

JOURNEY TO THE CENTRE OF THE
OPENGL
MEMORY MODEL

John Wickerson
Imperial College London

COLLABORATORS



Mark Batty




George Constantinides



Nadesh Ramanathan



Alastair Donaldson



Tyler Sorensen



Brad Beckmann

THIS TALK

THIS TALK

- Weak memory

THIS TALK

- Weak memory
- Formalising the OpenCL memory model

THIS TALK

- Weak memory
- Formalising the OpenCL memory model
- Implementing the OpenCL memory model on GPUs

THIS TALK

- Weak memory
- Formalising the OpenCL memory model
- Implementing the OpenCL memory model on GPUs
- Implementing the OpenCL memory model on FPGAs

WEAK MEMORY

WEAK MEMORY

```
MOV [x] 1
```

```
MOV [y] 1
```

```
MOV r0 [y]
```

```
MOV r1 [x]
```

WEAK MEMORY

<code>MOV [x] 1</code>	<code>MOV [y] 1</code>
<code>MOV r0 [y]</code>	<code>MOV r1 [x]</code>

`r0=1`

`r1=1`

WEAK MEMORY

MOV [x] 1	MOV [y] 1
MOV r0 [y]	MOV r1 [x]

r0=1

r1=1

r0=0

r1=1

WEAK MEMORY

```
MOV [x] 1
```

```
MOV [y] 1
```

```
MOV r0 [y]
```

```
MOV r1 [x]
```

```
r0=1
```

```
r0=0
```

```
r0=1
```

```
r1=1
```

```
r1=1
```

```
r1=0
```

WEAK MEMORY

MOV [x] 1	MOV [y] 1
MOV r0 [y]	MOV r1 [x]

r0=1

r0=0

r0=1

r1=1

r1=1

r1=0

SC

WEAK MEMORY

<code>MOV [x] 1</code>	<code>MOV [y] 1</code>
<code>MOV r0 [y]</code>	<code>MOV r1 [x]</code>

`r0=1`

`r0=0`

`r0=1`

`r0=0`

`r1=1`

`r1=1`

`r1=0`

`r1=0`

SC

WEAK MEMORY

<code>MOV [x] 1</code>	<code>MOV [y] 1</code>
<code>MOV r0 [y]</code>	<code>MOV r1 [x]</code>

`r0=1`

`r0=0`

`r0=1`

`r0=0`

`r1=1`

`r1=1`

`r1=0`

`r1=0`

SC

x86

WEAK MEMORY

```
MOV [x] 1 | MOV [y] 1
MOV r0 [y] | MOV r1 [x]
```

r0=1

r0=0

r0=1

r0=0

r1=1

r1=1

r1=0

r1=0

SC

x86



WEAK MEMORY

```
MOV [x] 1 | MOV [y] 1
MOV r0 [y] | MOV r1 [x]
```

r0=1

r0=0

r0=1


r0=0

r1=1


r1=1

r1=0

r1=0



x86





STACK MEMORY

```

MOV [x] 1 | MOV [y] 1
MOV r0 [y] | MOV r1 [x]

```

r0=1
r1=1


r0=0
r1=1


r0=1
r1=0

r0=0
r1=0



x86





```

MOV [x] 1 | MOV [y] 1
MOV r0 [y] | MOV r1 [x]

```

r0=1
r1=1


r0=0
r1=1


r0=1
r1=0

r0=0
r1=0



x86





```

MOV [x] 1 | MOV [y] 1
MOV r0 [y] | MOV r1 [x]


```

r0=1
r1=1


r0=0
r1=1



r0=1
r1=0

r0=0
r1=0




x86





```
MOV [x] 1
MOV r0 [y] | MOV
```




r0=1
r1=1


r0=0
r1=1





r0=
r1=0

r1=0







x86





MOV [x] 1
MOV [y] 1
=0 r0=
=1 r1=0
r1=1 r1=0

x86




WEAK MEMORY IS HARD!


- x86 proved tricky to formalise correctly [*Sarkar et al., POPL'09; Owens et al., TPHOLs'09*]
- Bug found in deployed "Power 5" processors [*Alglave et al., CAV'10*]
- C++ specification did not guarantee its own key property [*Batty et al., POPL'11*]
- Routine compiler optimisations are invalid under Java and C++ memory models [*Sevcik, PLDI'11; Vafeiadis et al. POPL'15*]
- Behaviour of NVIDIA graphics processors contradicted NVIDIA's programming guide [*Alglave et al., ASPLOS'15*]


THIS TALK


- ~~Weak memory~~
- Formalising the OpenCL memory model
- Implementing the OpenCL memory model on GPUs
- Implementing the OpenCL memory model on FPGAs












For an atomic operation **B** that reads the value of an atomic object **M**, if there is a *memory_order_seq_cst* fence **X** sequenced-before **B**, then **B** observes either the last *memory_order_seq_cst* modification of **M** preceding **X** in the total order **S** or a later modification of **M** in its modification order. [OpenCL 2.0 standard, 2015]



For an atomic operation **B** that reads the value of an atomic object **M**, if there is a *memory_order_seq_cst* fence **X** sequenced-before **B**, then **B** observes either the last *memory_order_seq_cst* modification of **M** preceding **X** in the total order **S** or a later modification of **M** in its modification order. [OpenCL 2.0 standard, 2015]

"OpenCL"

withoutsc

```
let mo = co & ((!nonatomicloc)^2)
```

```
let sb = po
```

```
let rb = (rf^-1; mo) \ id
```

```
(* Access modes *)
```

```
let mo_acq = memory_order_acquire
```

```
let mo_rel = memory_order_release
```

```
let mo_acq_rel = memory_order_acq_rel
```

```
let mo_rlx = memory_order_relaxed
```

```
let mo_sc = memory_order_seq_cst
```

```
(* Scope annotations *)
```

```
let s_wi = memory_scope_work_item
```

```
let s_wg = memory_scope_work_group
```

```
let s_dev = memory_scope_device
```

```
let s_all = memory_scope_all_svm_devices
```

```
(* Synchronisation *)
```

```
(*****)
```

```
let acq = (mo_acq | mo_acq_rel | mo_sc) & (R | F)
```

```
let rel = (mo_rel | mo_acq_rel | mo_sc) & (W | F)
```

```
(* Fences sequenced before or after *)
```

```
let Fsb = [F]; sb
```

```
let sbF = sb; [F]
```

```
(* Release sequence *)
```

```
let rs' = wi | (unv; [R & W])
```

```
let rs = mo & rs' & -((mo & ~rs'); mo)
```

```
(* Inclusive scopes *)
```

```
let incl = wg & s_wg^2 | dev & s_dev^2 | s_all^2
```

```
(* Inclusive scopes, less conservative  
(unused) version *)
```

```
let incl1 = ([s_wg];wg) | ([s_dev];dev)
```

```
          | ([s_all];unv)
```

```
let incl' = incl1 & incl1^1
```

```
(* Release-acquire synchronisation *)
```

```
let ra_sw(r) =
```

```
  [[r & rel]; Fsb?; [W \ s_wi]; rs?; [r]; rf;
```

```
  [R \ s_wi]; sbF?; [acq & r] & incl & ~wi
```

```
(* Barrier synchronisation *)
```

```
let bar_sw(r) = (entry_fence * exit_fence) &  
              same_B & ~wi & wg & r^2
```

```
(* Allowed to synchronise on the other region *)
```

```
let scf = mo_sc^2 | (G & L & F)^2
```

```
(* Global and local synchronises-with *)
```

```
let gsw = ra_sw(G) | bar_sw(G) | (scf & ra_sw(L))
```

```
let lsw = ra_sw(L) | bar_sw(L) | (scf & ra_sw(G))
```

```
(* Happens-before *)
```

```
(*****)
```

```
(* Global and local happens-before *)
```

```
let ghb = (((G^2) & (sb | (I * !I))) | gsw)+
```

```
let lhb = (((L^2) & (sb | (I * !I))) | lsw)+
```

```
show ghb
```

```
show lhb
```

```
irreflexive ghb as 0-HbG
```

```
irreflexive lhb as 0-HbL
```

```
(* Coherence *)
```

```
(*****)
```

```
let coh(hb) = (rf^-1)?; mo; rf?; hb
```

```
irreflexive coh(ghb) as 0-CohG
```

```
irreflexive coh(lhb) as 0-CohL
```

```
(* Consistency of reads *)
```

```
(*****)
```

```
(* A load can only read from a store that already  
happened. *)
```

```
irreflexive rf; (ghb | lhb) as 0-Rf
```

```
(* Visible side effects *)
```

```
let vis(hb) = (W * R) & hb & loc &
```

```
              ~((hb & loc); [W]; hb)
```

```
(* A non-atomic load can only read from a store  
that is visible. *)
```

```
empty (rf;[G & nonatomicloc])\vis(ghb) as 0-NaRfG
```

```
empty (rf;[L & nonatomicloc])\vis(lhb) as 0-NaRfL
```

```
(* Consistency of RMWs *)
```

```
irreflexive rf | (mo;mo;rf^-1) | (mo;rf) as 0-Rmw
```



```

"OpenCL"
withoutsc

let mo = co & (!!nonatomicloc)^2
let sb = po
let rb = (rf^-1; mo) \ id

(* Access modes *)
let mo_acq = memory_order_acquire
let mo_rel = memory_order_release
let mo_acq_rel = memory_order_acq_rel
let mo_rlx = memory_order_relaxed
let mo_sc = memory_order_seq_cst

(* Scope annotations *)
let s_wi = memory_scope_work_item
let s_wg = memory_scope_work_group
let s_dev = memory_scope_device
let s_all = memory_scope_all_svm_devices

(* Synchronisation *)
(*****)

let acq = (mo_acq | mo_acq_rel | mo_sc) & (R | F)
let rel = (mo_rel | mo_acq_rel | mo_sc) & (W | F)

(* Fences sequenced before or after *)
let Fsb = [F]; sb
let sbF = sb; [F]

(* Release sequence *)
let rs' = wi | (unv; [R & W])
let rs = mo & rs' & ~((mo & ~rs') ; mo)

(* Inclusive scopes *)
let incl = wg & s_wg^2 | dev & s_dev^2 | s_all^2

(* Inclusive scopes, less conservative
(unused) version *)
let incl1 = ([s_wg];wg) | ([s_dev];dev)
| ([s_all];unv)
let incl' = incl1 & incl1^-1

(* Release-acquire synchronisation *)
let ra_sw(r) =
  ([r & rel]; Fsb?; [W \ s_wi]; rs?; [r]; rf;

