code> cons++;</code> | 2.8 |
| 1.9 | <code> }</code> | <code> }</code> | 2.9 |
| 1.10 | <code>}</code> | <code>}</code> | 2.10 |



Throughput results



Circuit area results



Conclusion

- First implementation of weak atomics in a hardware synthesis tool
- Implementing atomics using scheduling constraints seems more efficient than using locks
- **Limitations:** no support for RMW operations; small and artificial benchmarks; only on-chip memory
- **Next steps:** add support for loop pipelining

Hardware Synthesis of Weakly Consistent C Concurrency



John Wickerson
Imperial College London

a talk based on joint work with



Nadesh Ramanathan
Imperial College London



Shane T. Fleming
Imperial College London



George A. Constantinides
Imperial College London