

TAPAS: Generating Parallel Accelerators from Parallel Programs


Steven Margerm¹, **Amirali Sharifian**¹, Apala Guha¹
Gilles Pokam², Arrvindh Shriraman¹



<https://github.com/sfu-arch/tapas>

SFU

Simon Fraser University¹, Intel Corp.²



Motivation




FPGAs are everywhere

- Lots parallelism
 - 150\$ Cyclone V SoC - 60 stencil tasks
- 10s of cycles for invoking a hardware “task”
- Fine-grain parallelism
 - Cyclone V. 512 arithmetic ops



High Level Synthesis



Clang



HLS Compiler



- Mixes schedule and algorithm
 - #pragma
- Static schedule
 - limited concurrency control
- Domain specific templates
 - generalizable ?

TAPAS: Auto generating Parallel Dataflow Accelerator




- **MIT's parallel compiler(TAPIR)**




- **Hardware component library:**
 - like UCB Rocket, but for accelerators



- **Generator**
 - synthesizing RTL from compiler IR



TAPAS



Parallel Accelerator

Cilk/Go/OpenMP

Overview

- HLS Challenge: Static Parallelism
- TAPAS : modular high level synthesis
- TAPAS: generating task units

HLS Challenge: Static Parallelism

```
for(i = 0 until n){  
  if(node[i].valid){  
    compute(&node[i]);  
  }  
}
```

Loop bound is
unknown

Conditional
Body

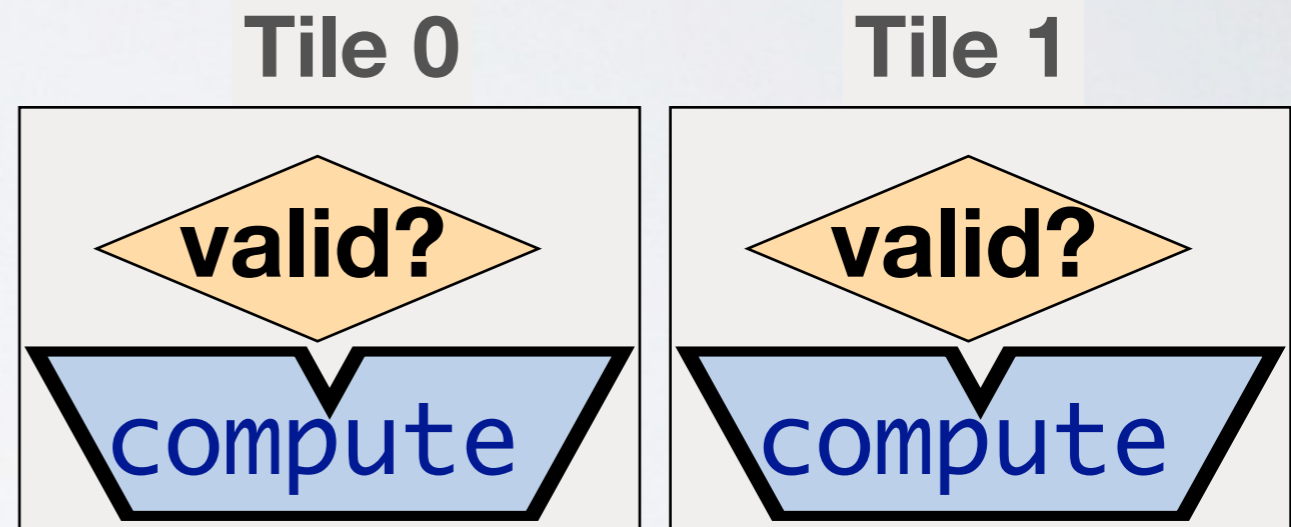
Non-
deterministic
latency

HLS Challenge: Static Parallelism

Unrolled Program

```
#pragma UNROLL 2  
for(i = 0 until n){  
  if(node[i].valid){  
    compute(&node[i]);  
  }  
}
```

Hardware



Worst case schedule \rightarrow Low utilization

Our Approach: Dynamic Parallelism

Task Program


```
for(i = 0 until n){  
  if(node[i].valid){  
Spawn compute(&node[i]);  
  }  
Sync;
```

Task Hardware

Root

```
for (); if (valid)
```

Spawn



Sync

Child

compute


Run time schedule —> High utilization

Compilation Flow


Cilk/Go



Task Graph Representation



Task-Level RTL



TAPAS Accelerator

Compilation Flow

Cilk/Go

Task extraction

Task Graph Representation

Top RTL Generation

Task-Level RTL

Generate TXUs

TAPAS Accelerator

```
for(i = 0 until n){  
  if(node[i].valid){  
    compute(&node[i])  
  }  
}
```

Compilation Flow

Cilk/Go

Static Task Graph



Task extraction

Task Graph Representation

Top RTL Generation

Task-Level RTL

Generate TXUs

TAPAS Accelerator

Root

Child

Compilation Flow

Cilk/Go

Task extraction

Task Graph Representation

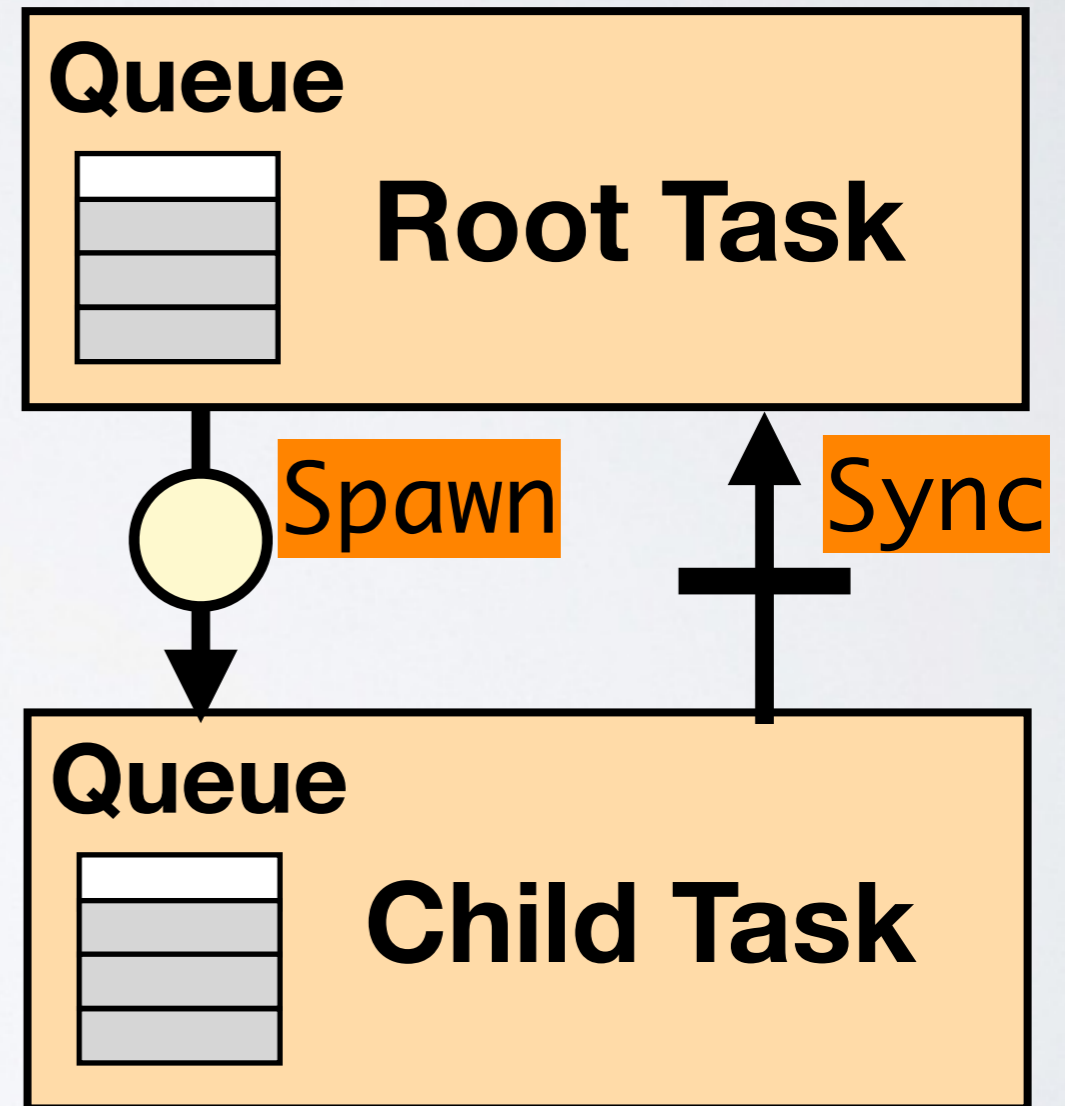
Top RTL Generation

Task-Level RTL

Generate TXUs

TAPAS Accelerator

Task-level
Architecture



Compilation Flow

Cilk/Go

Task extraction

Task Graph Representation

Top RTL Generation

Task-Level RTL

Generate TXUs

TAPAS Accelerator

Task Execution Unit

Root

```
for (); if (valid)
```

Spawn



Sync

Child

compute

Static Task Graph

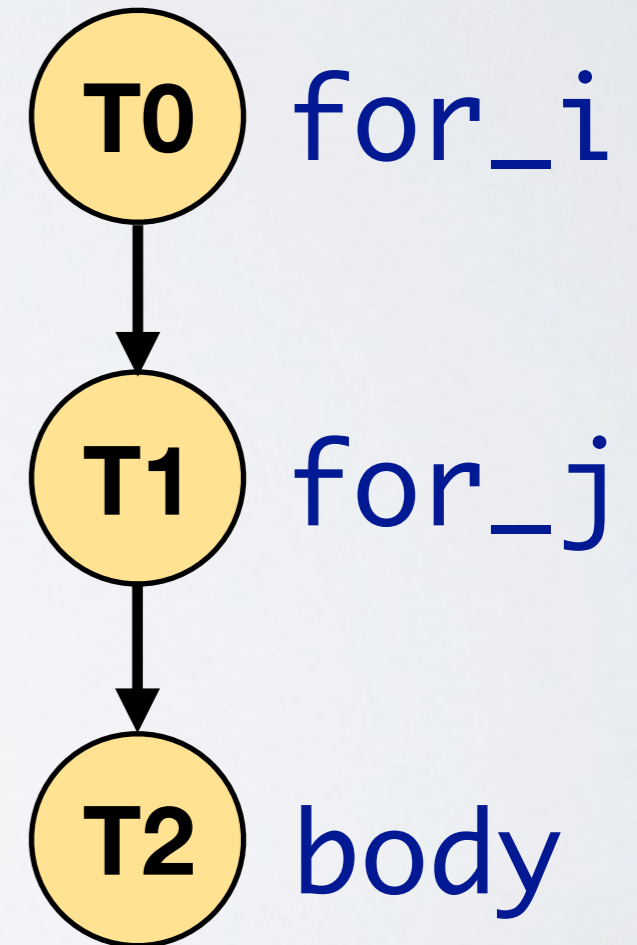
```
cilk_for(i = 0 until n){  
  cilk_for(j = 0 until n){  
    c[i][j] =  
      a[i][j]+b[i][j];  
  }  
}
```