```

```

(* Barrier synchronisation *)
let bar_sw(r) = (entry_fence * exit_fence) &
  same_B & ~wi & wg & r^2

(* Allowed to synchronise on the other region *)
let scf = mo_sc^2 | (G & L & F)^2

(* Global and local synchronises-with *)
let gsw = ra_sw(G) | bar_sw(G) | (scf & ra_sw(L))
let lsw = ra_sw(L) | bar_sw(L) | (scf & ra_sw(G))

(* Happens-before *)
(*****)

(* Global and local happens-before *)
let ghb = (((G^2) & (sb | (I * !I))) | gsw)+
let lhb = (((L^2) & (sb | (I * !I))) | lsw)+
show ghb
show lhb
irreflexive ghb as 0-HbG
irreflexive lhb as 0-HbL

(* Coherence *)
(*****)

let coh(hb) = (rf^-1)?; mo; rf?; hb
irreflexive coh(ghb) as 0-CohG
irreflexive coh(lhb) as 0-CohL

(* Consistency of reads *)
(*****)


(* A load can only read from a store that already
happened. *)
irreflexive rf; (ghb | lhb) as 0-Rf

(* Visible side effects *)
let vis(hb) = (W * R) & hb & loc &
  ~((hb & loc); [W]; hb)

(* A non-atomic load can only read from a store
that is visible. *)
empty (rf;[G & nonatomicloc])\vis(ghb) as 0-NaRfG
empty (rf;[L & nonatomicloc])\vis(lhb) as 0-NaRfL

(* Consistency of RMWs *)
(*****)
empty (rf;[G & nonatomicloc])\vis(ghb) as 0-RmwG
empty (rf;[L & nonatomicloc])\vis(lhb) as 0-RmwL


```





For an atomic operation **B** that reads the value of an atomic object **M**, if there is a *memory_order_seq_cst* fence **X** sequenced-before **B**, then **B** observes either the last *memory_order_seq_cst* modification of **M** preceding **X** in the total order **S** or a later modification of **M** in its modification order. [OpenCL 2.0 standard, 2015]

THIS TALK

- ~~• Weak memory~~
- ~~• Formalising the OpenCL memory model~~
- Implementing the OpenCL memory model on GPUs
- Implementing the OpenCL memory model on FPGAs









SyState

DvState

WgState

ThState

Reg	Val
-----	-----

rmw: Lock

L1: Cache

Addr	Val	Hygiene	Freshness
------	-----	---------	-----------

Fifo

L2: Cache


Addr	Val	Hygiene	Freshness
------	-----	---------	-----------

Fifo


Addr | Lock

Global

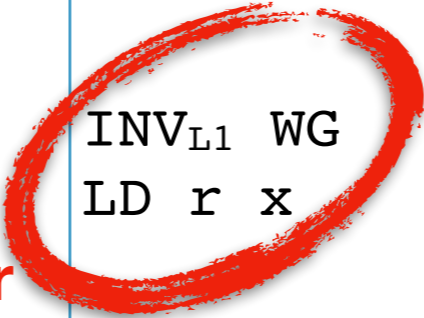
Addr | Val

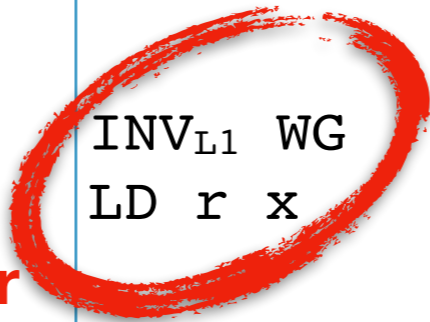
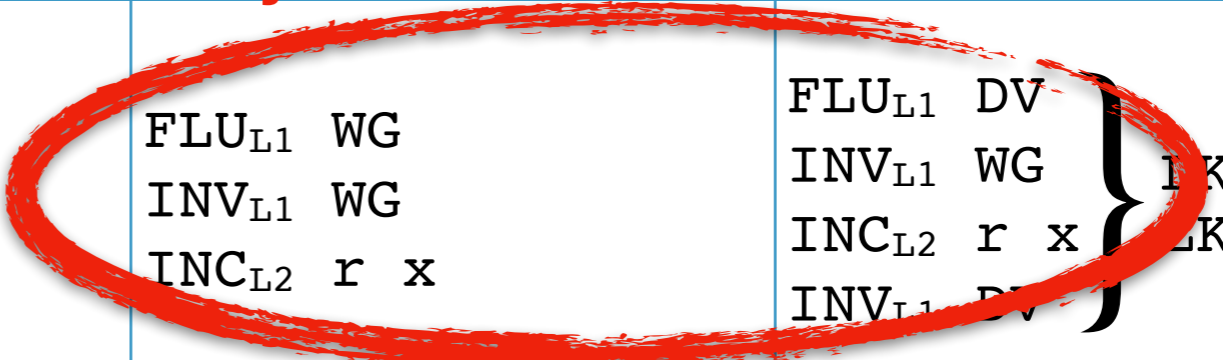


	na or WG	DV (not remote)	DV (remote)
<code>r=load(x)</code>	LD r x	INV _{L1} WG LD r x	FLU _{L1} DV INV _{L1} WG LD r x } LK x
<code>store(x,r)</code>	ST r x	FLU _{L1} WG ST r x	FLU _{L1} WG ST r x INV _{L1} DV } LK x
<code>r=fetch_inc(x)</code>	INC _{L1} r x	FLU _{L1} WG INV _{L1} WG INC _{L2} r x	FLU _{L1} DV INV _{L1} WG INC _{L2} r x INV _{L1} DV } LK x LK _{rmw}



	na or WG	DV (not remote)	DV (remote)
<code>r=load(x)</code>	LD r x	INV _{L1} WG LD r x	FLU _{L1} DV INV _{L1} WG LD r x } LK x
<code>store(x,r)</code>	ST r x	FLU _{L1} WG ST r x	FLU _{L1} WG ST r x INV _{L1} DV } LK x
<code>r=fetch_inc(x)</code>	INC _{L1} r x	FLU _{L1} WG INV _{L1} WG INC _{L2} r x	FLU _{L1} DV INV _{L1} WG INC _{L2} r x INV _{L1} DV } LK x LK _{rmw}

	na or WG	DV (not remote)	DV (remote)
r=load(x) message passing error	LD r x	 INV _{L1} WG LD r x	FLU _{L1} DV INV _{L1} WG LD r x } LK x
store(x,r)	ST r x	FLU _{L1} WG ST r x	FLU _{L1} WG ST r x INV _{L1} DV } LK x
r=fetch_inc(x)	INC _{L1} r x	FLU _{L1} WG INV _{L1} WG INC _{L2} r x	FLU _{L1} DV INV _{L1} WG INC _{L2} r x INV _{L1} DV } LK x LK _{rmw}

