

Scalable Performance for Scala Message-Passing Concurrency

Andrew Bate

Department of Computer Science
University of Oxford

Motivation


Multi-core commodity hardware

Non-uniform shared memory

Expose potential parallelism

Correctness and formal verification

Compatibility



- 1 Embedded DSL
- 2 Bytecode rewriting
- 3 Channels
- 4 Scheduler
- 5 Deadlock detection

EMBEDDED DOMAIN-SPECIFIC LANGUAGE

Why an Embedded DSL?

Ease of implementation

Leverage existing tools

Leverage known syntax

Higher-order functions

Rich type system

Lightweight syntax

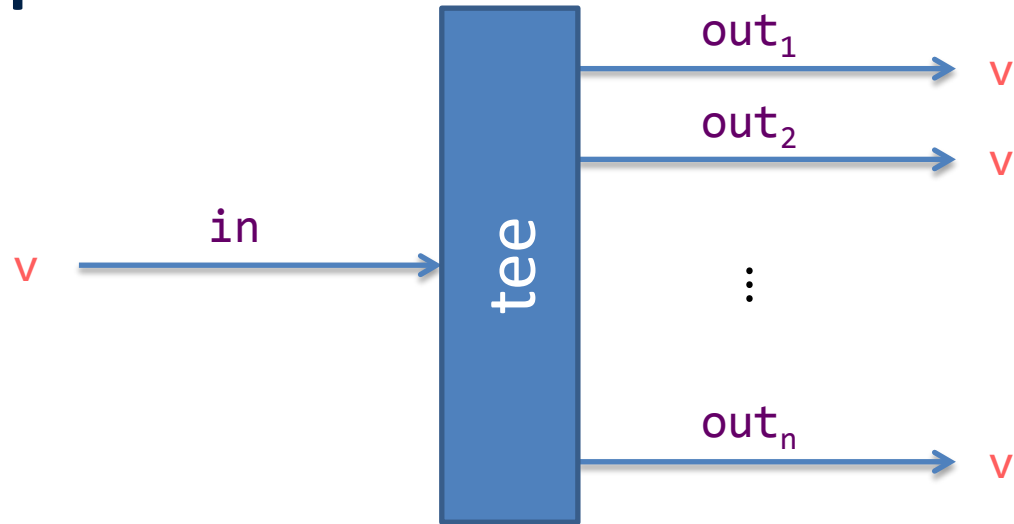
Compile-time macros

Examples



```
def map[I, O](f: I => O)(in: ?[I], out: ![O]) =  
  proc {  
    repeat { out ! (f(in?)) }  
    run (proc { in.closein } || proc { out.closeout })  
  }
```