- **Parallel Compiler**

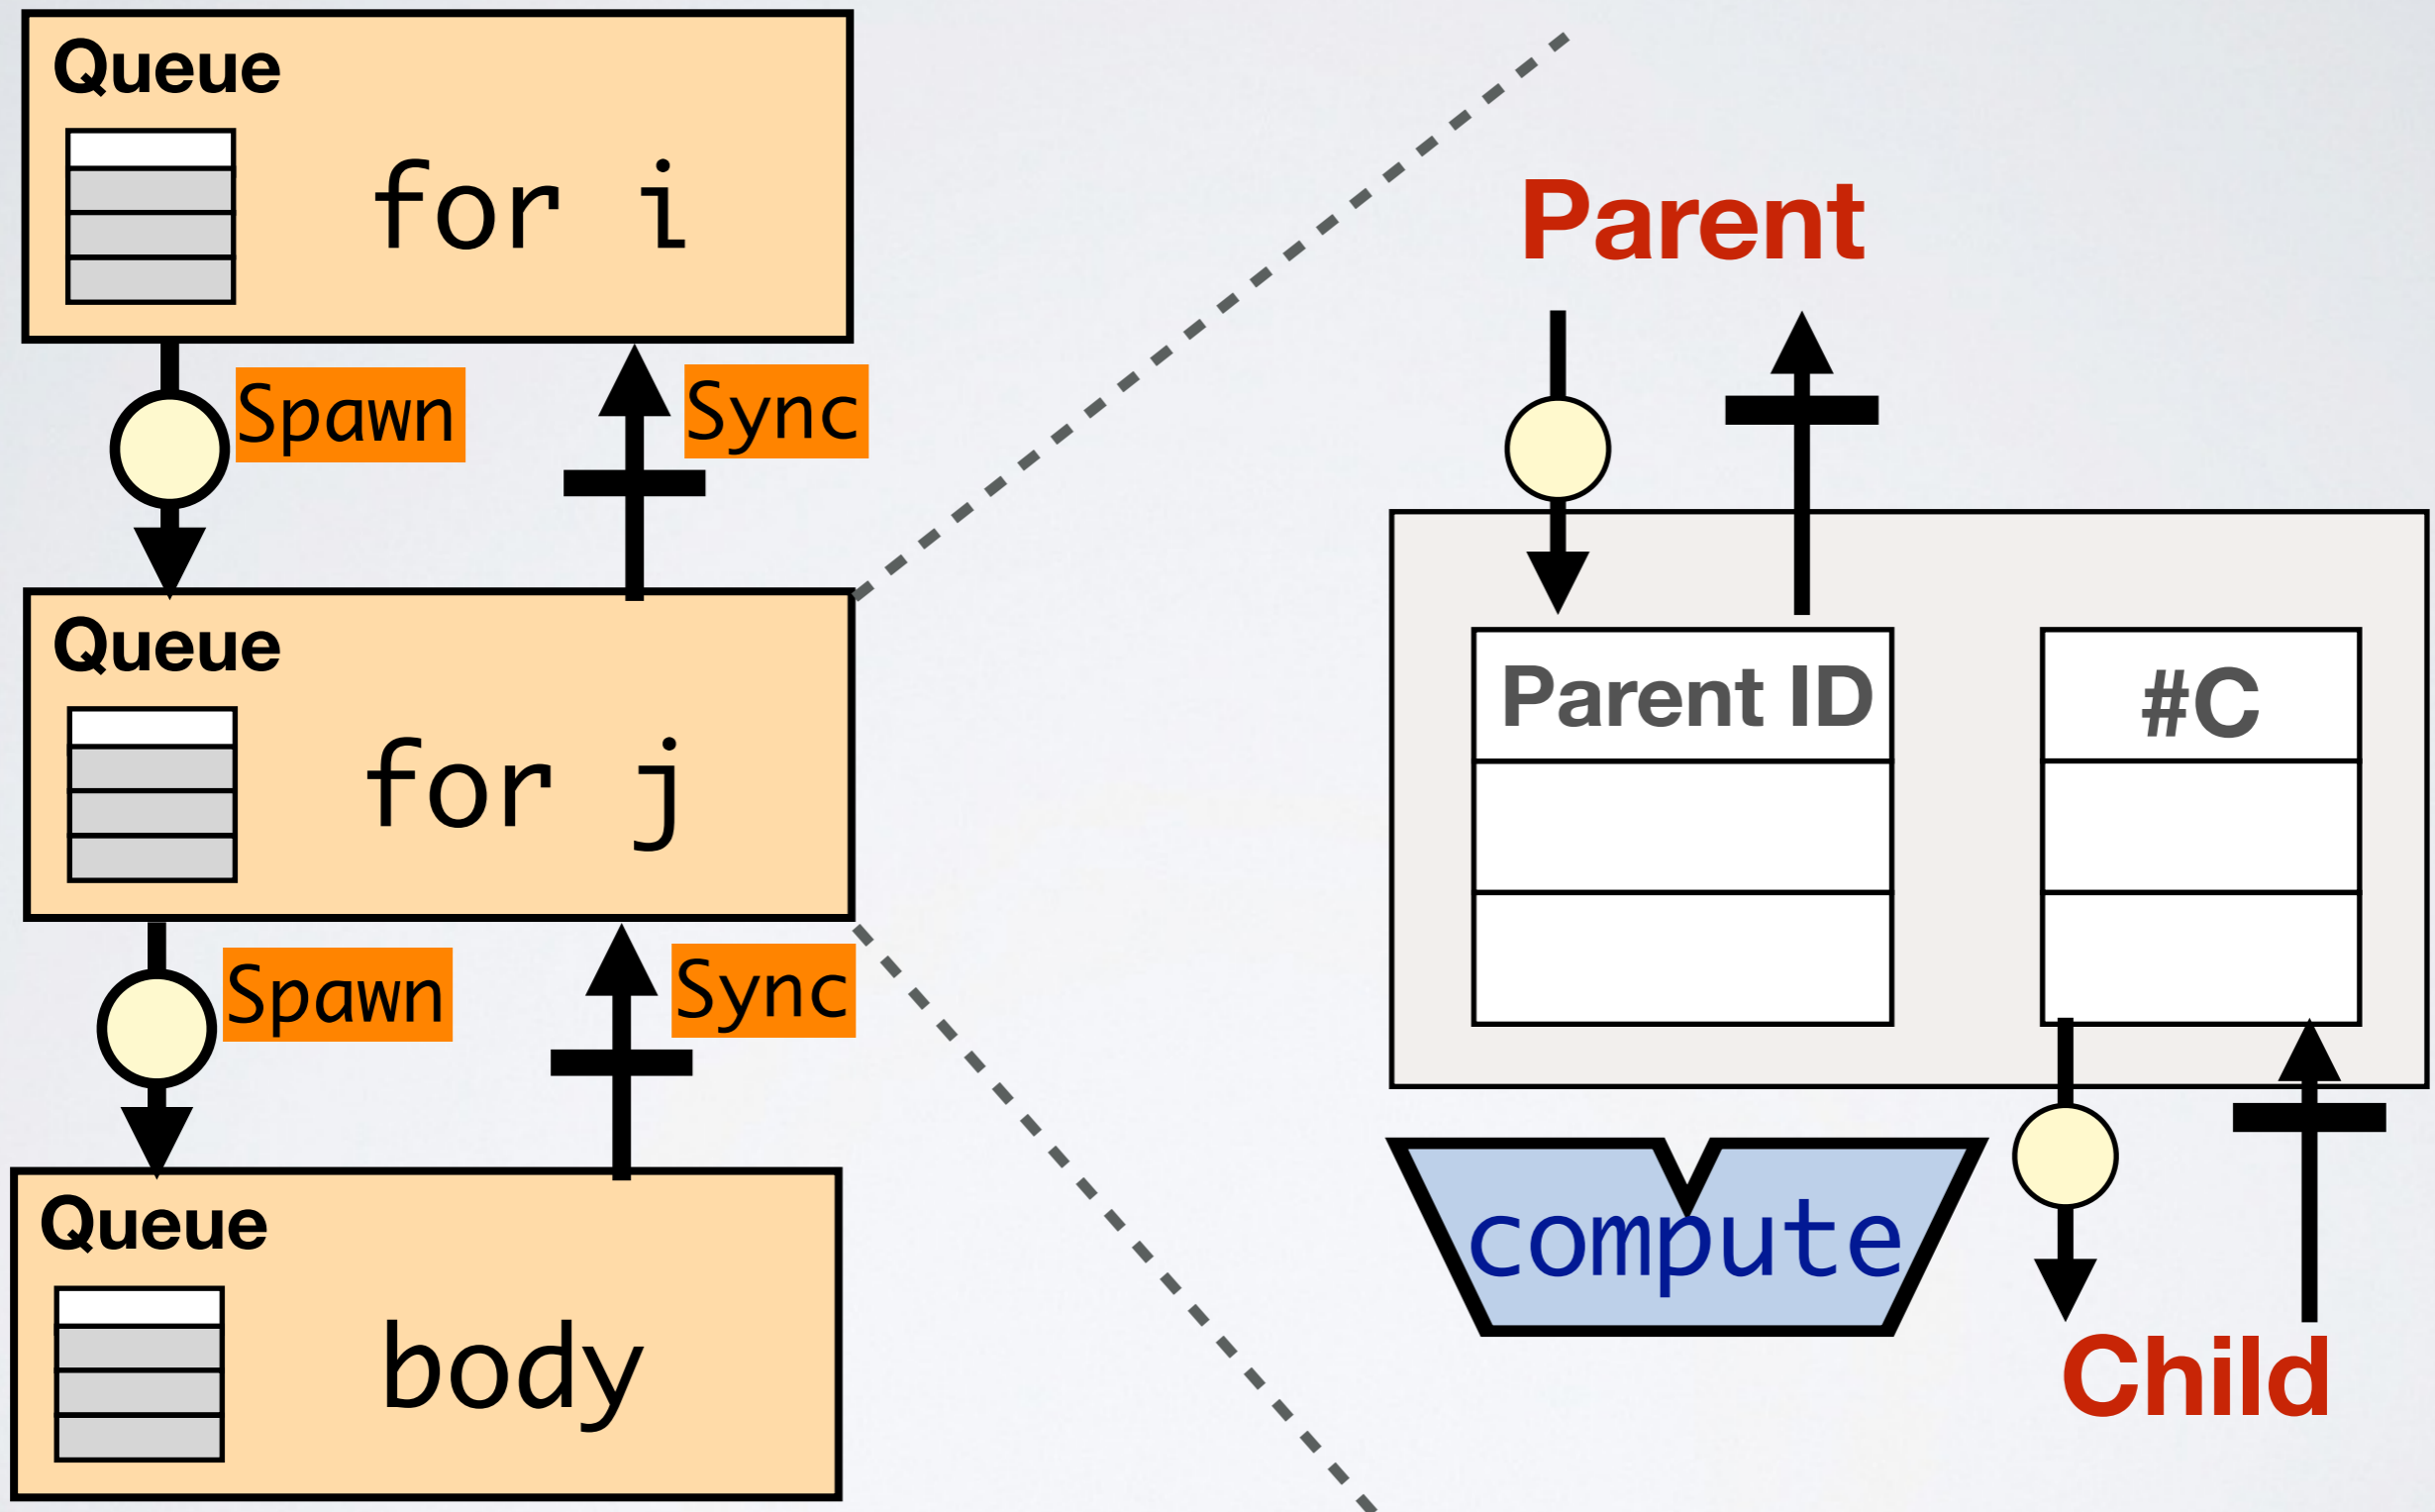
- Captures Spawn and Sync from IR

- **Task Extractor:**

- Wraps each task in a first class entity



Task-Level Architecture




Heterogeneous!

Nested Parallel!

Asynchronous!