	na or WG	DV (not remote)	DV (remote)
r=load(x)	LD r x	 INV _{L1} WG LD r x	FLU _{L1} DV } INV _{L1} WG } LK x LD r x
store(x,r)	ST r x	FLU _{L1} WG ST r x	FLU _{L1} WG } ST r x } LK x INV _{L1} DV
r=fetch_inc(x)	INC _{L1} r x	 FLU _{L1} WG INV _{L1} WG INC _{L2} r x	FLU _{L1} DV } INV _{L1} WG } LK x INC _{L2} r x } LK _{rmw} INV _{L1} DV

message passing error



RMW atomicity error

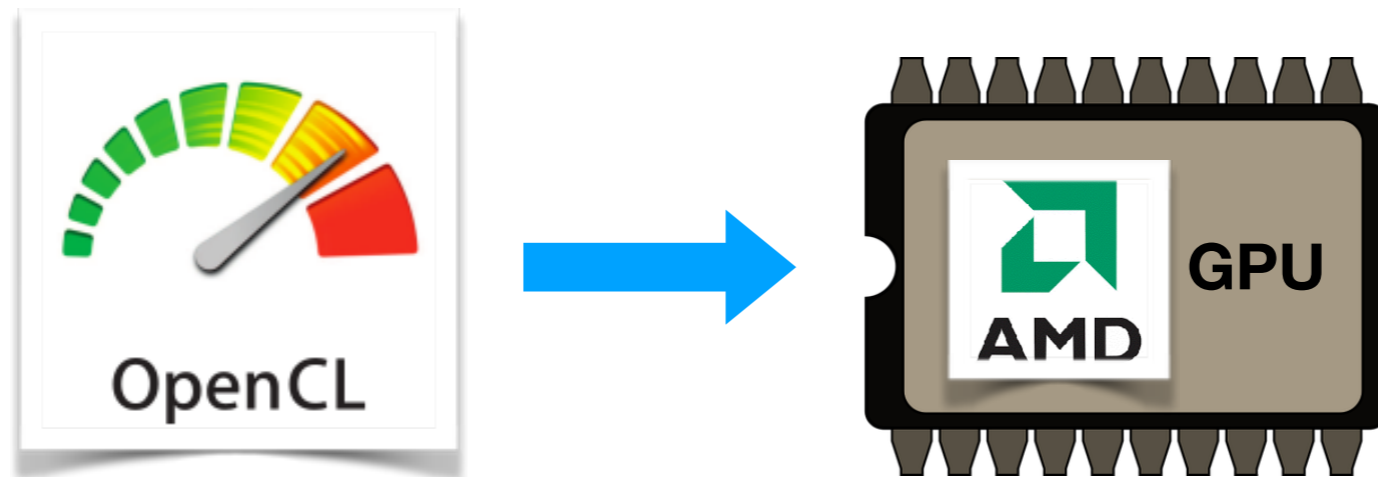
	na or WG	DV (not remote)	DV (remote)
r=load(x)	LD r x	INV _{L1} WG LD r x	unnecessary locking FLU _{L1} DV INV _{L1} WG LD r x LK x
store(x,r)	ST r x	FLU _{L1} WG ST r x	FLU _{L1} WG ST r x INV _{L1} DV LK x
r=fetch_inc(x)	INC _{L1} r x	FLU _{L1} WG INV _{L1} WG INC _{L2} r x	FLU _{L1} DV INV _{L1} WG INC _{L2} r x INV _{L1} DV LK x LK _{rmw}

message passing error


RMW atomicity error

	na or WG	DV (not remote)	DV (remote)
<code>r=load(x)</code>	LD r x	LD r x INV _{L1} WG	LD r x FLU _{L1} DV INV _{L1} WG
<code>store(x,r)</code>	ST r x	FLU _{L1} WG ST r x	FLU _{L1} WG INV _{L1} DV ST r x } LK _{rmw}
<code>r=fetch_inc(x)</code>	INC _{L1} r x	FLU _{L1} WG INC _{L2} r x INV _{L1} WG	FLU _{L1} WG INV _{L1} DV INC _{L2} r x FLU _{L1} DV INV _{L1} WG } LK _{rmw}












"Only after the work by John and others in the programming languages community, have programmers gained confidence in using fine-grain inter-thread communication and synchronization on GPUs." *Dr Brad Beckmann, Principal Researcher, AMD*











"Only after the work by John and others in the programming languages community, have programmers gained confidence in using fine-grain inter-thread communication and synchronization on GPUs." *Dr Brad Beckmann, Principal Researcher, AMD*