Examples



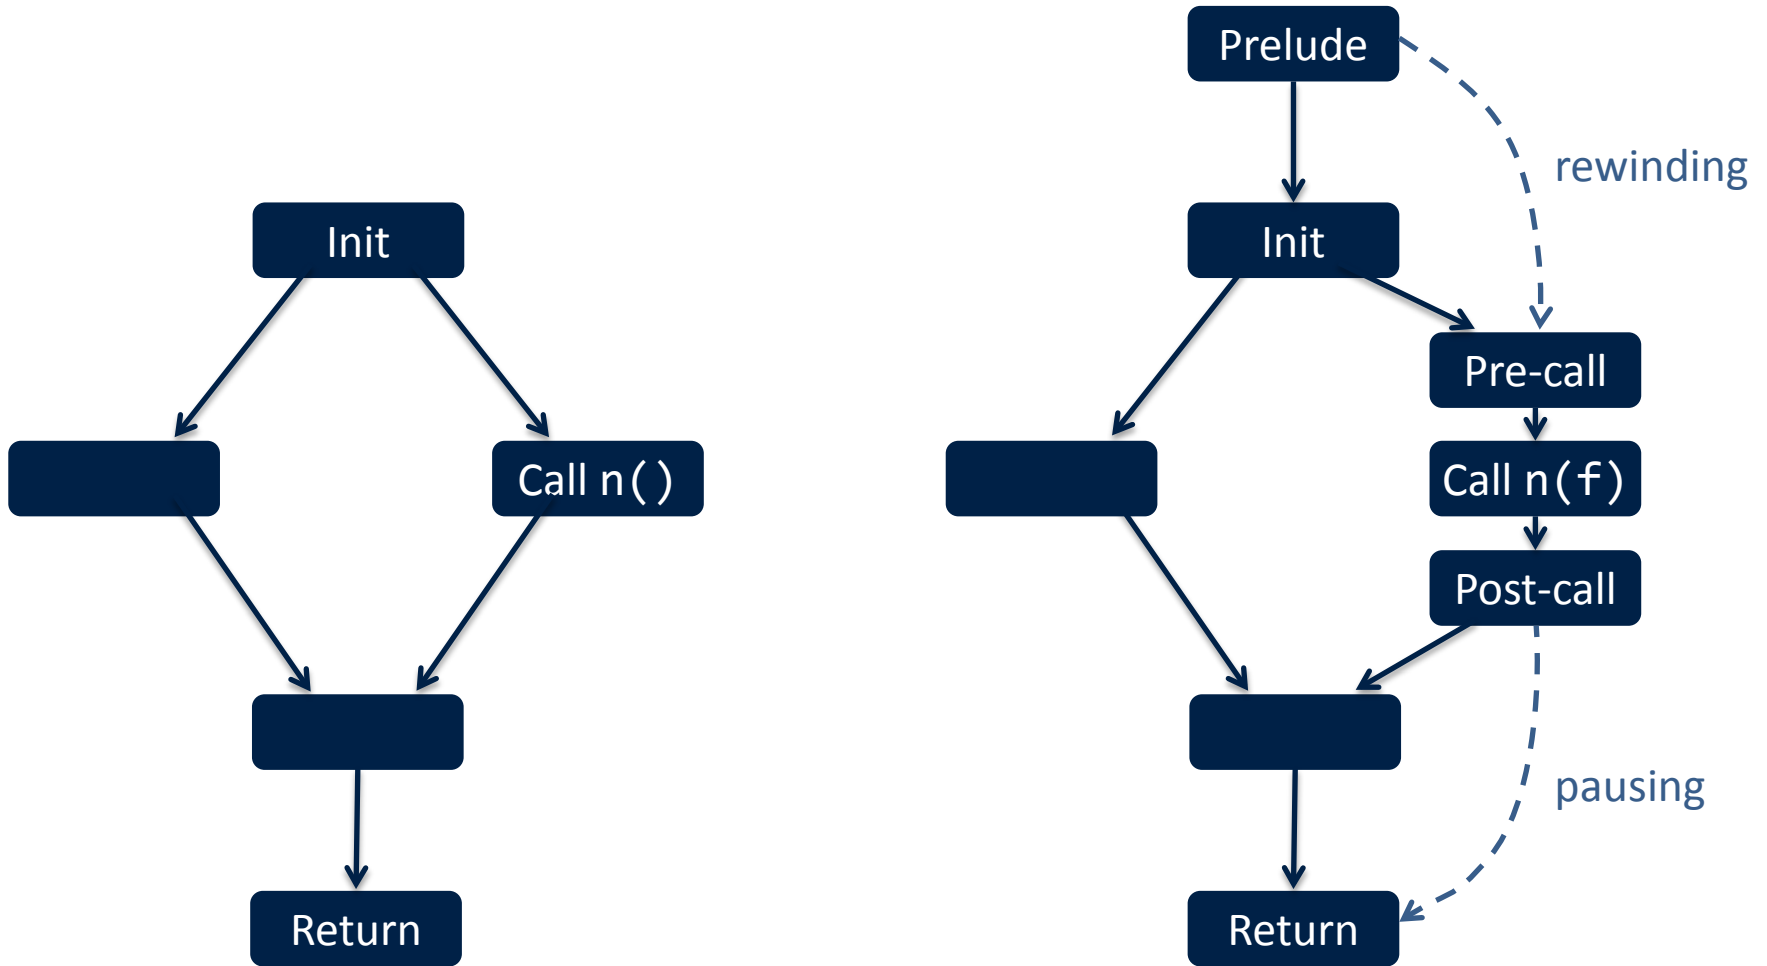
```
def tee[@specialized I](in: ?[I], outs: Seq![I]) =  
  proc {  
    var v = null  
    val outputs = (|| (out <- outs) proc { out ! v })  
    repeat { v = in?; run outputs }  
    run (proc { in.closein } || (|| (out <- outs) proc { out.closeout } ))  
  }
```

- 1 Embedded DSL
- 2 Bytecode rewriting
- 3 Channels
- 4 Scheduler
- 5 Deadlock detection


SCALABLE PERFORMANCE

through bytecode rewriting

CPS Transformation



Analysing the call graph



Engineering

Live variable analysis

Lazy load and store

Constant inlining

Functional Expressions

```
for (i <- 0 until n; j <- i until n) println(i)
```



```
intWrapper(0).until(n).foreach(  
  i: Int => intWrapper(i).until(n).foreach(j: Int => println(i))  
)
```



```
var i = 0  
while (i < n) {  
  var j = i  
  while (j < n) { println(i); j += 1 }  
  i += 1  
}
```

More Features

Tail call optimisations

Shared memory

SBT plugin support

- 1 Embedded DSL
- 2 Bytecode rewriting
- 3 Channels**
- 4 Scheduler
- 5 Deadlock detection