Task Level Execution



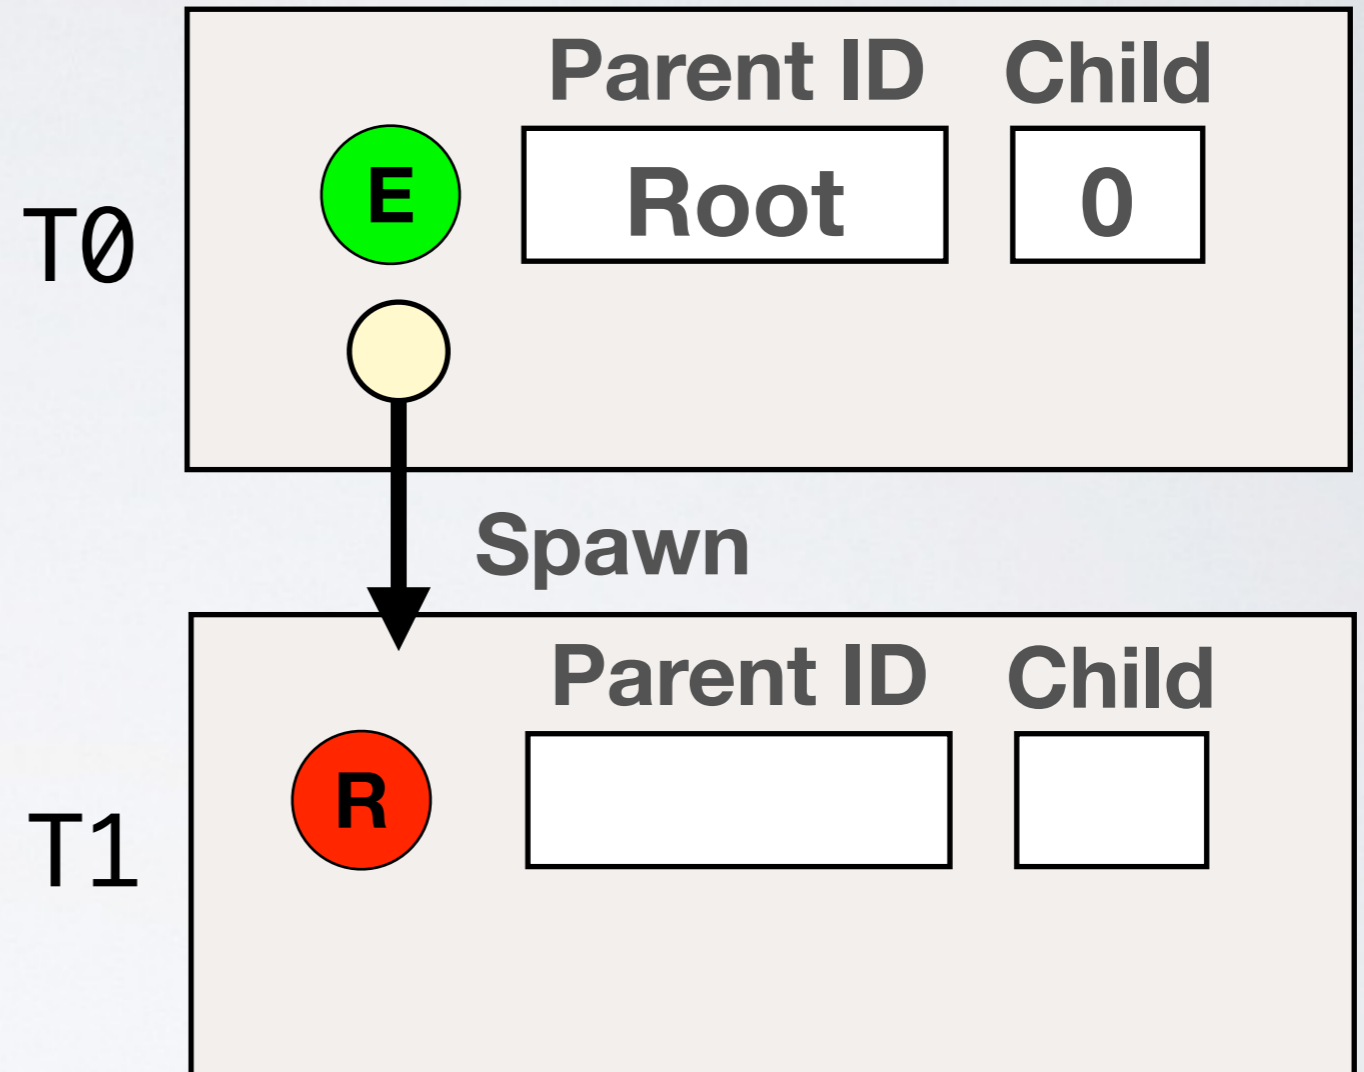

T0

	Parent ID	Child
E	Root	0


T1

	Parent ID	Child
R		

Task Level Execution



Task Level Execution



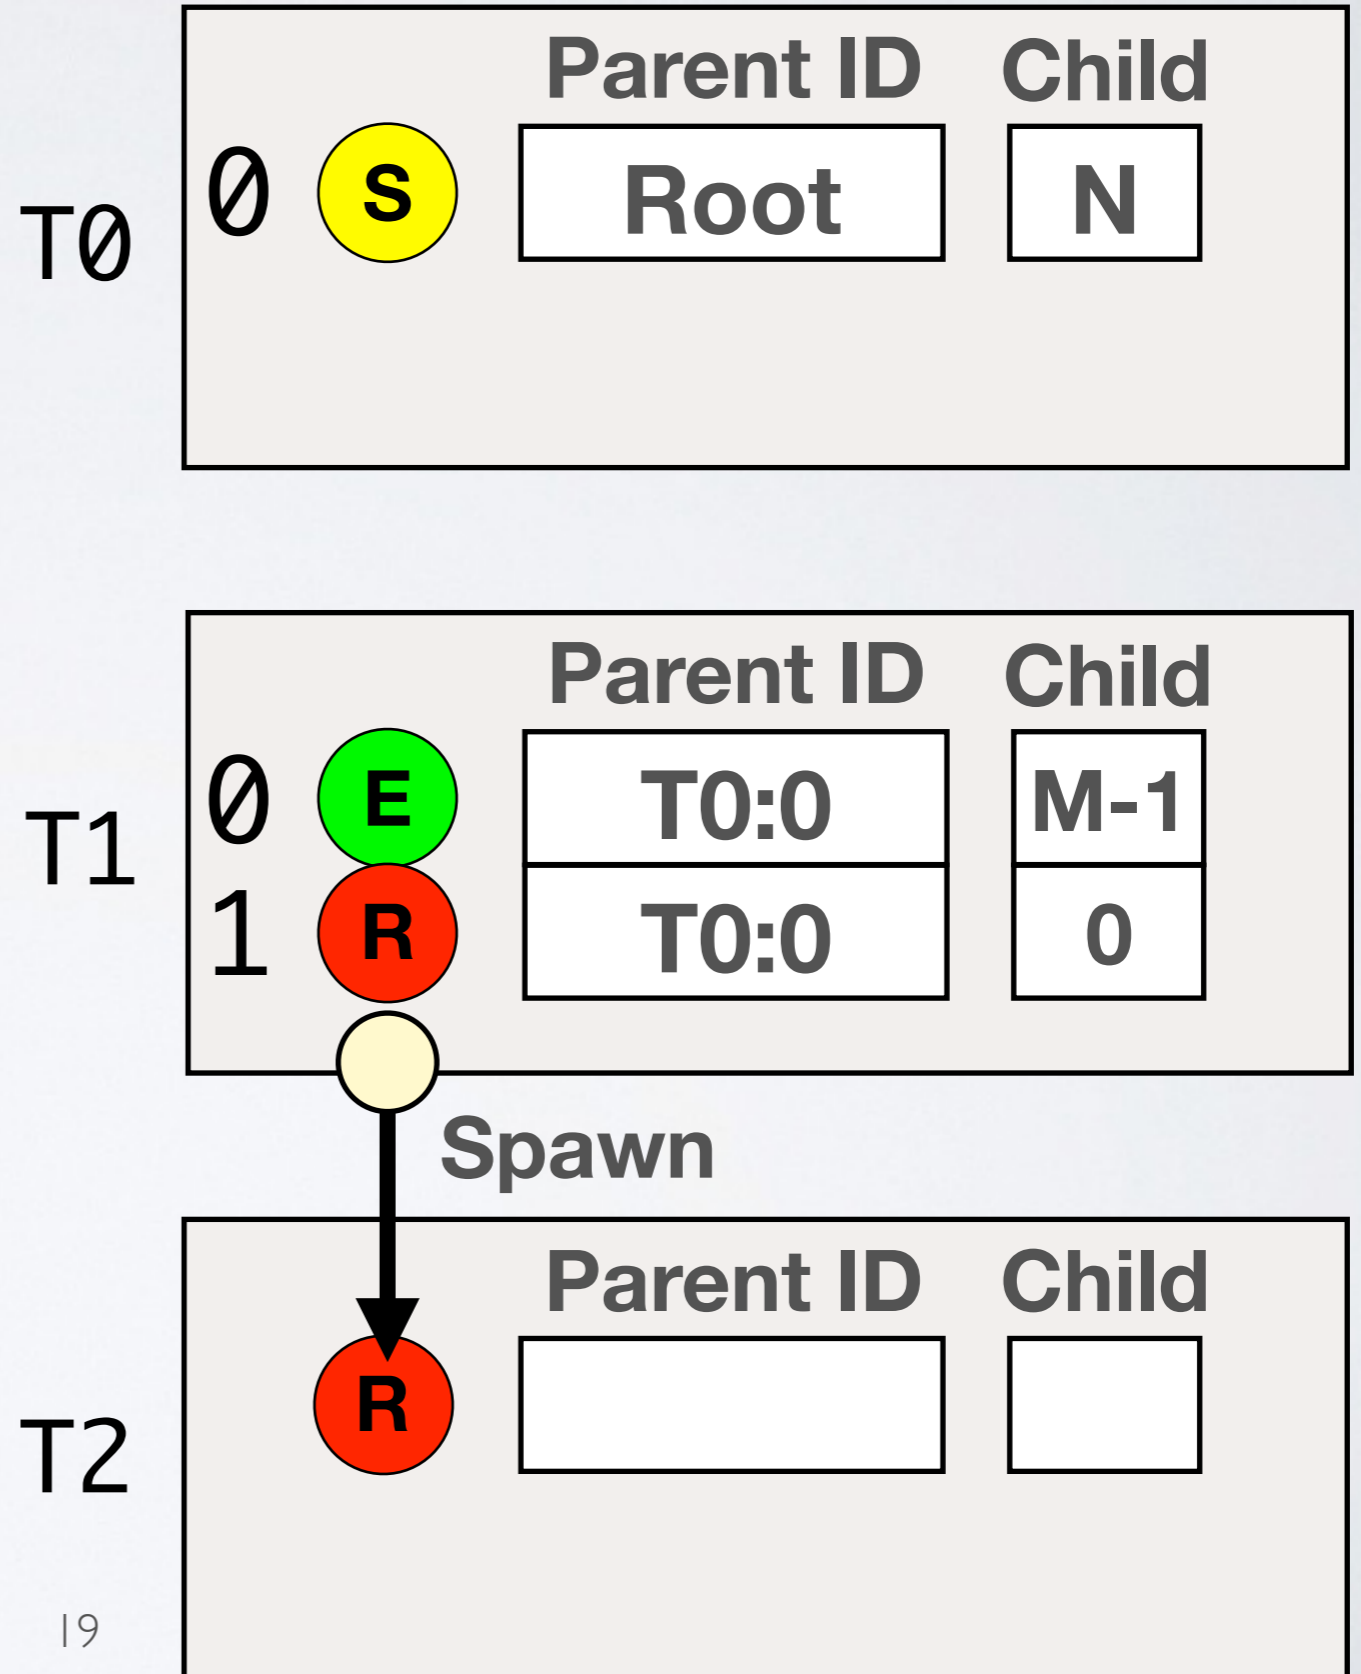
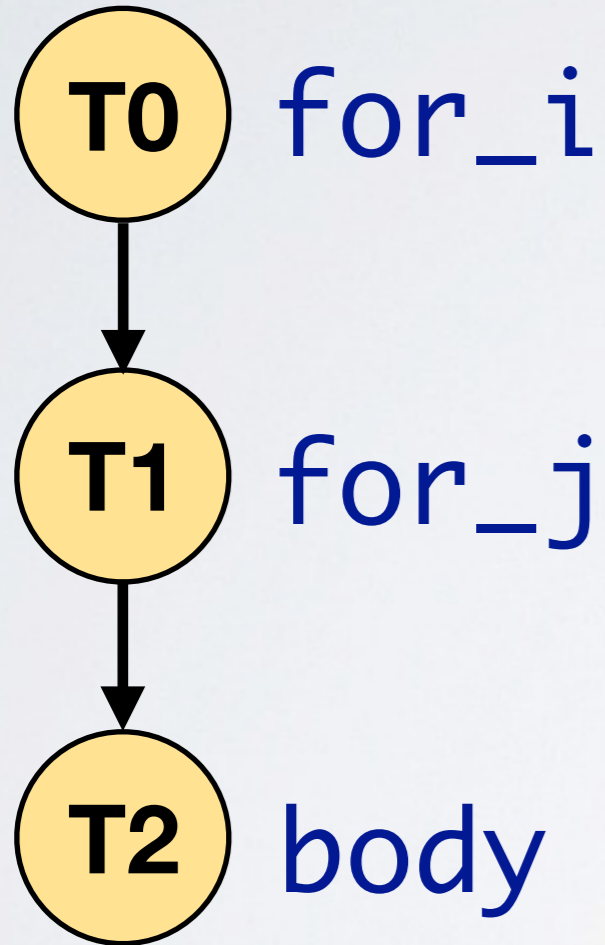
T0