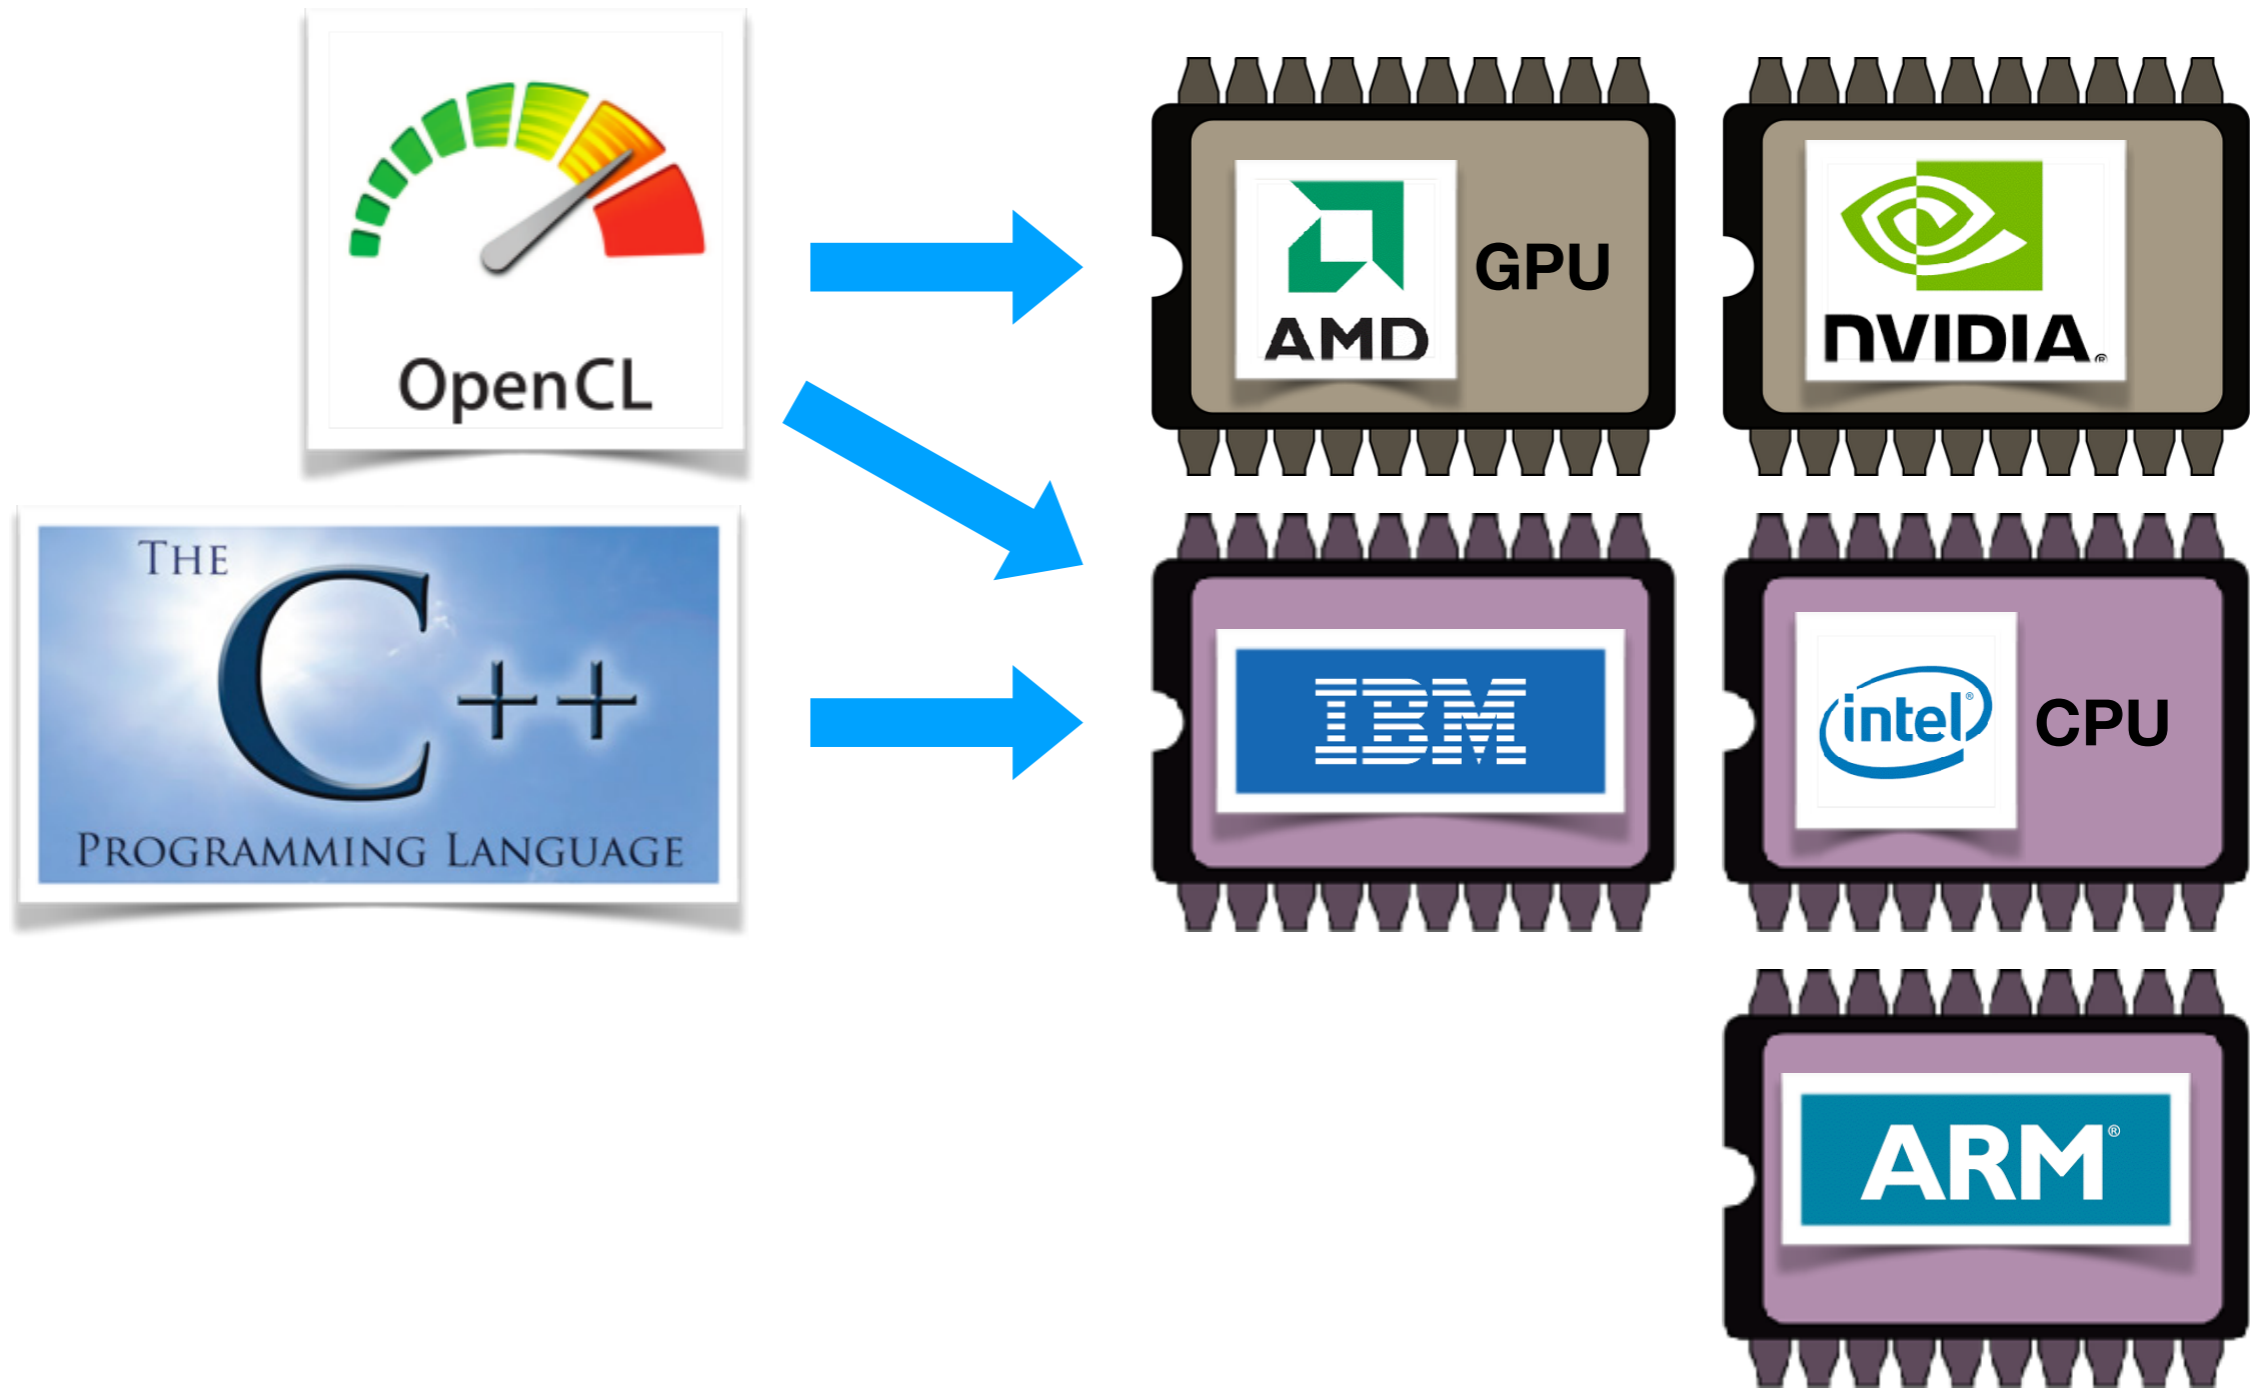

















Model of
the ARM
architecture



Model of
the ARM
architecture

Model of the
'transactional
lock elision'
process

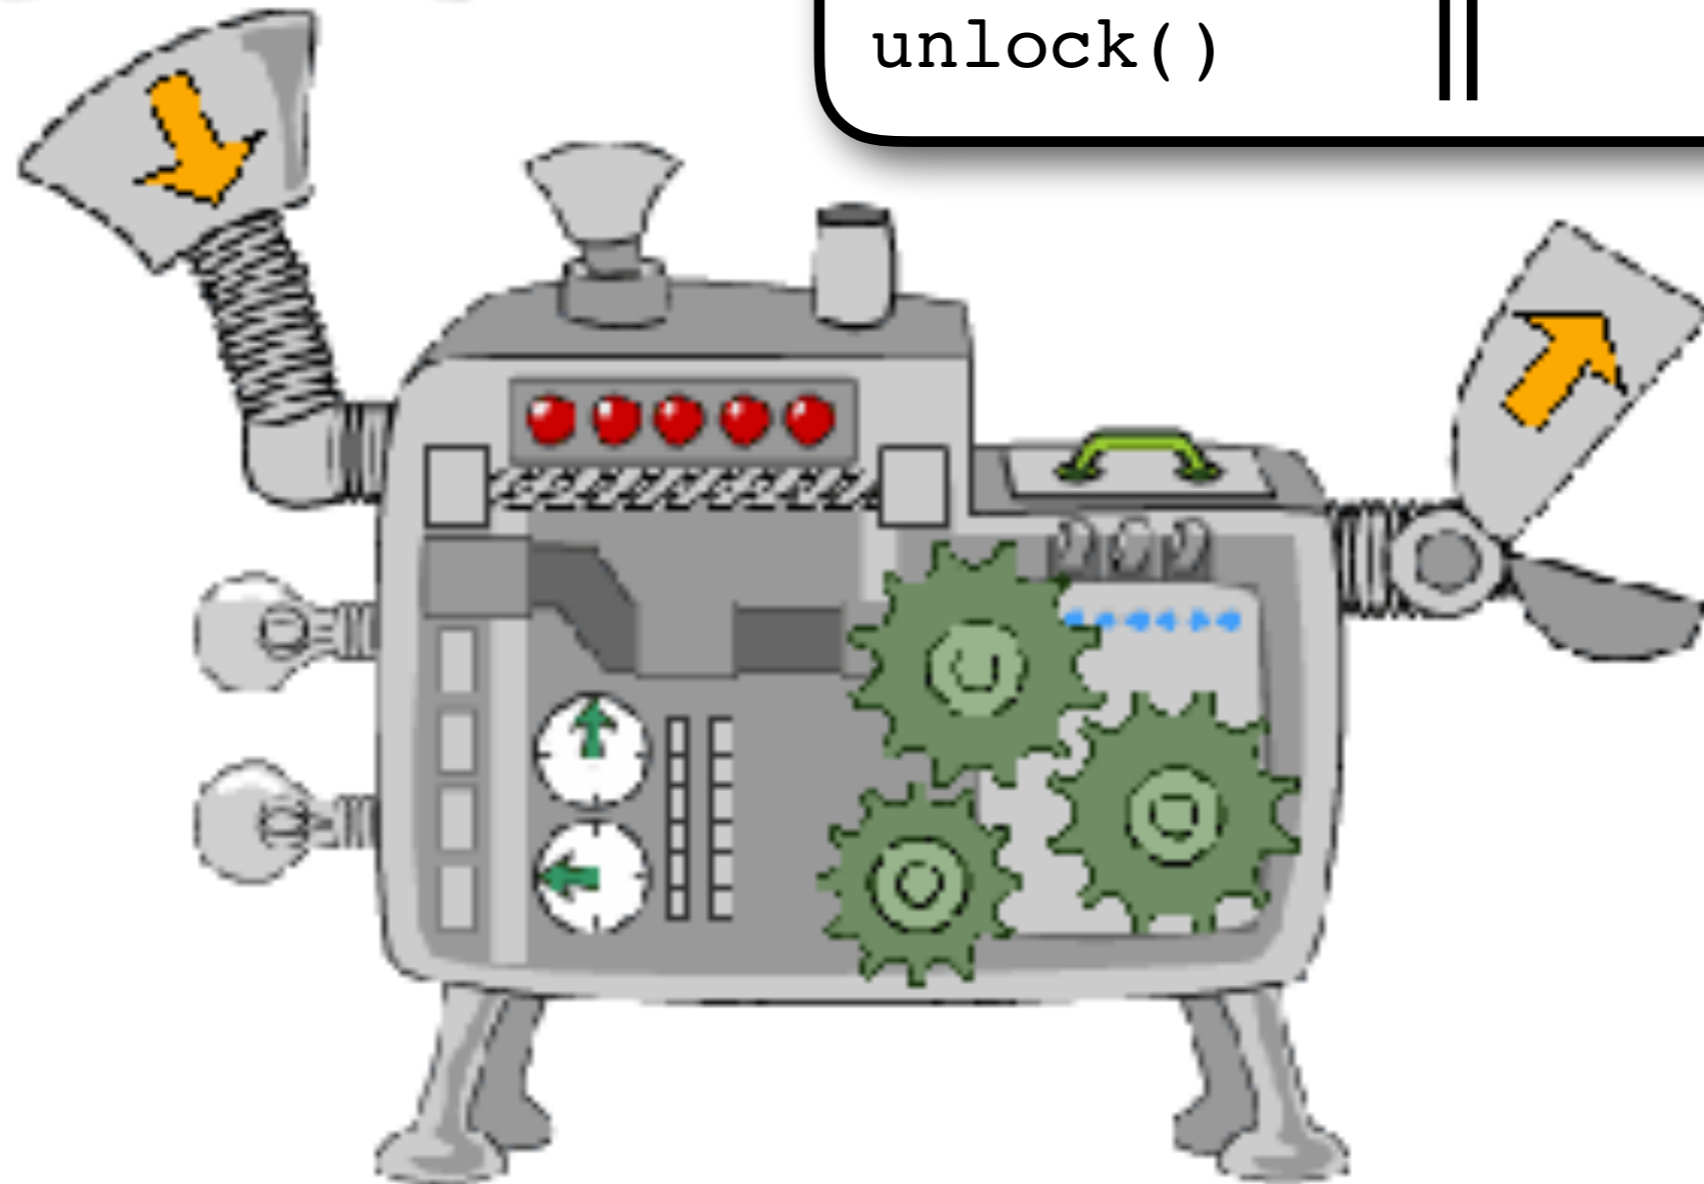


Model of
the ARM
architecture

Model of the
'transactional
lock elision'
process

```
lock()  
LDR W5,[X0]  
ADD W5,W5,#2  
STR W5,[X0]  
unlock()
```

```
lock()  
MOV W7,#1  
STR W7,[X0]  
unlock()
```




Model of
the ARM
architecture

Model of the
'transactional
lock elision'
process


```
lock()  
LDR W5,[X0]  
ADD W5,W5,#2  
STR W5,[X0]  
unlock()
```


```
lock()  
MOV W7,#1  
STR W7,[X0]  
unlock()
```





THIS TALK


- ~~• Weak memory~~
- ~~• Formalising the OpenCL memory model~~
- ~~• Implementing the OpenCL memory model on GPUs~~
- Implementing the OpenCL memory model on FPGAs













```
r = atomic_load(&y,  
memory_order_acquire);
```



```
r = atomic_load(&y,  
memory_order_acquire);
```

****not supported****



```
r = atomic_load(&y,  
memory_order_acquire);
```

```
lock();  
r = y;  
unlock();
```

SCHEDULING ATOMICS

Clock cycle:	1	2	3	4	5	6
<code>r1 = x</code>	load x					