CHANNELS

More Features

Generalised alt

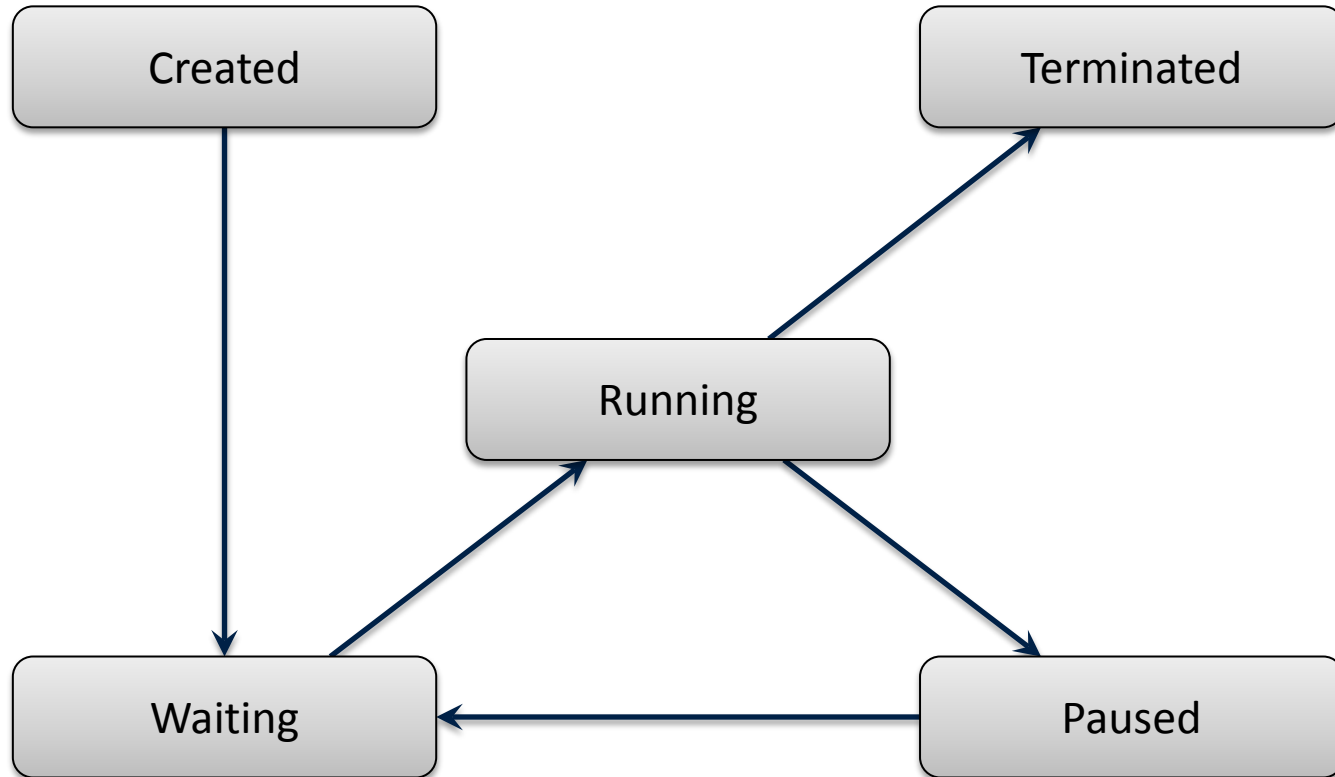
Specialization for primitives

Optimised extended rendezvous


- 1 Embedded DSL
- 2 Bytecode rewriting
- 3 Channels
- 4 Scheduler**
- 5 Deadlock detection