		Parent ID	Child
0	E	Root	1


T1

		Parent ID	Child
0	E	T0:0	0

Task-Level Execution



Task Level Execution



T0

		Parent ID	Child
0	S	Root	N

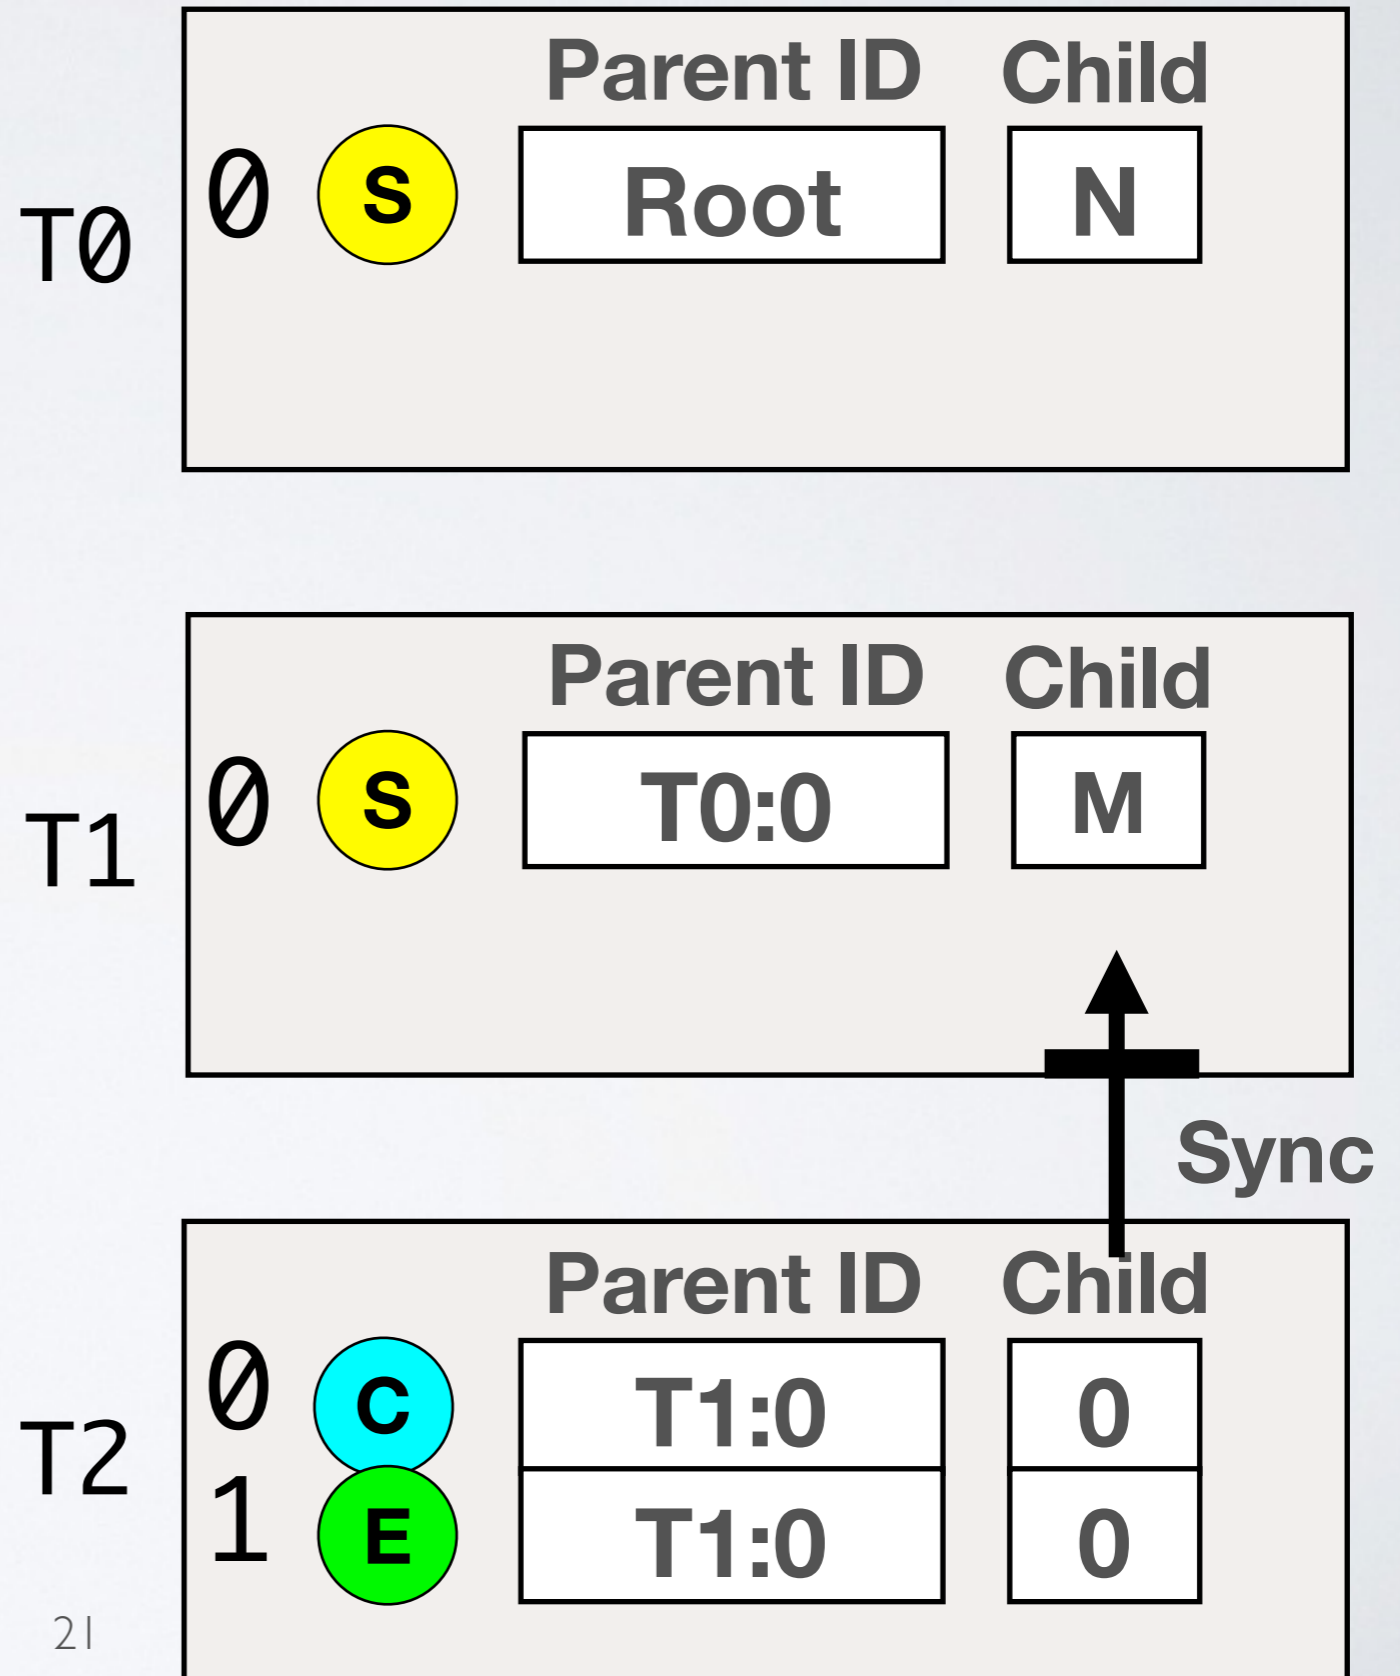

T1

		Parent ID	Child
0	S	T0:0	M
1	E	T0:0	0


T2

		Parent ID	Child
0	E	T1:0	0

Task Level Execution



Task Level Execution



T0

		Parent ID	Child
0	S	Root	N

T1

		Parent ID	Child
0	S	T0:0	M-1

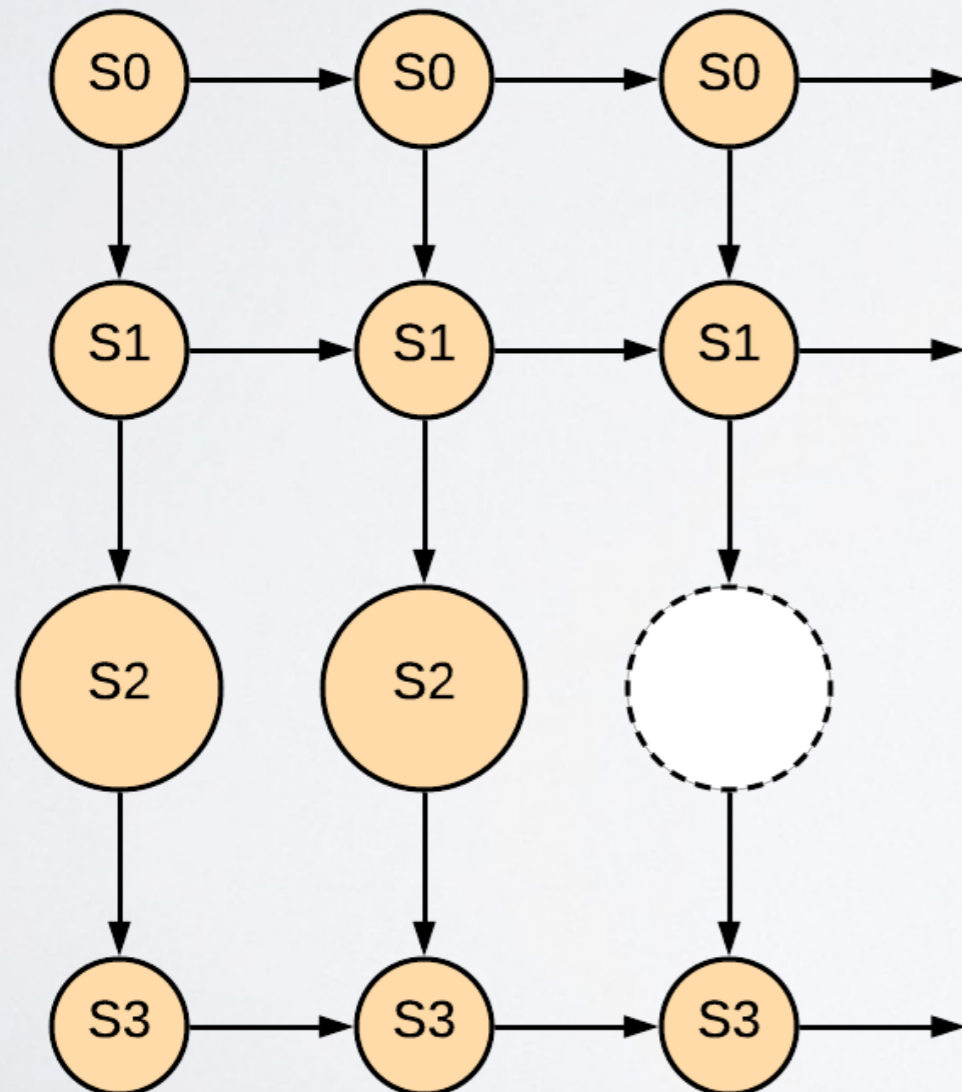
T2

		Parent ID	Child
0	E	T1:0	0

Supporting Dynamic Task Graphs

- What does task pipelining look like in TAPAS?


Dedup



Supporting Dynamic Task Graphs

- What does task pipelining look like in TAPAS?

Dedup

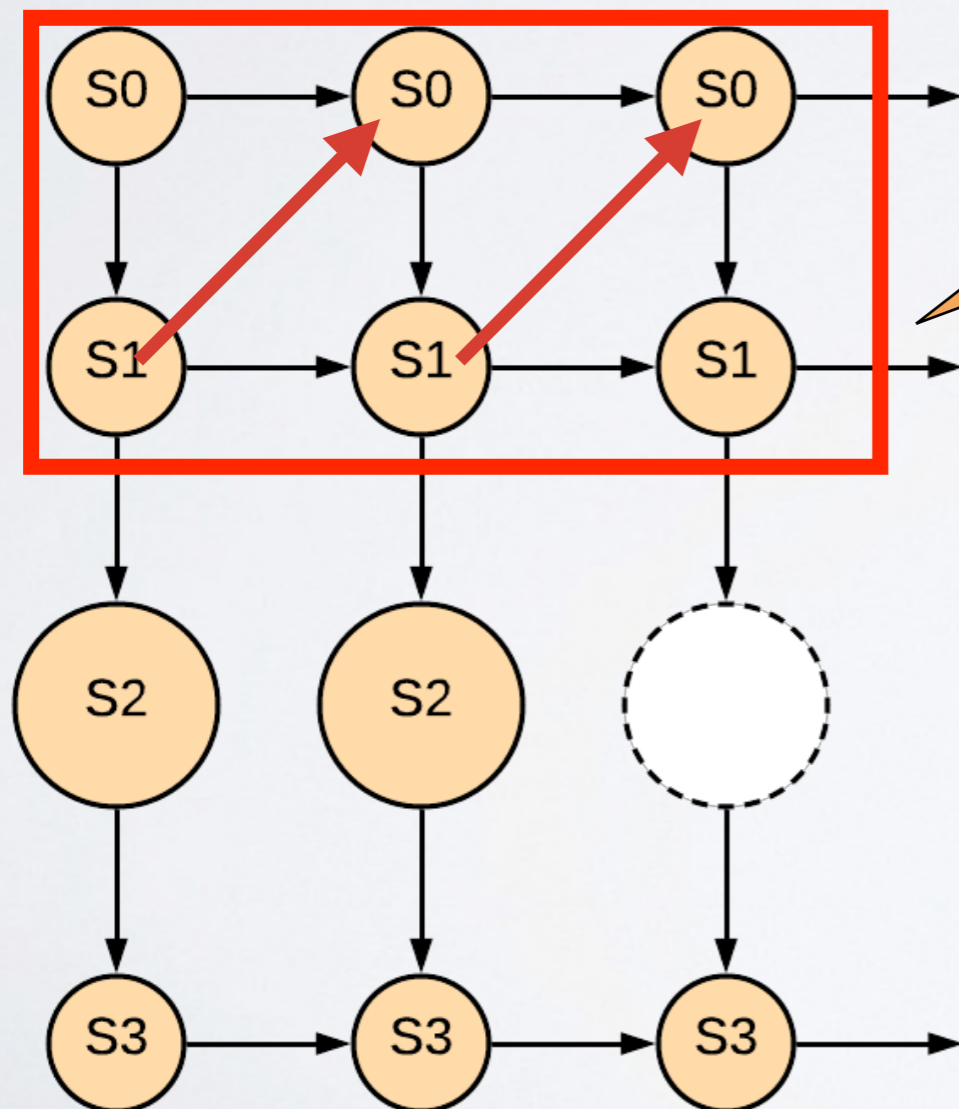


Regular task pipeline

Supporting Dynamic Task Graphs

- What does task pipelining look like in TAPAS?