<code>r2 = atomic_load(&y, acquire)</code>			load y			
<code>r3 = z</code>					load z	

SCHEDULING ATOMICS

Clock cycle:	1	2	3	4	5	6
<code>r1 = x</code>	load x					
<code>r2 = atomic_load(&y, acquire)</code>			load y			
<code>r3 = z</code>					load z	

"Too conservative!"

SCHEDULING ATOMICS

Clock cycle:	1	2	3	4	5	6
<code>r1 = x</code>	load x					
<code>r2 = atomic_load(&y, acquire)</code>	load y					
<code>r3 = z</code>	load z					

SCHEDULING ATOMICS

Clock cycle:	1	2	3	4	5	6
<code>r1 = x</code>	load x					
<code>r2 = atomic_load(&y, acquire)</code>	load y					
<code>r3 = z</code>	load z					

"Too aggressive!"

SCHEDULING ATOMICS

Clock cycle:	1	2	3	4	5	6
<code>r1 = x</code>	load x					
<code>r2 = atomic_load(&y, acquire)</code>	load y					
<code>r3 = z</code>			load z			

SCHEDULING ATOMICS

Clock cycle:	1	2	3	4	5	6
<code>r1 = x</code>	load x					
<code>r2 = atomic_load(&y, acquire)</code>	load y					
<code>r3 = z</code>			load z			

"Just right!"

SCHEDULING ATOMICS

Clock cycle:	1	2	3	4	5	6
<code>r1 = x</code>	load x					
<code>r2 = atomic_load(&y, acquire)</code>	load y					
<code>r3 = z</code>	load z					


"Too aggressive!"