SCHEDULER

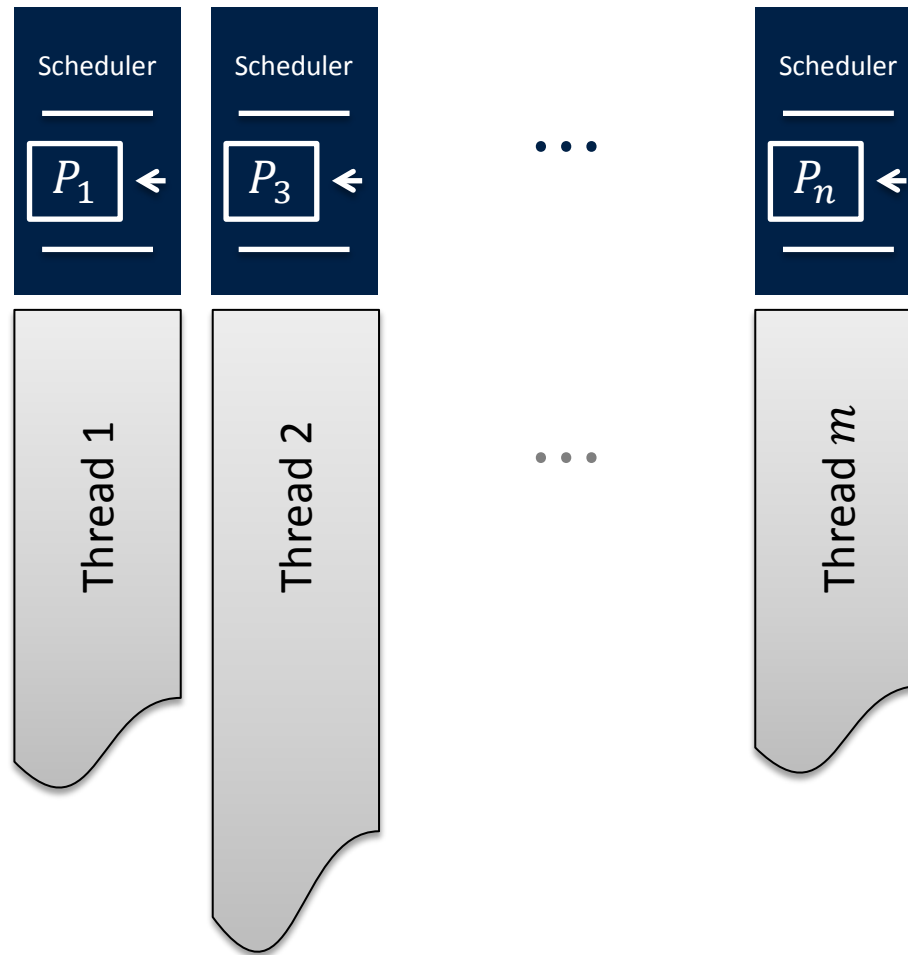
Scheduler States



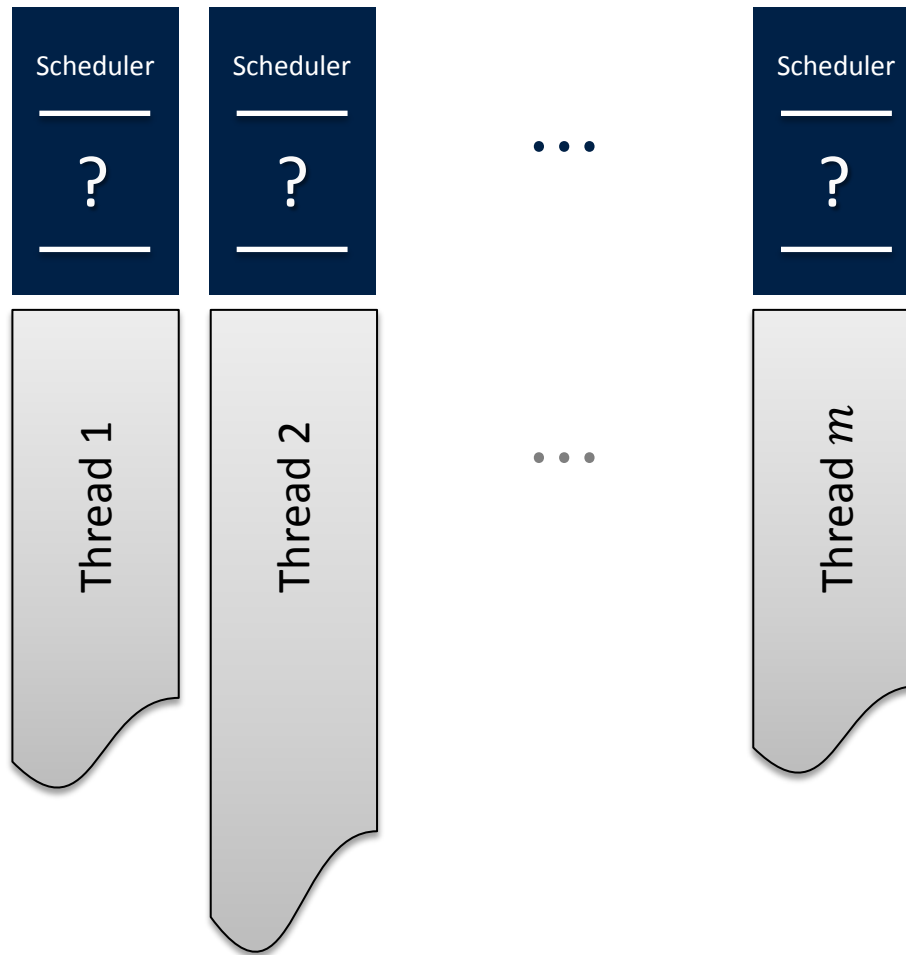
Scheduling: Central FIFO



Scheduling: FIFO per thread




Scheduling: Batches per thread



Scheduling: Batches per thread

Scheduler

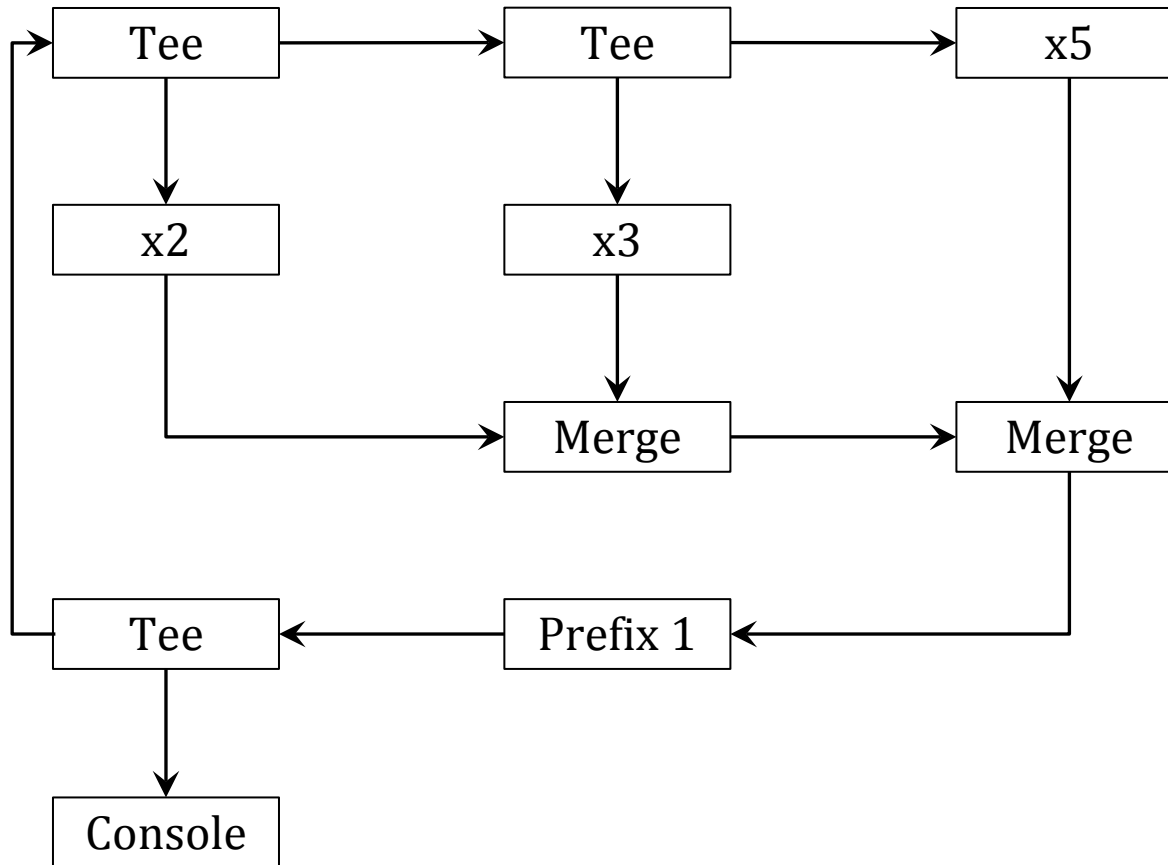
Dispatch Count = $\max(\text{const} \times \text{Batch Length}, \text{Dispatch Limit})$