Dedup

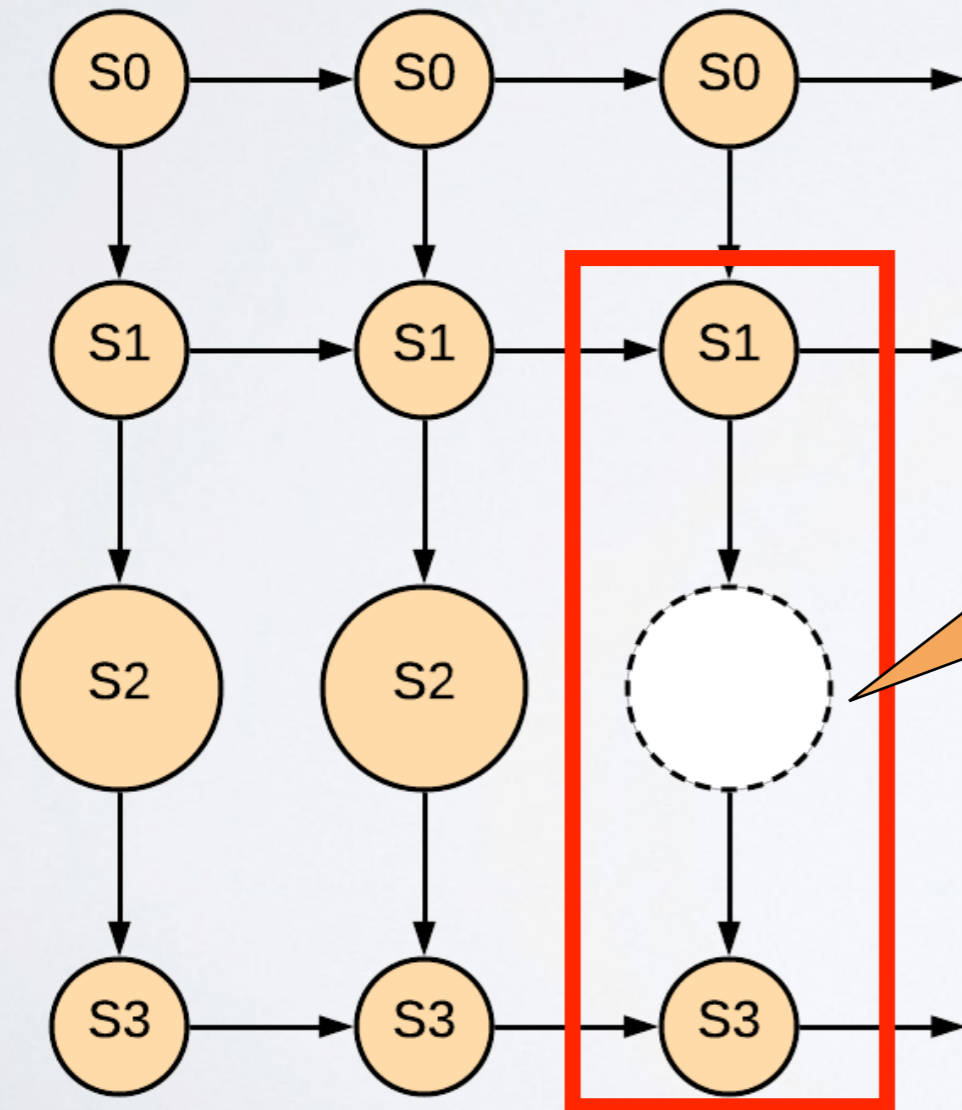


Irregular task pipeline

Supporting Dynamic Task Graphs

- What does task pipelining look like in TAPAS?

Dedup



Conditional task pipeline

**TAPAS supports
recursive tasks as well
(Paper: §5)**

Compilation Flow

Cilk/Go

Task extraction

Task Graph Representation

Top RTL Generation

Task-Level RTL

Generate TXUs

TAPAS Accelerator

Task Execution Unit

Root

```
for (); if (valid)
```

Spawn



Sync



Child


compute

Task Execution Unit (TXU)

- What are the element inside each TXU?




Task Queue

Asynchronous!



Available task!