SCHEDULING ATOMICS

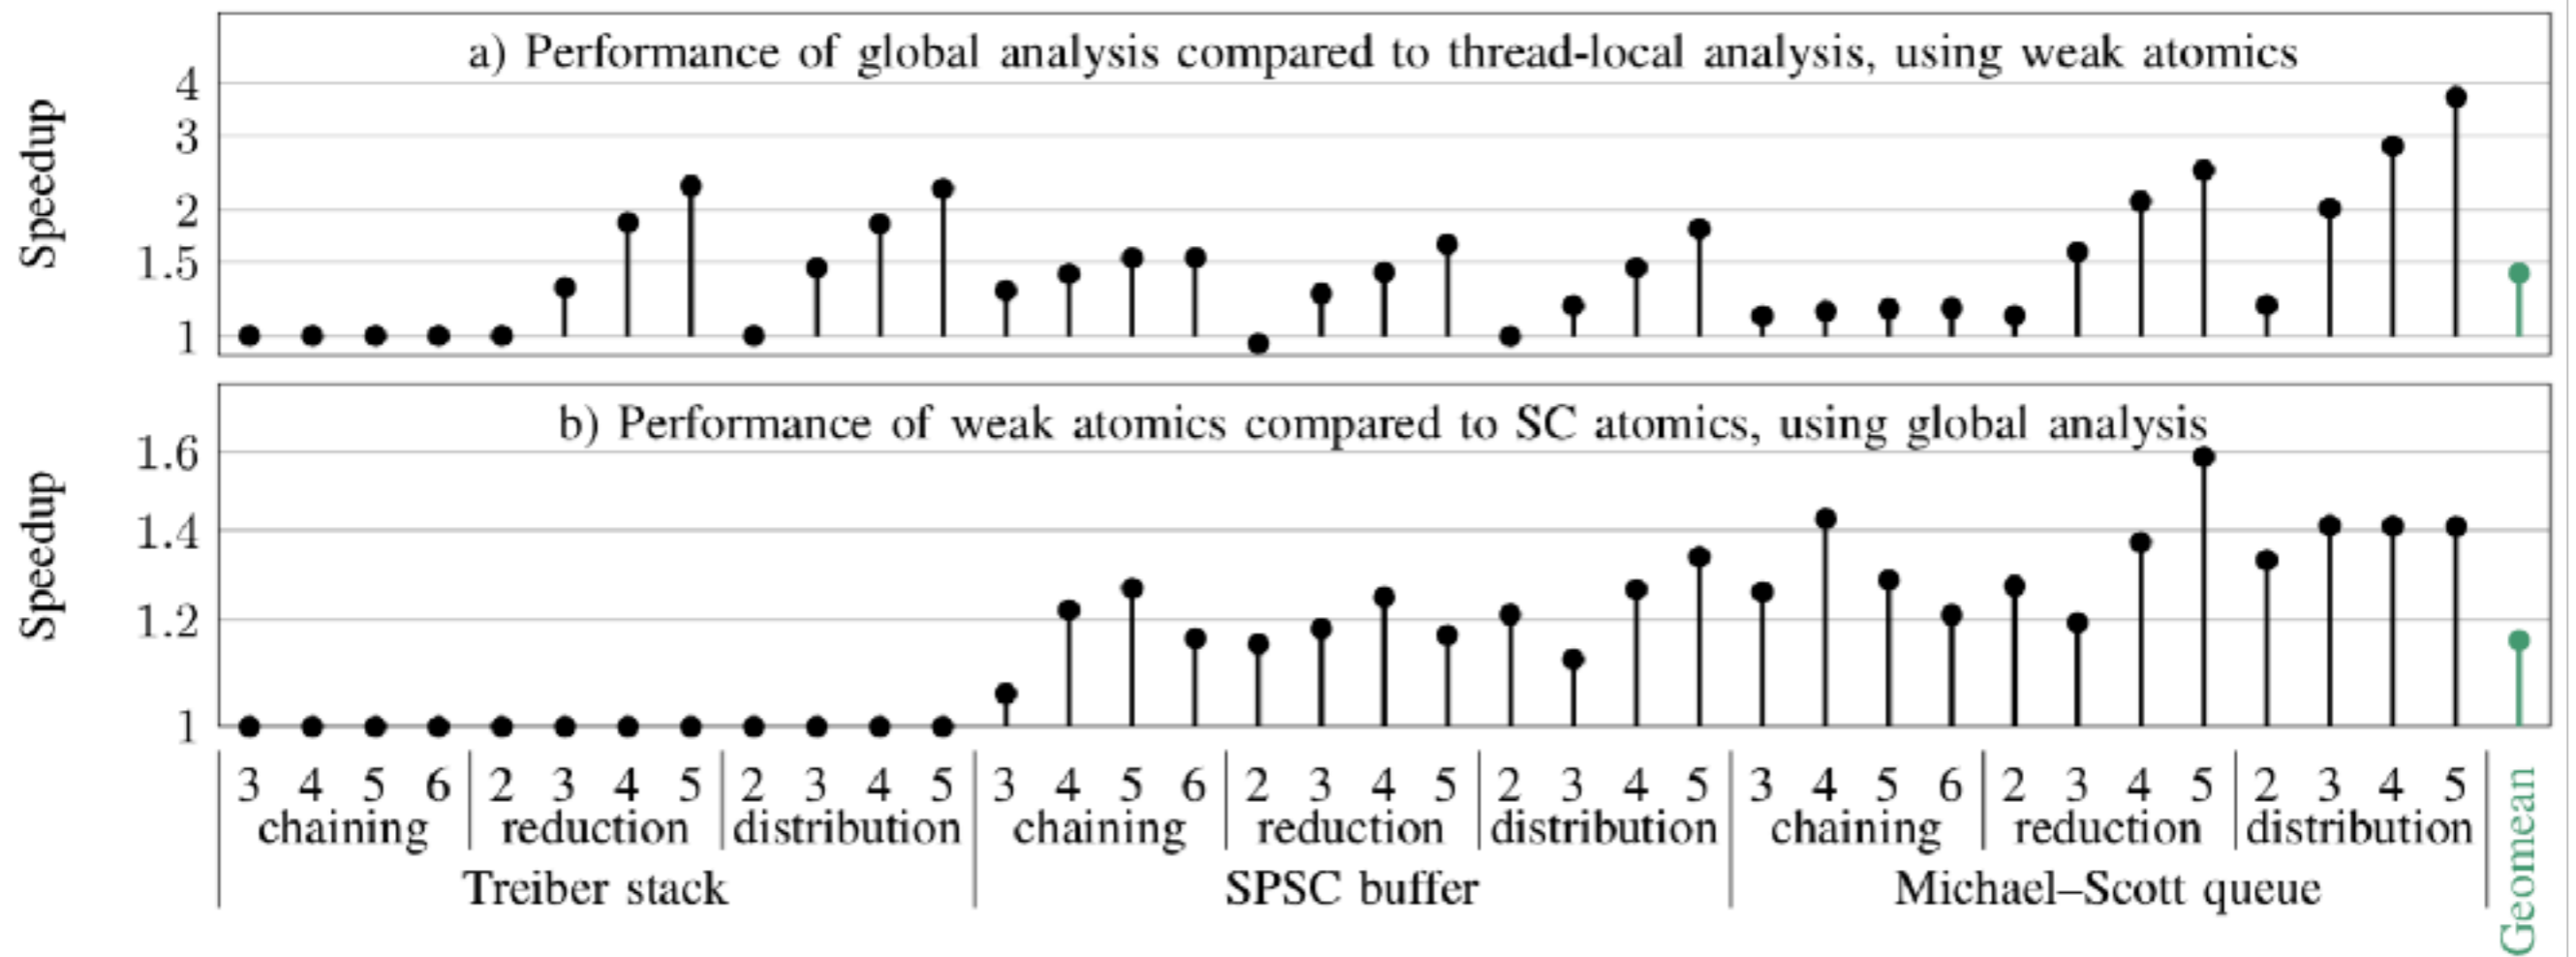
Clock cycle:	1	2	3	4	5	6
<code>r1 = x</code>	load x					
<code>r2 = atomic_load(&y, acquire)</code>	load y					
<code>r3 = z</code>	load z					

"Too aggressive!" ... or is it?