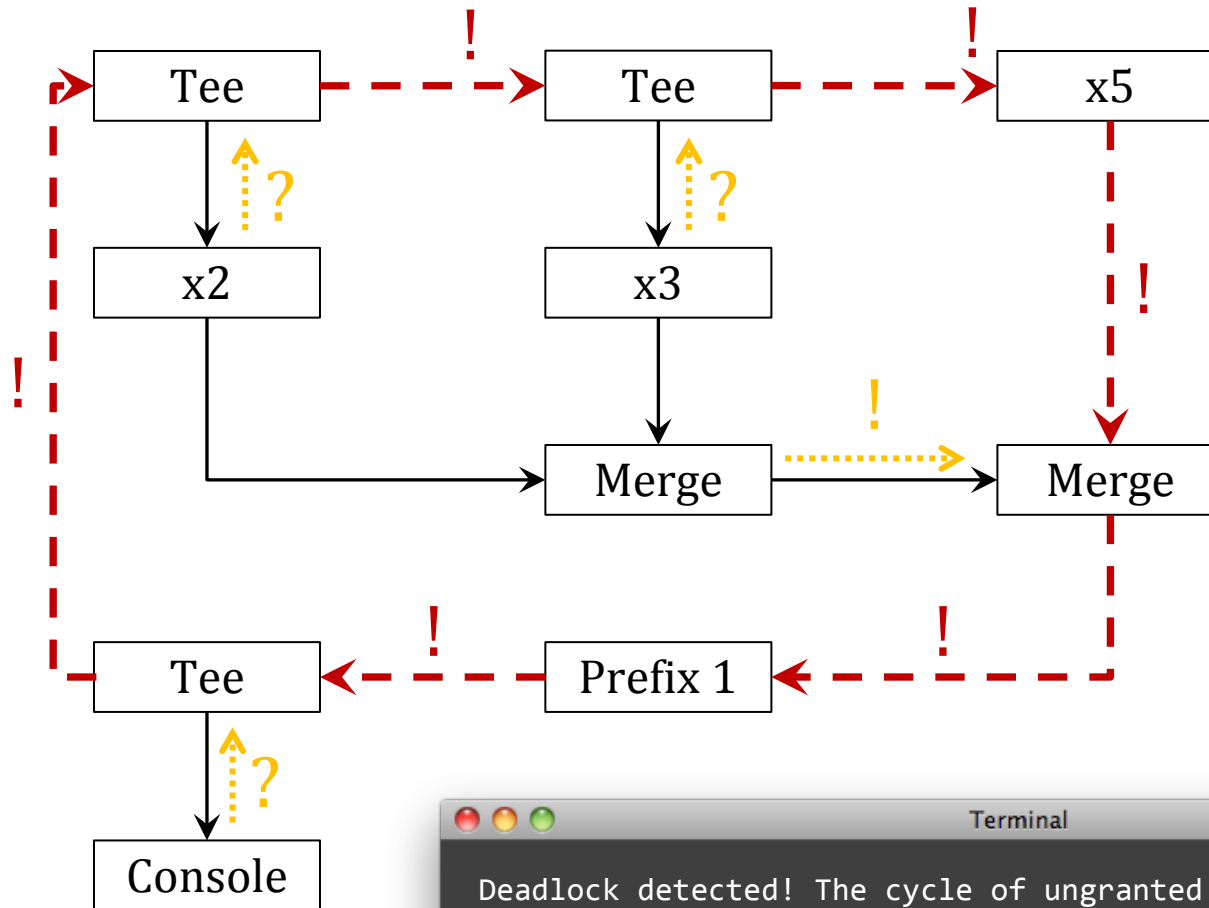
- 1 Embedded DSL
- 2 Bytecode rewriting
- 3 Channels
- 4 Scheduler
- 5 Deadlock detection

DEADLOCK DETECTION

Example



Example





```
Terminal
Deadlock detected! The cycle of ungranted requests is:
Prefix1 -!-> Tee1      Tee3 -!-> x5
Tee1 -!-> Tee2        x5 -!-> Merge2
Tee2 -!-> Tee3        Merge2 -!-> Prefix1
```



- 1 Embedded DSL
- 2 Bytecode rewriting
- 3 Channels
- 4 Scheduler
- 5 Deadlock detection