Dynamic issue

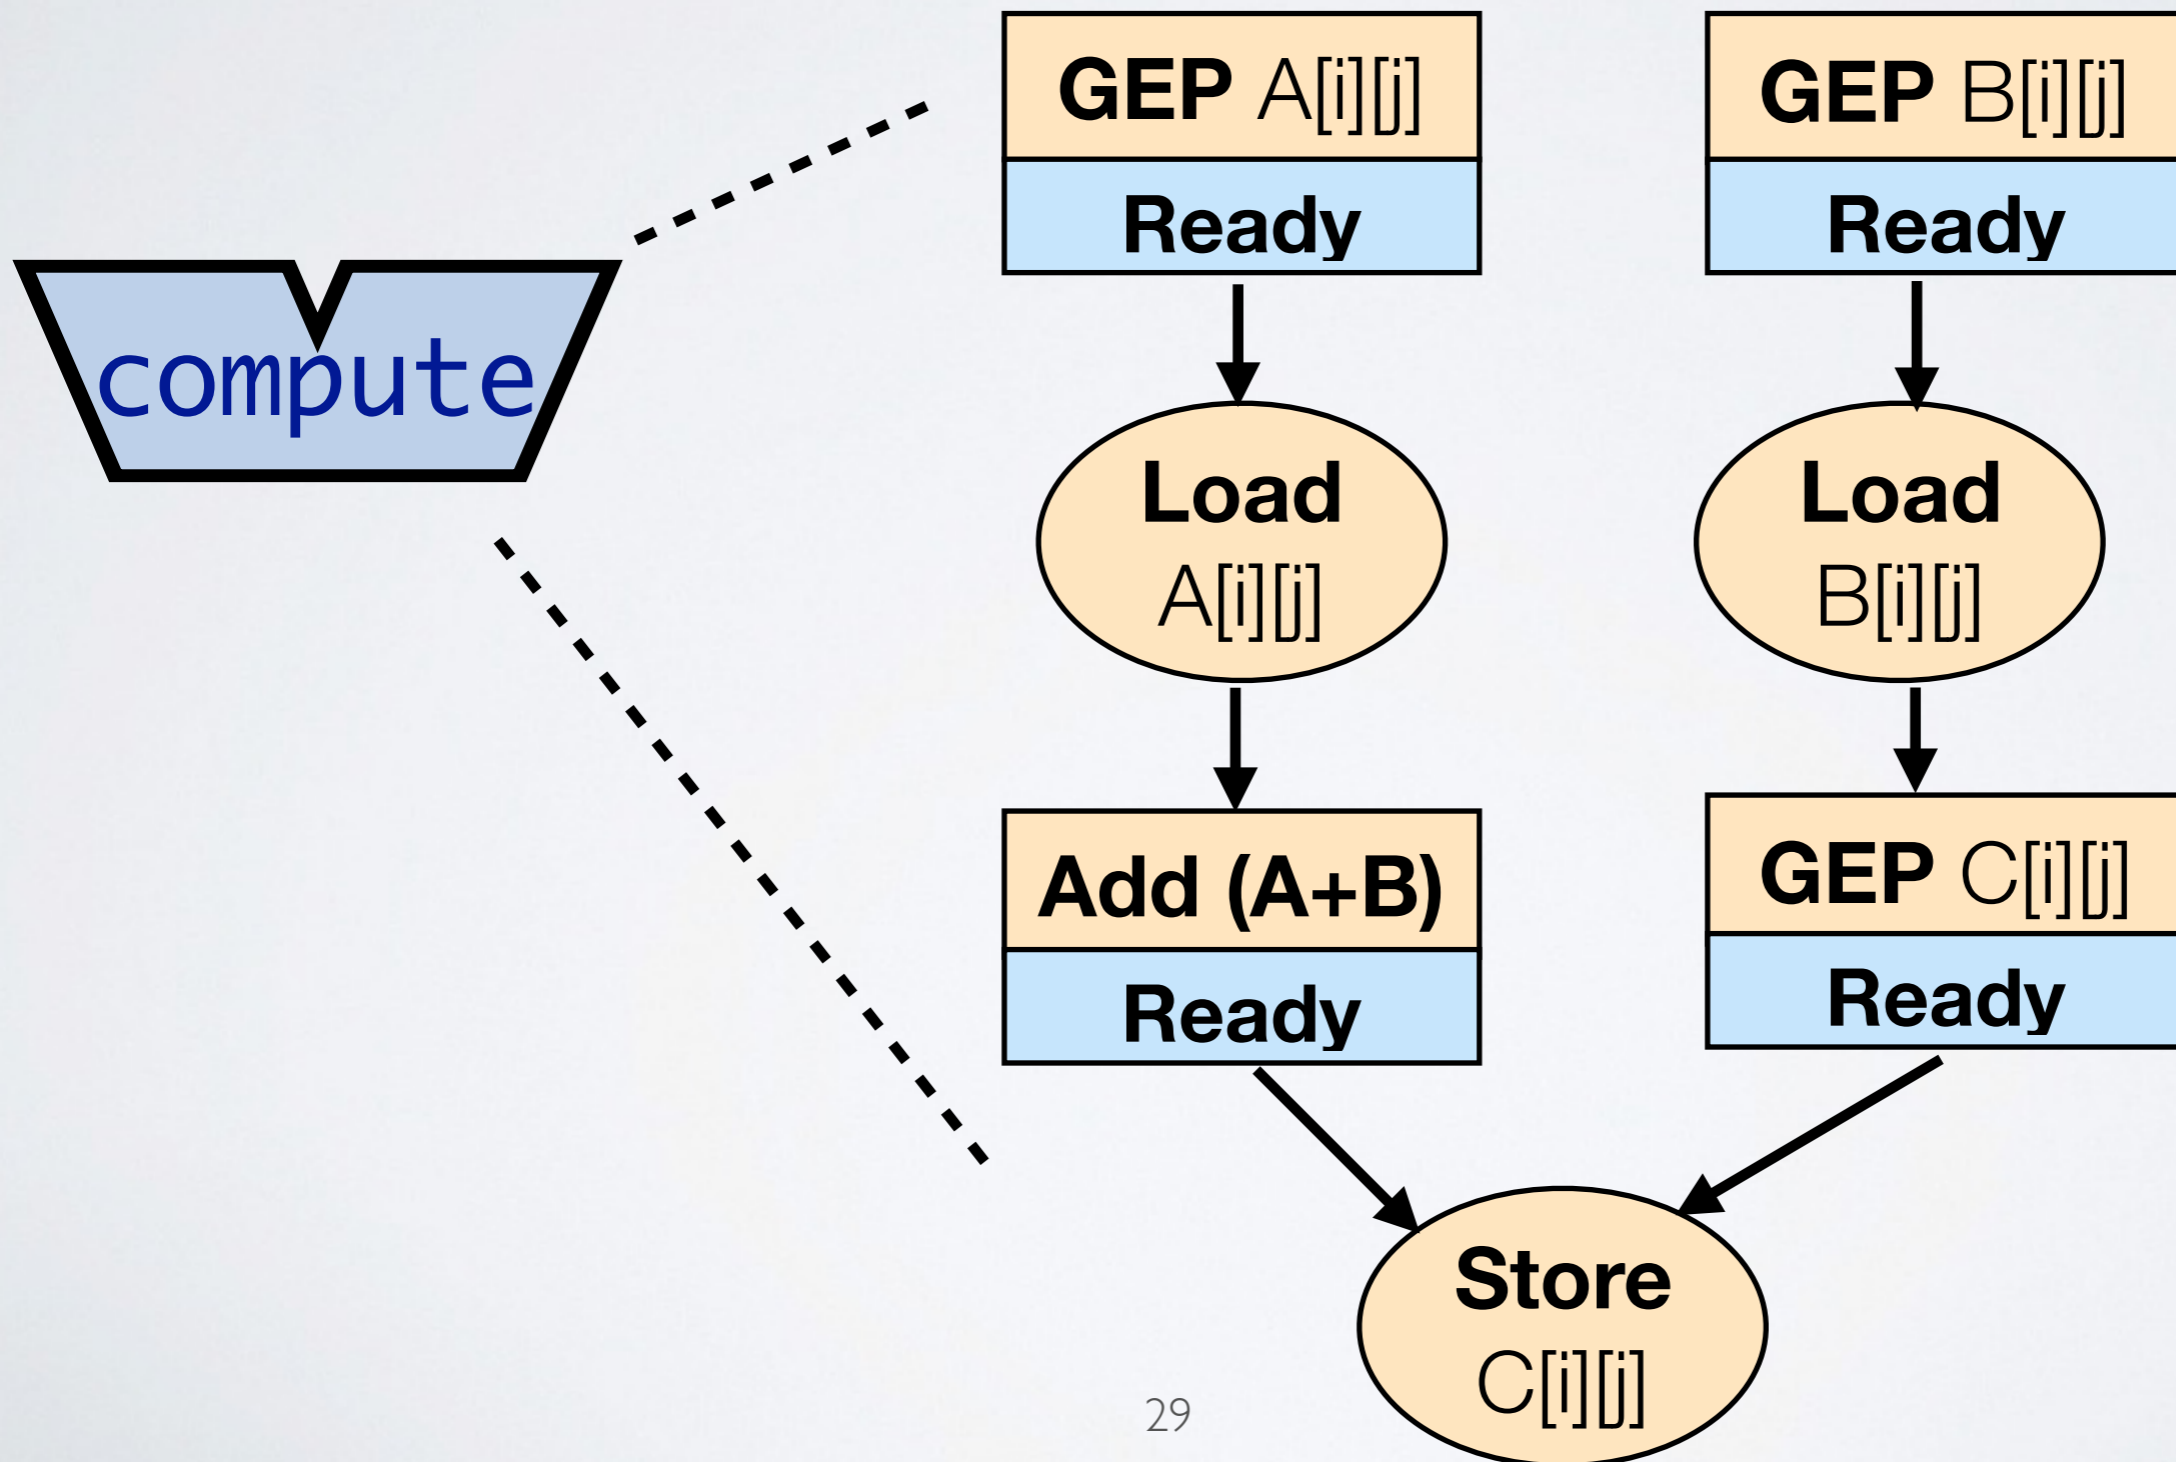


Multi-core

Ready

Task Execution Unit (TXU)

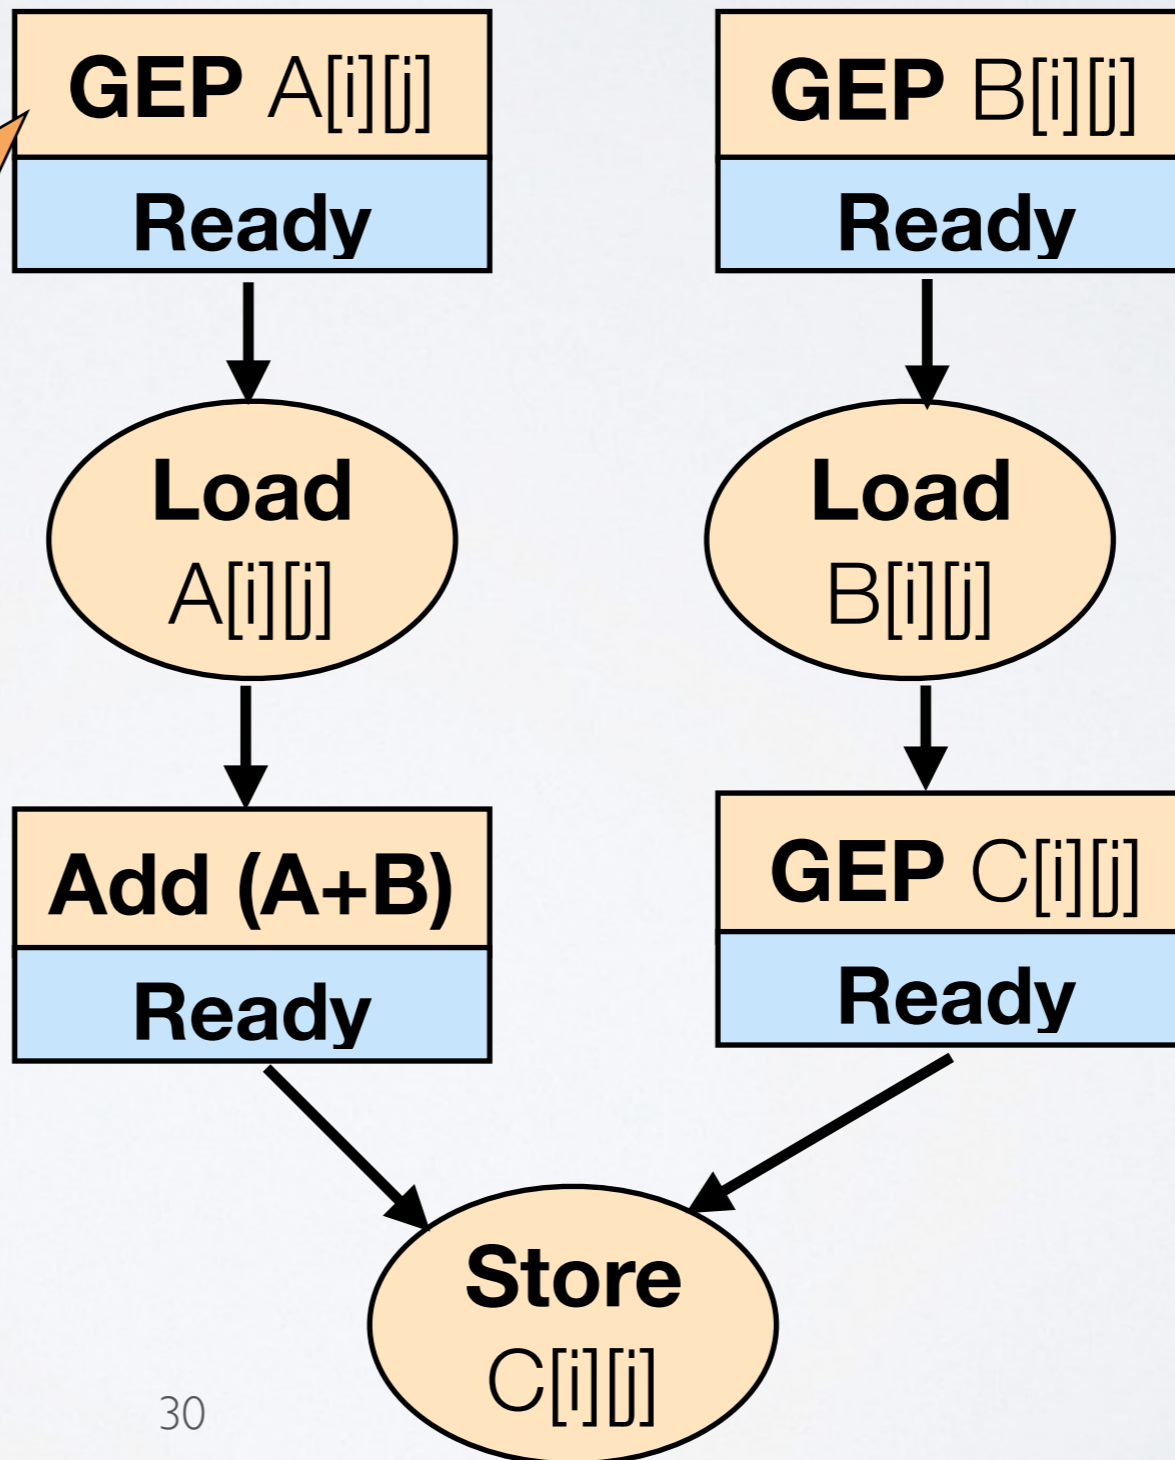
$$c[i][j] = a[i][j] + b[i][j];$$



Task Execution Unit (TXU)