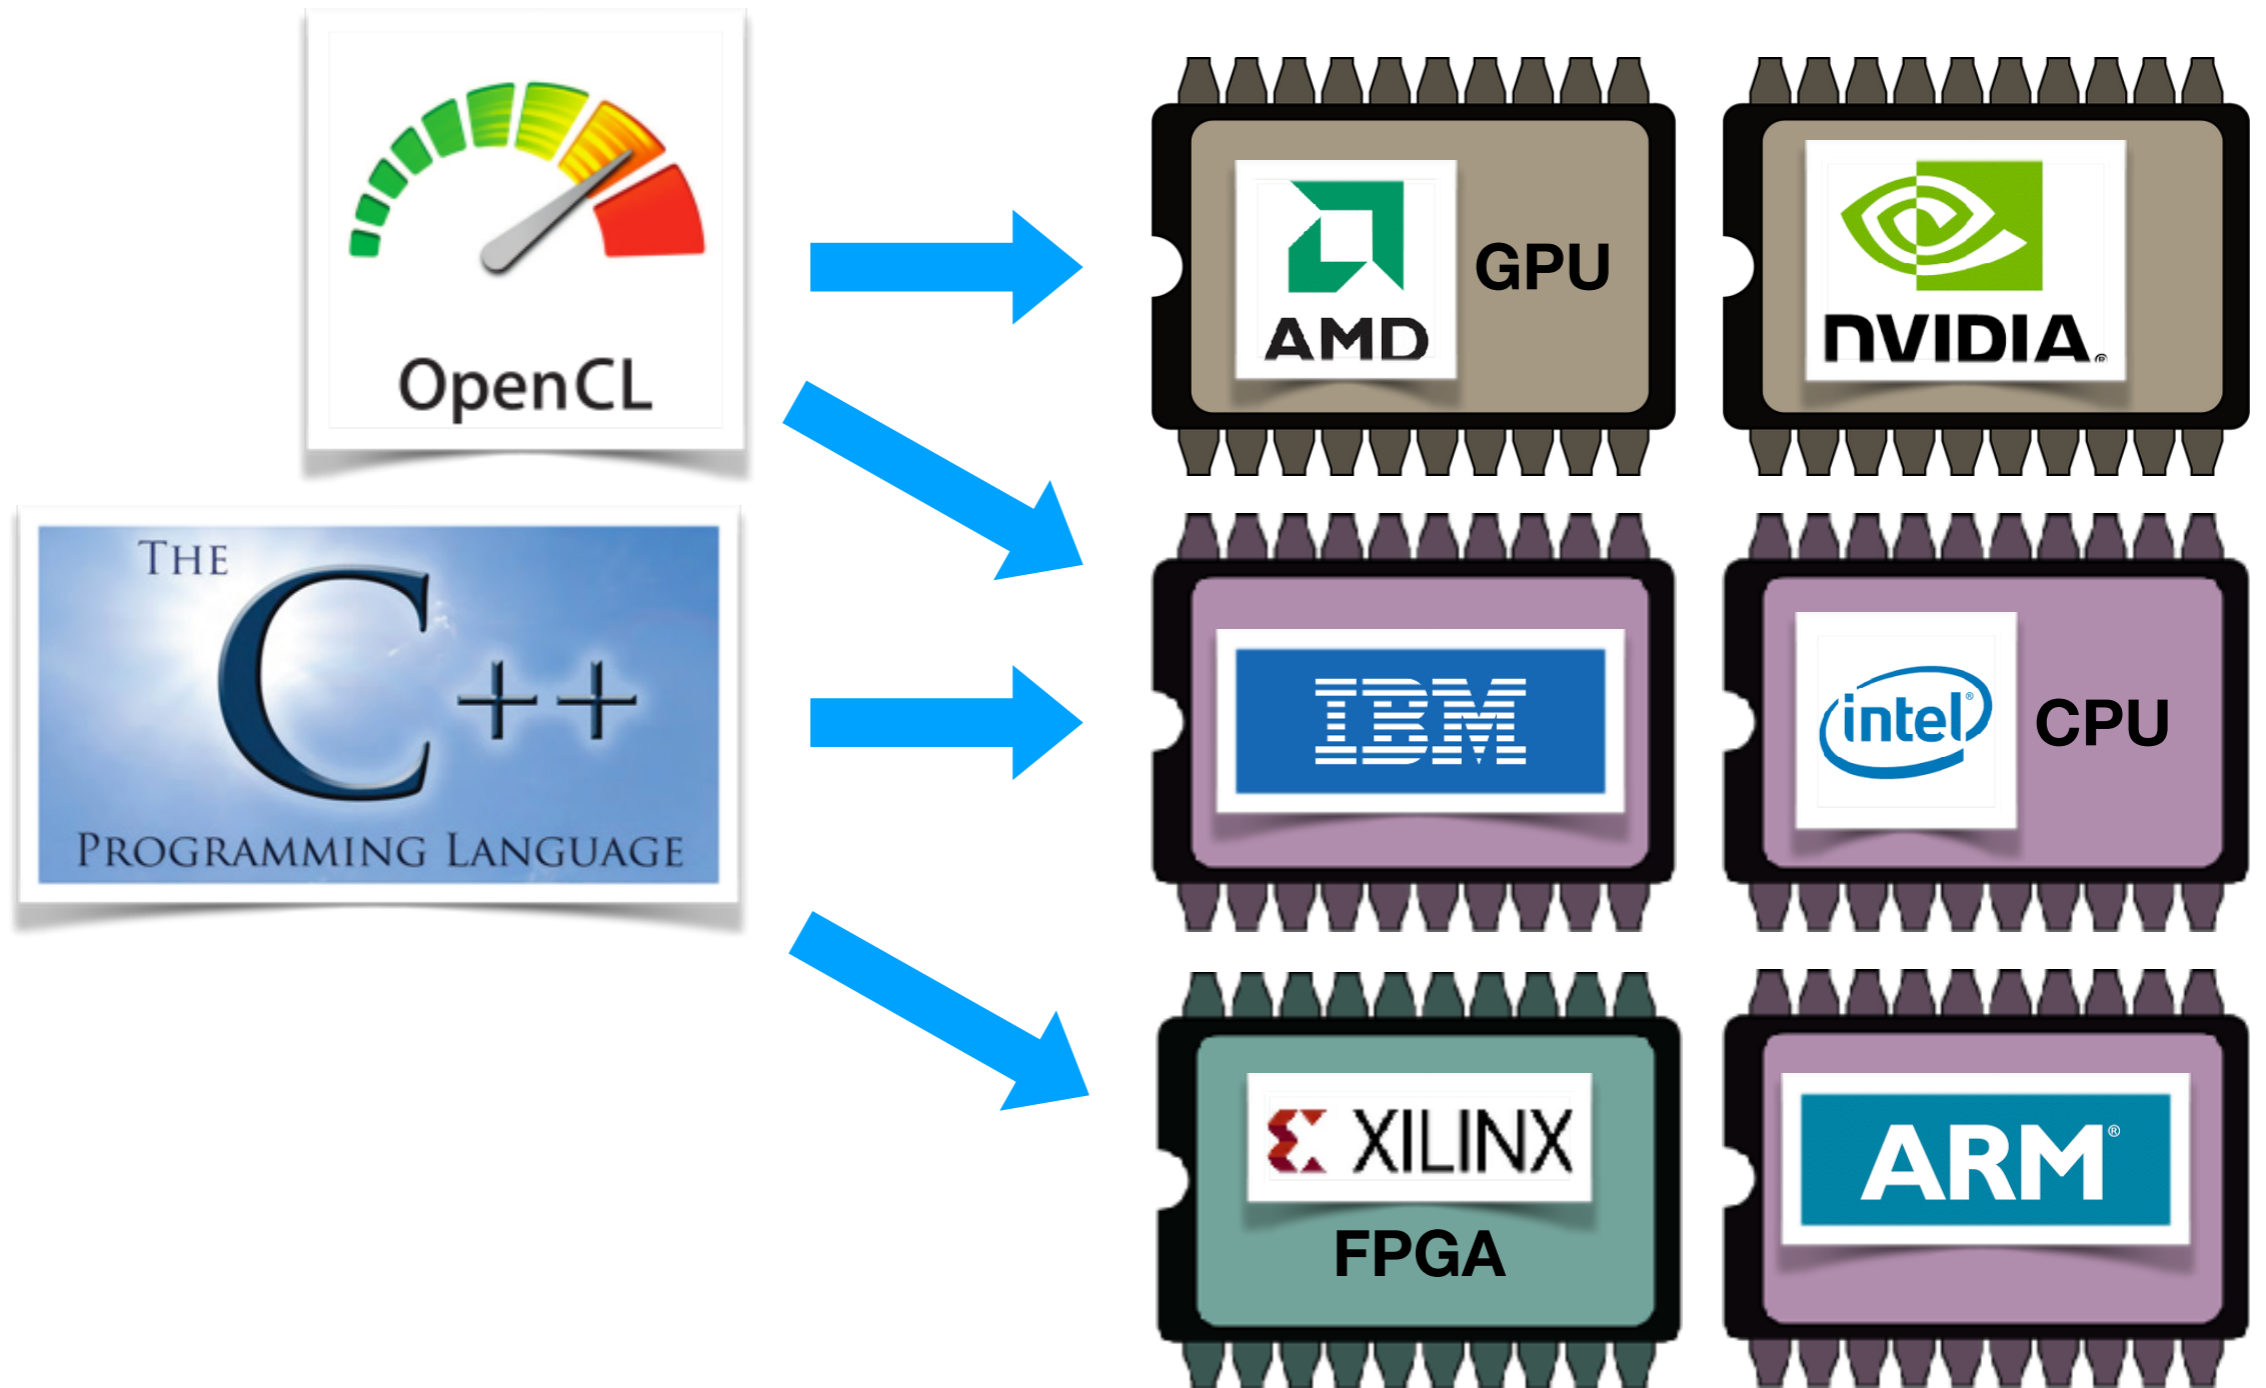
SCHEDULING ATOMICS



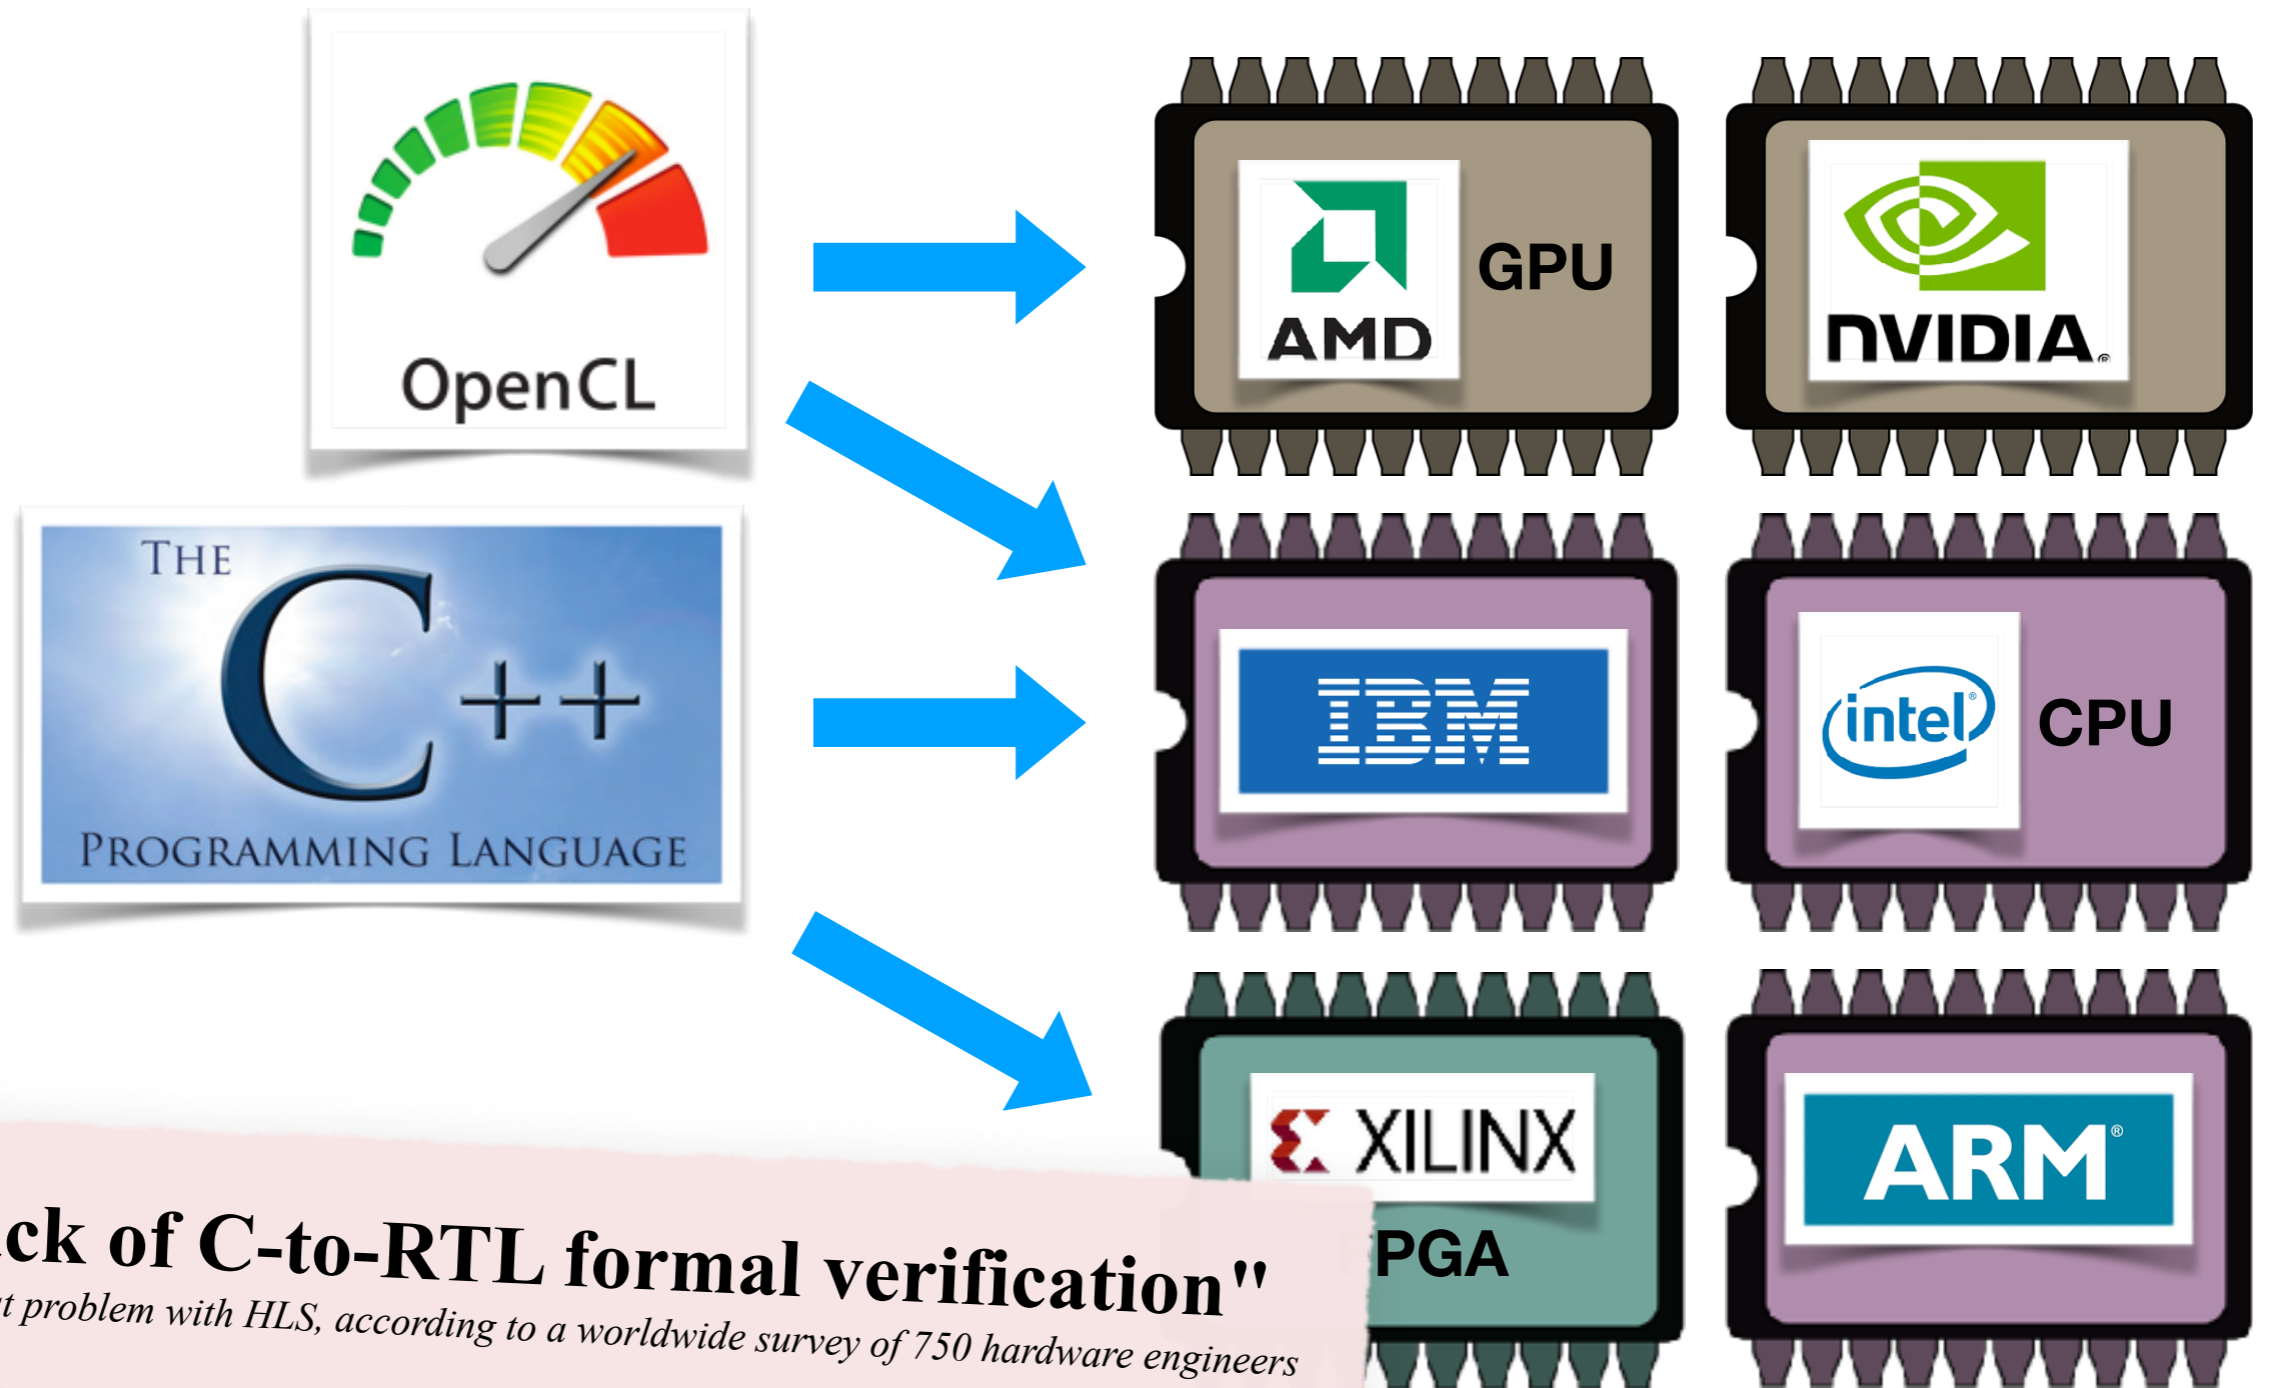
SCHEDULING ATOMICS



THE FUTURE?



THE FUTURE?



"Lack of C-to-RTL formal verification"

The biggest problem with HLS, according to a worldwide survey of 750 hardware engineers