PERFORMANCE EVALUATION

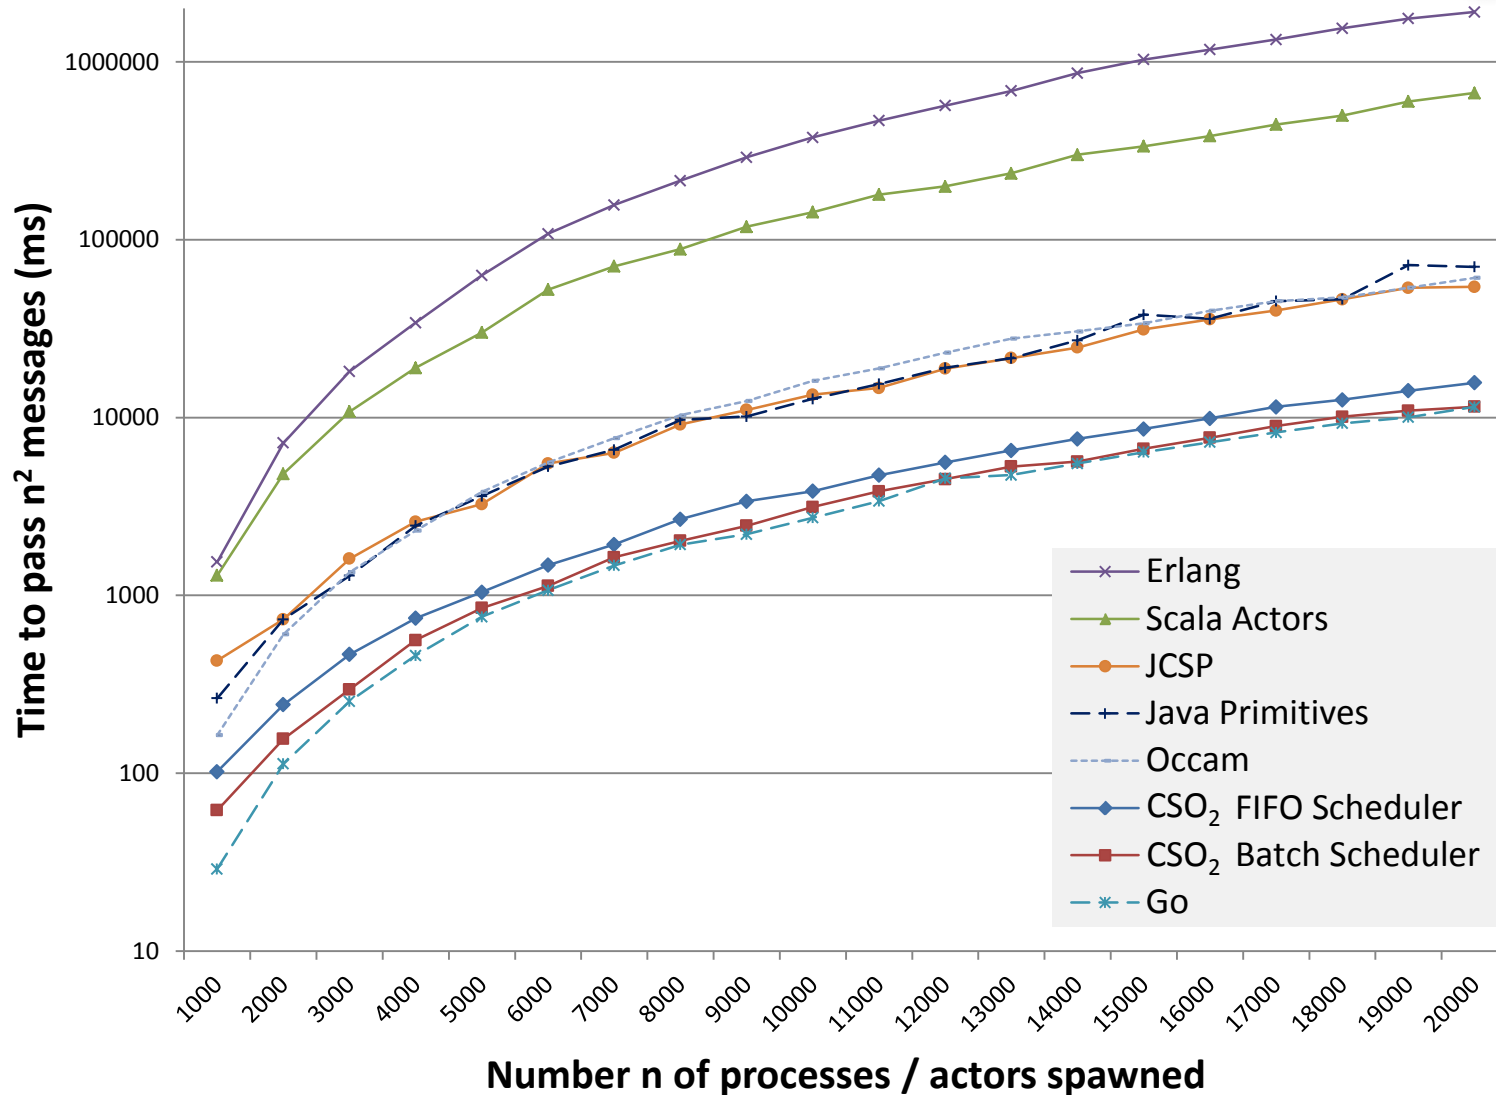

Ring topology



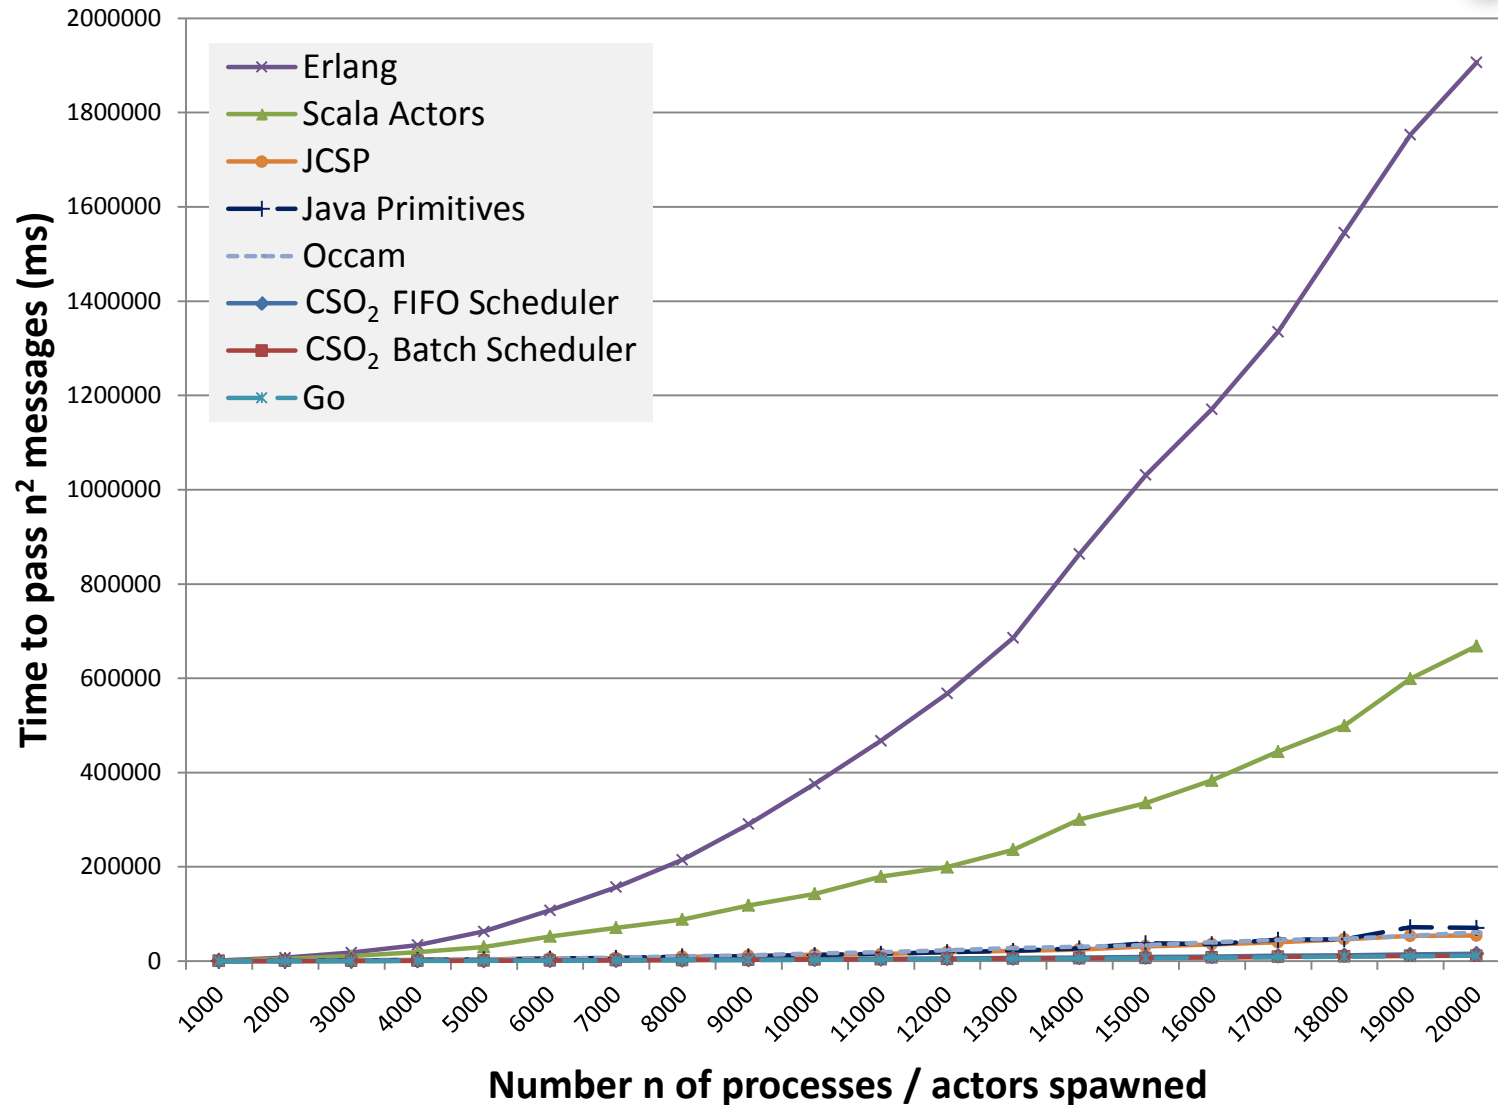

Ring topology