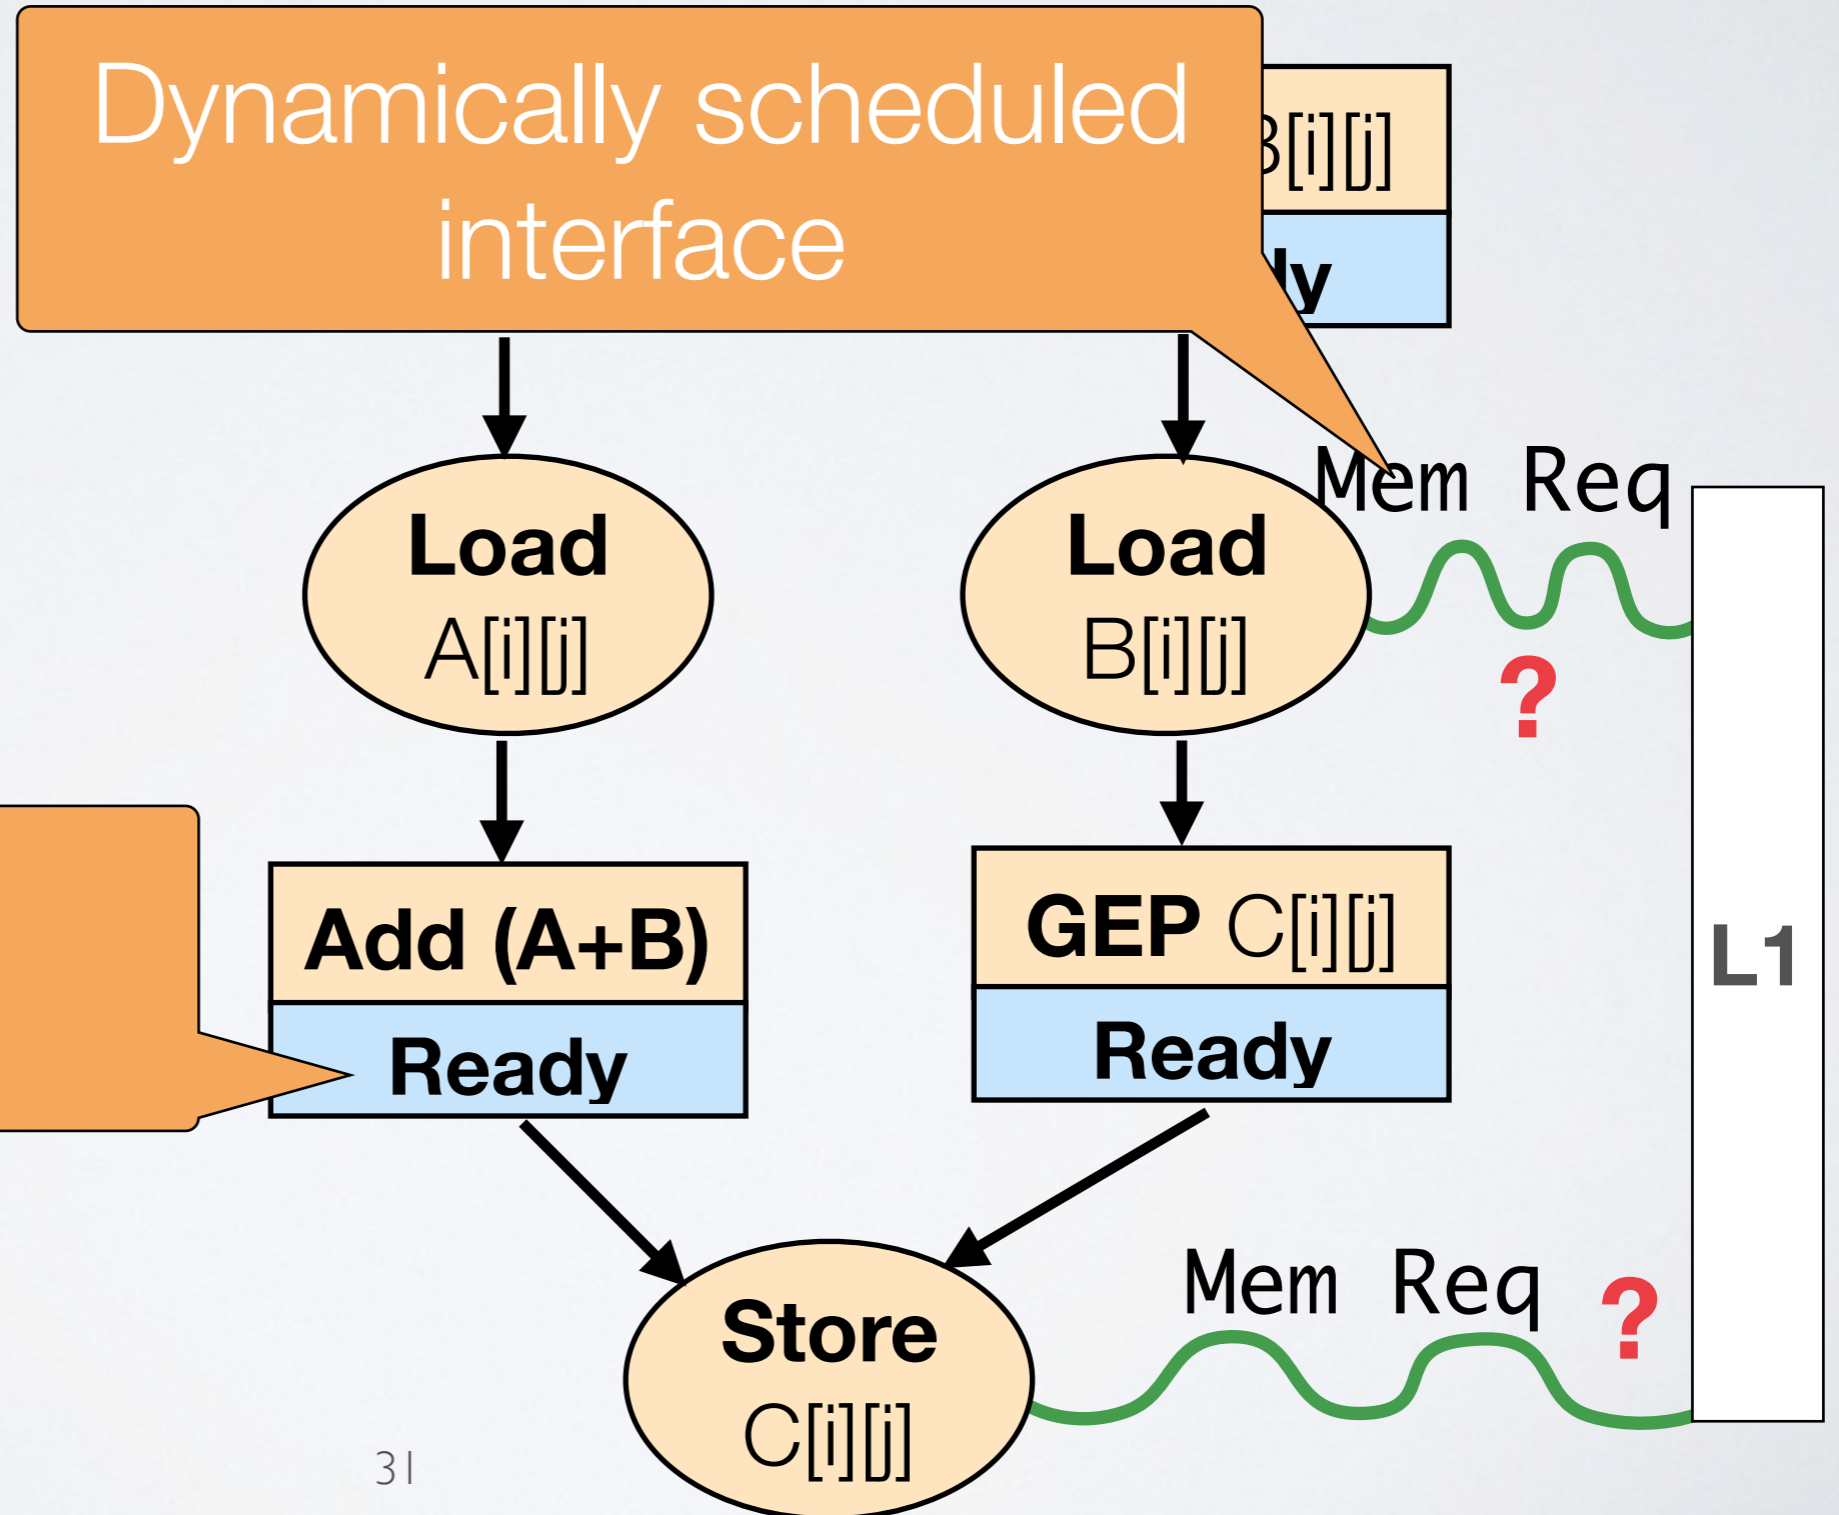
$c[i][j] = a[i][j] + b[i][j];$

Support for all LLVM semantics



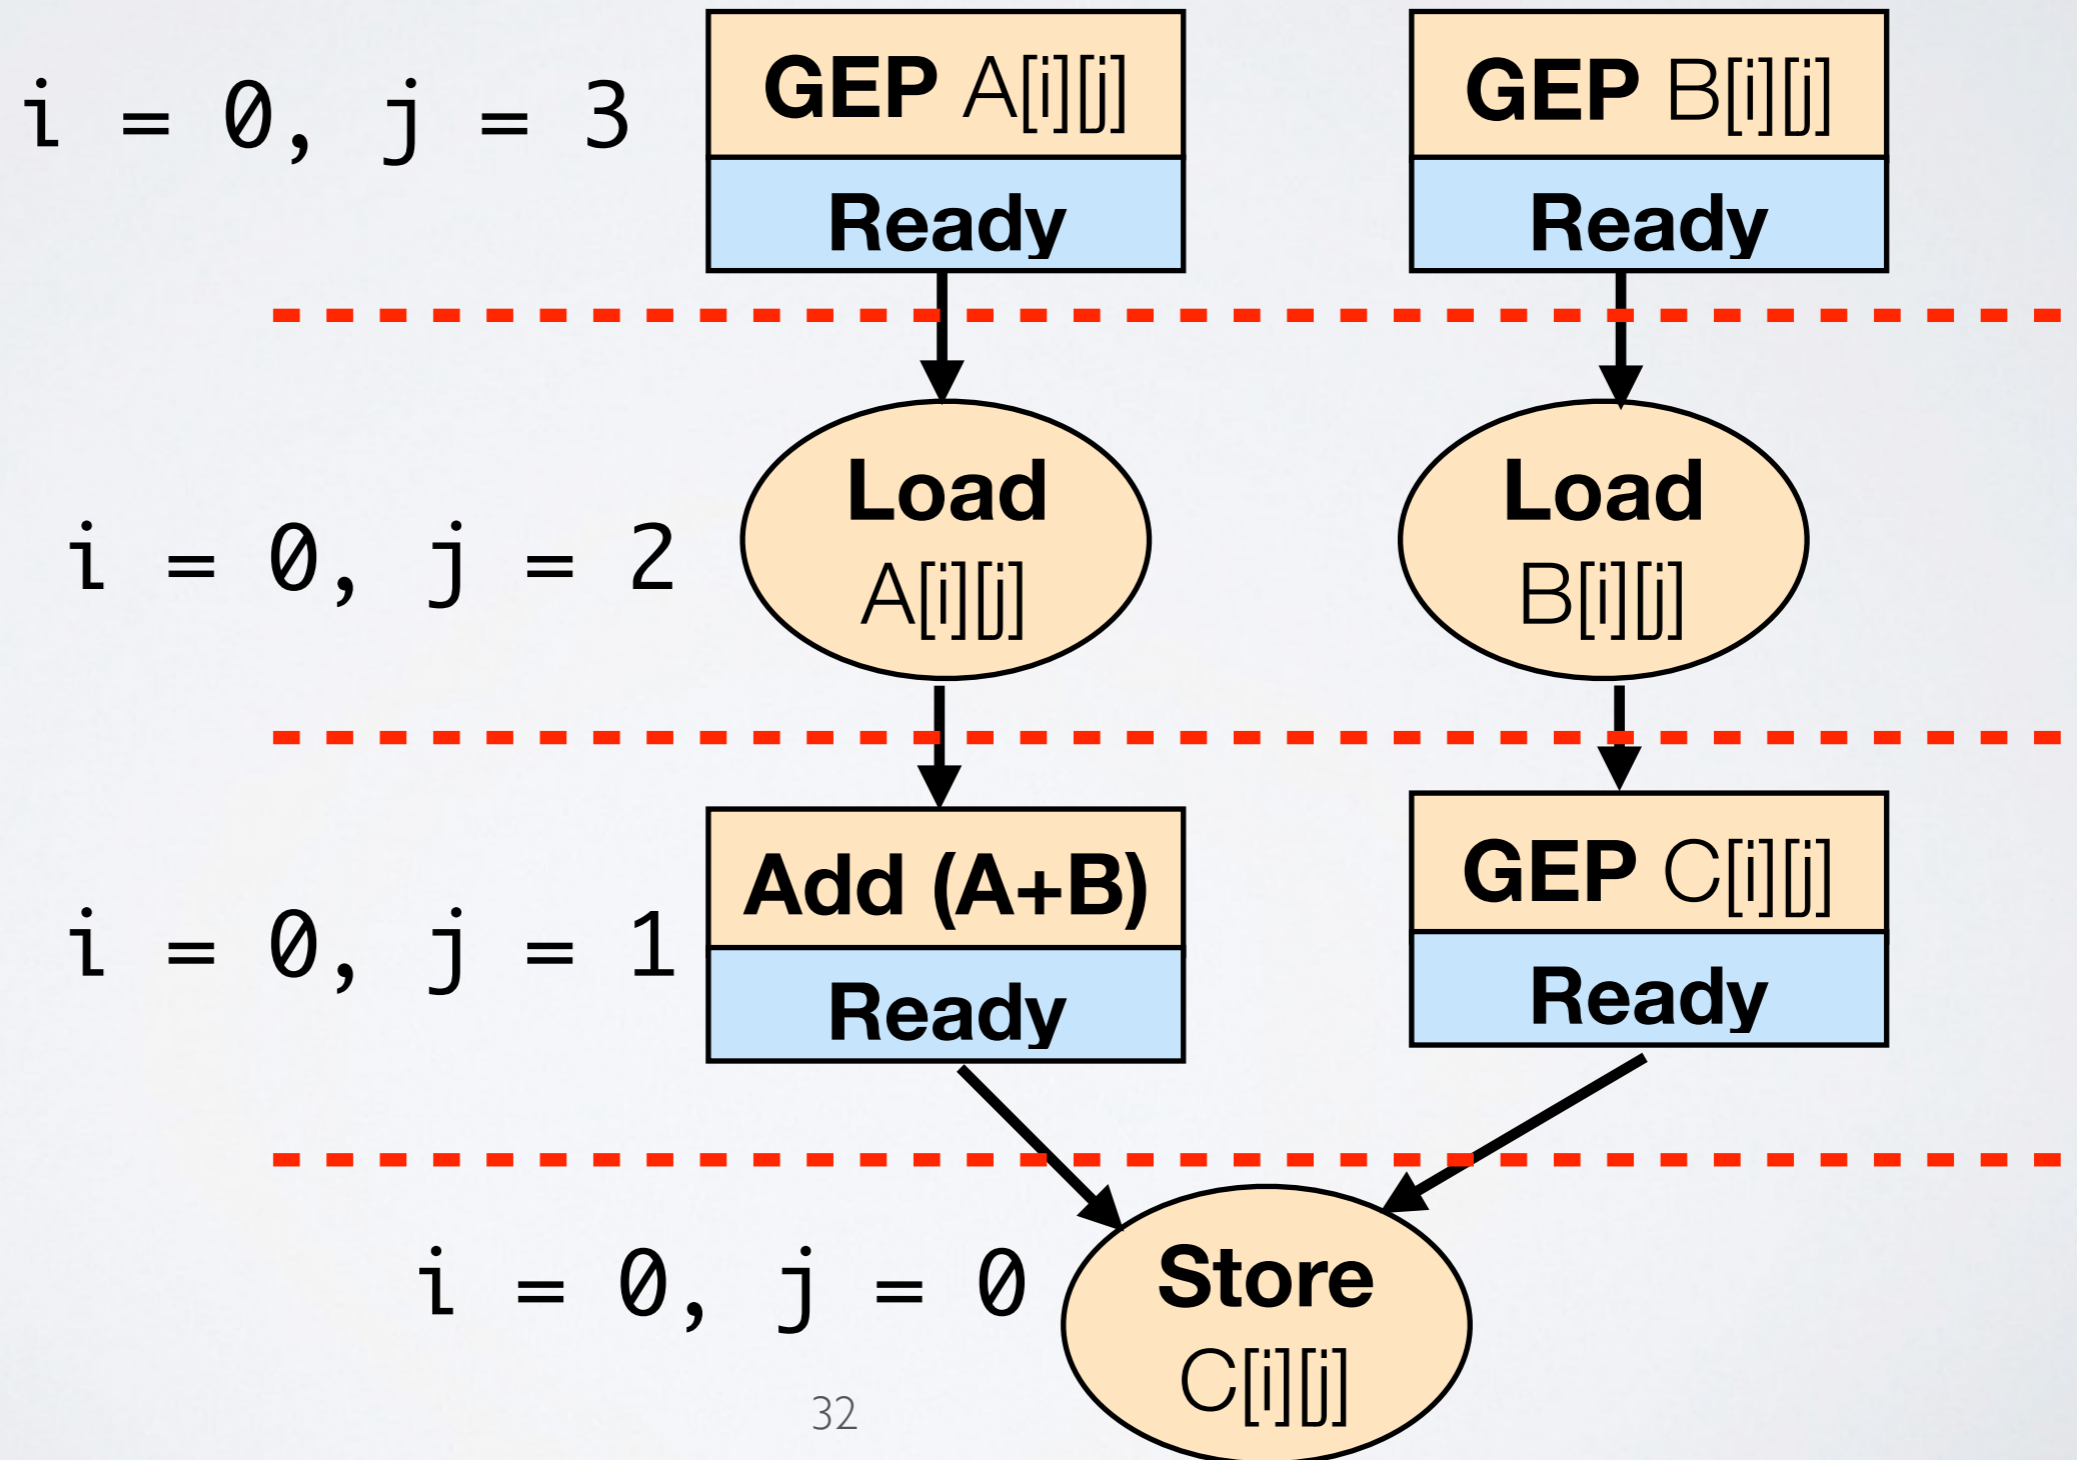
Task Execution Unit (TXU)