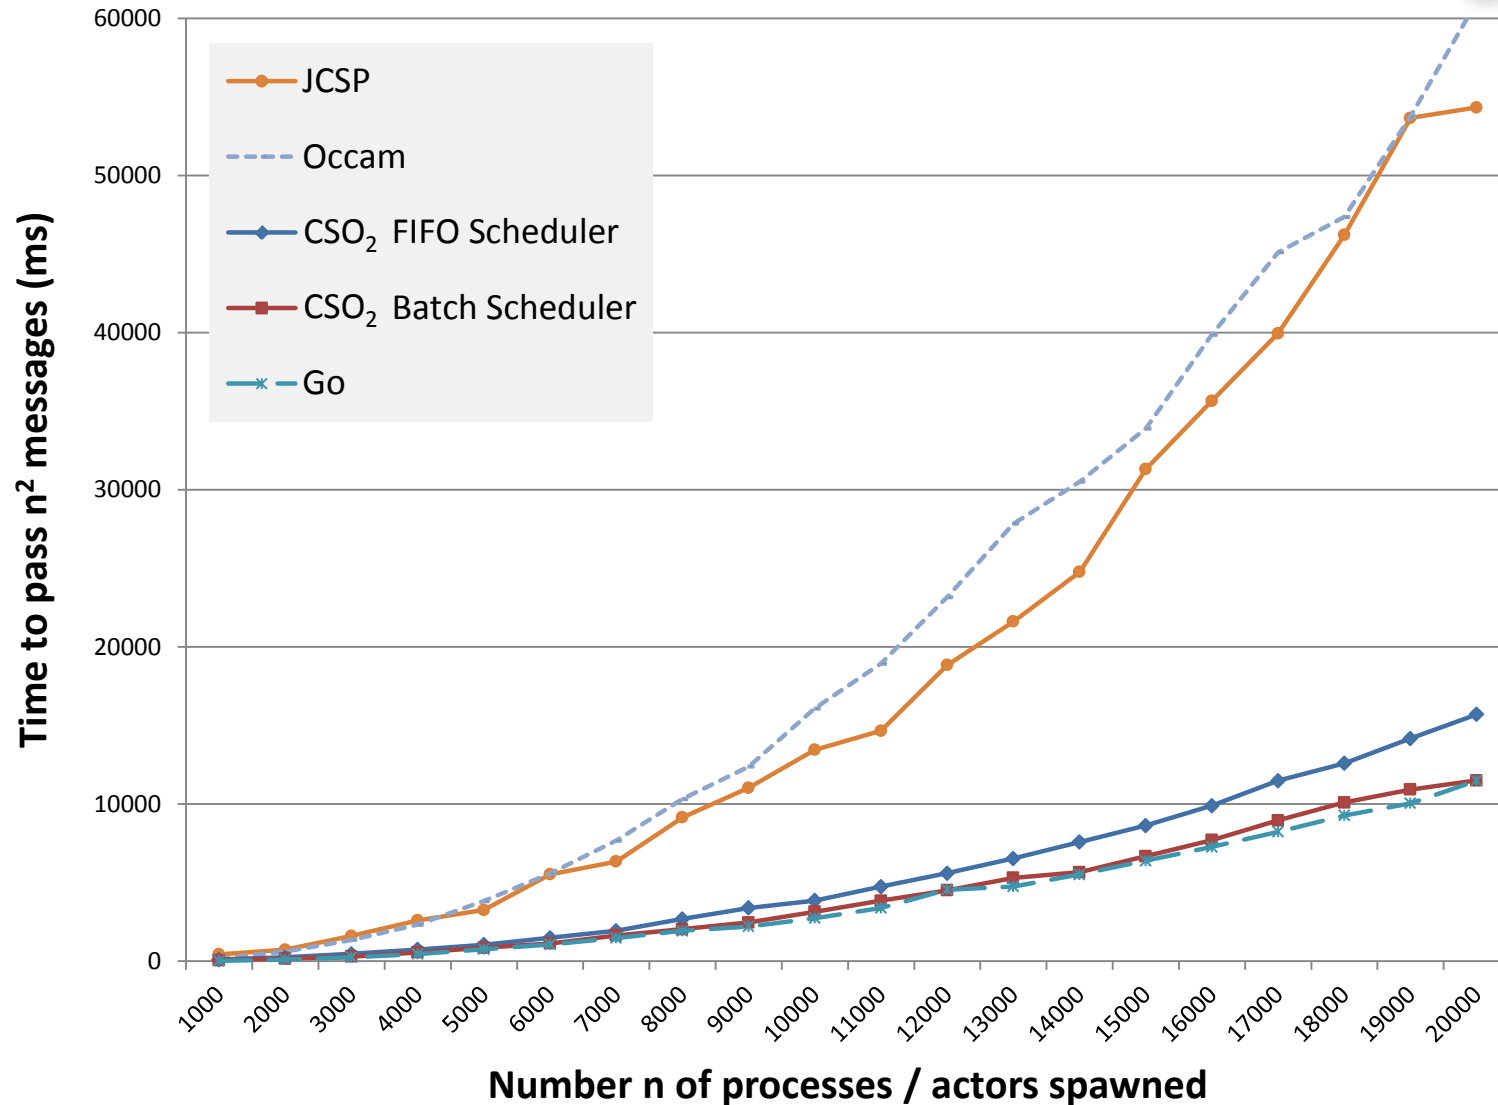

Fully connected topology





Fully connected topology



Fully connected topology



Fully connected topology



Summary

- High performance library for building massively concurrent systems on the JVM
- Deadlock detection
- Outperforms Java primitives, JCSP, Scala Actors, Occam, and very close to Go