$$c[i][j] = a[i][j] + b[i][j];$$



Task Execution Unit (TXU)

$c[i][j] = a[i][j] + b[i][j];$




Experiment

- Board:
 - Arria 10 SOC
 - Intel core i7
- Execution time reported
 - Number of Cycles
- Goal:
 - Performance/watt improvement
 - Reducing overhead of spawning tasks with few instructions




How does performance scale with workload size?

- Unlike a CPU, FPGA performance scales with **#TXU** even for ***fine grained*** parallelism.



Does performance scale with recursion?


- Performance scales with recursive algorithms



How does performance compare to CPU?

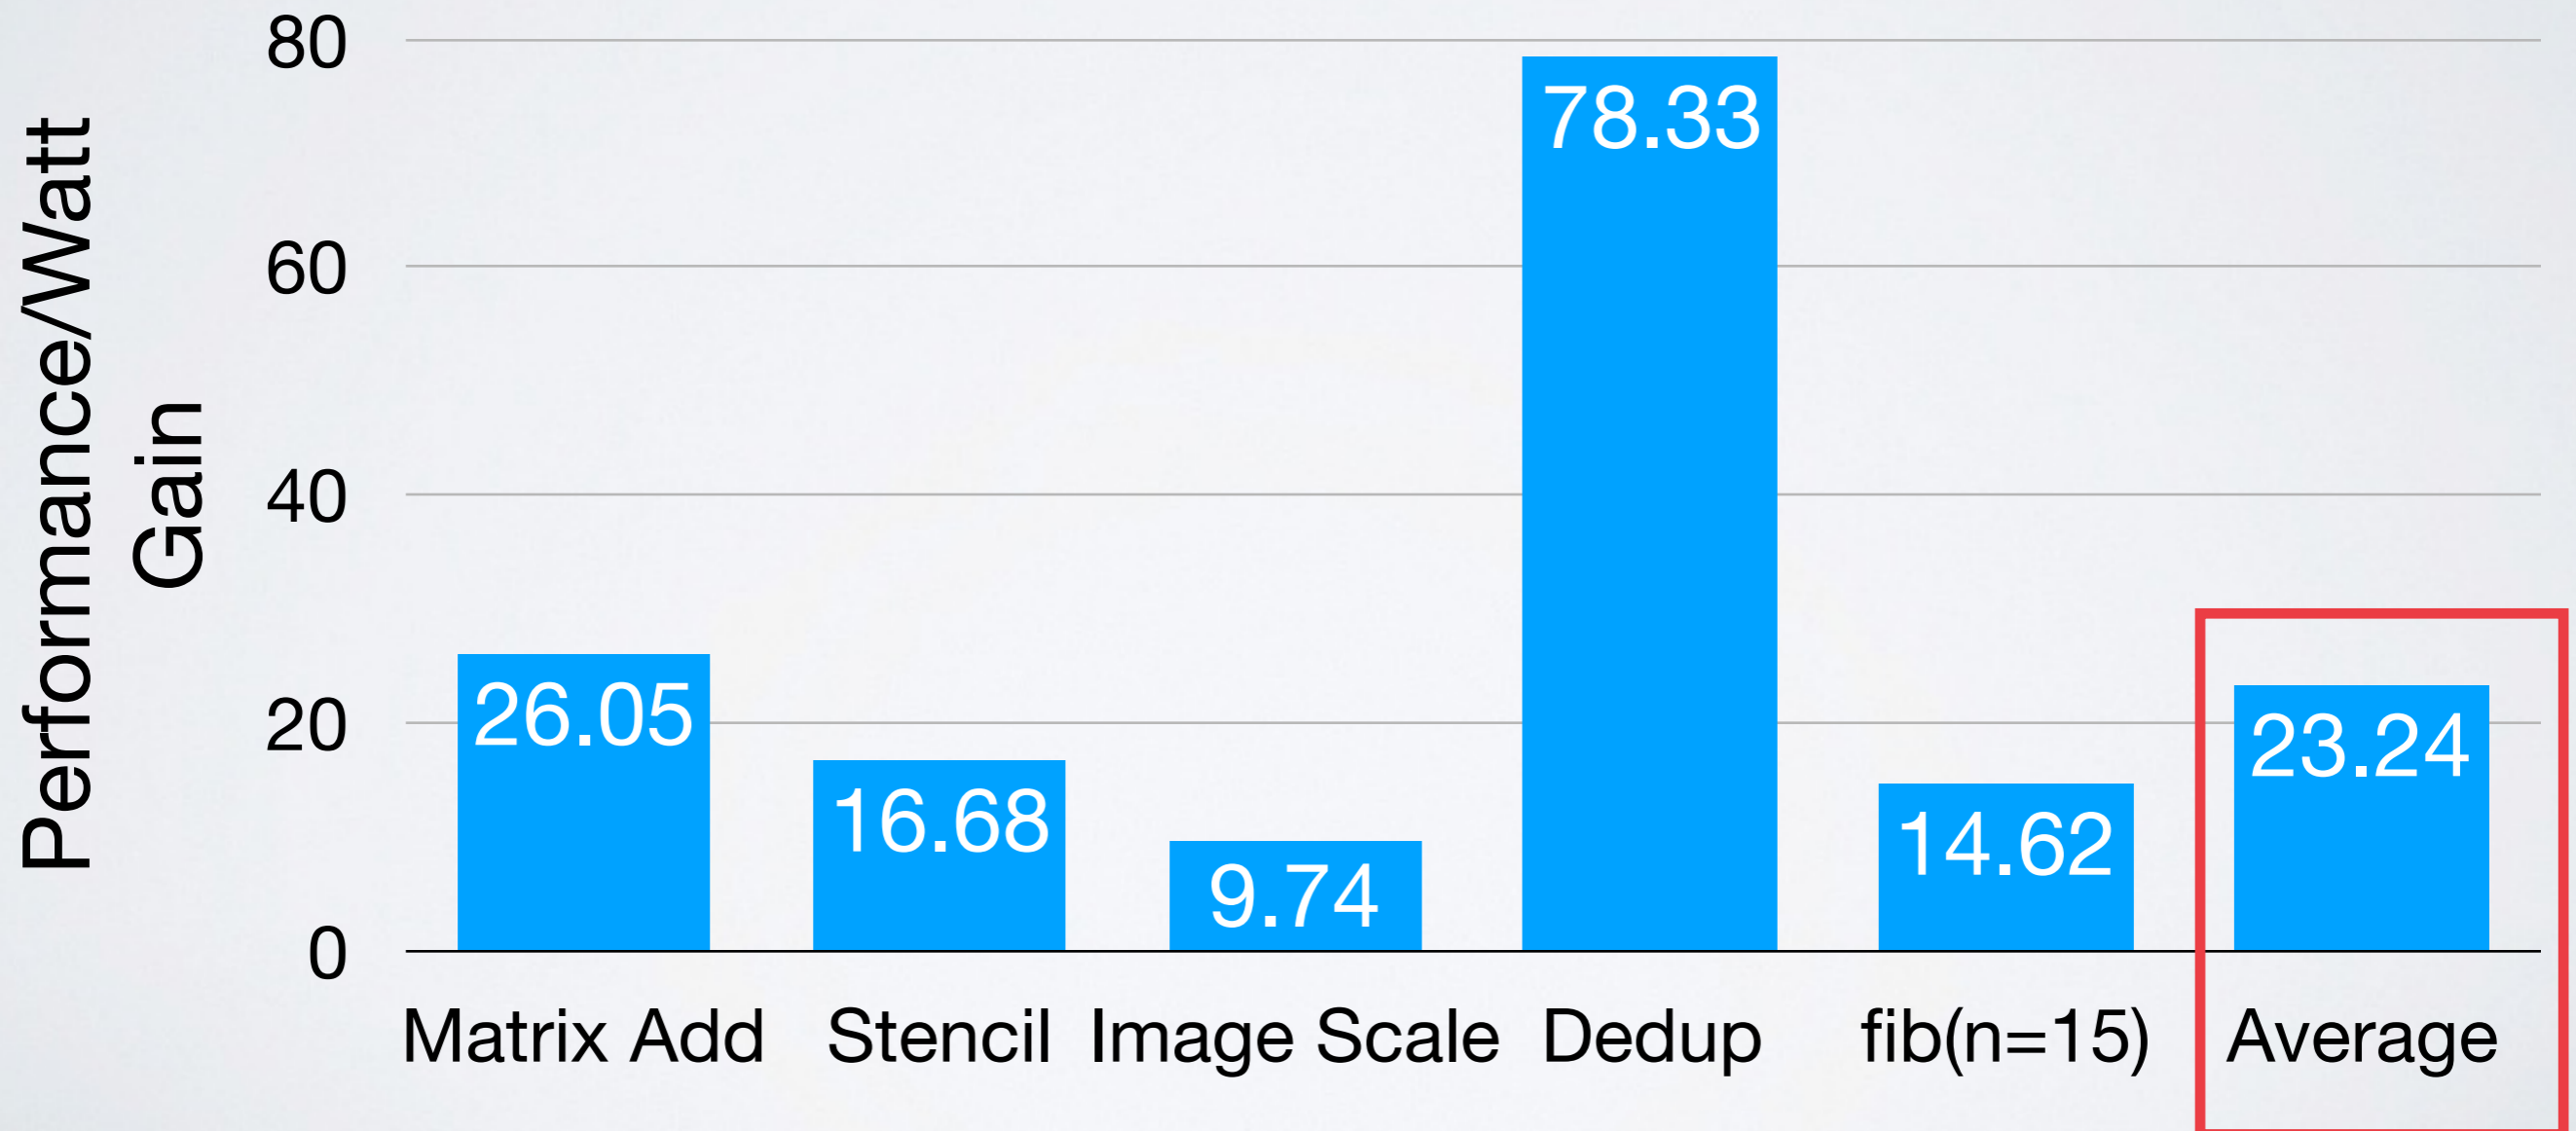
- Performance gain compare to a Intel core i7

(> 1) = FPGA Faster



How does Performance/Watt compare to CPU?


- Performance/Watt has significant improvement




What is the overhead of task controller?

- ALM Utilization by Sub-block

● Tiles ● Parallel for ● Task Ctrl ● Mem Arb ● Misc



1 Tile / 1 Ins



10 Tile / 1 Ins

Available now



<https://github.com/sfu-arch/tapas>

Thanks Chisel and Tapir folks

Shout out to related...

- An Architectural Framework for Accelerating Dynamic Parallel Algorithms on Reconfigurable Hardware (**MICRO51**)
- Dynamically scheduled high-level synthesis (**FPGA18`**)

Parametrization and Configuration

- TAPAS generated accelerator is Parametrizable and Configurable.
 - The number of TXUs can be set specifically for each task base on different criteria.
 - Datapath width can be set at this phase, supporting **mixed precision** as well.
 - Memory modules within each Task Unit are configurable like scratchpads, network and cache.