

DFScala: High Level Dataflow Support for Scala

Daniel Goodman, Salman Khan, Chris Seaton,
Yegor Gusgov, Behram Khan,
Mikel Luján and Ian Watson

Workshop on Data-Flow Execution Models for Extreme Scala Computing
23 September 2012
Minneapolis

TERAFLUX

EPSRC

Engineering and Physical Sciences
Research Council

The Teraflux project is funded by the European Commission Seventh Framework Programme under grant agreement number 249013. Chris Seaton is an EPSRC funded student. Mikel Luján is a Royal Society University Research Fellow.

Context

- Processors with thousands cores
- Simplifying the architecture by minimising cache coherency
- Teraflux is looking at hardware architectures combining dataflow and transactional memory
- This work is looking at implementing the programming model for non-specialist, industrial programmers

Why Dataflow?

- Simple programming model
- Implicit parallelism
- Deterministic by default
- No shared-state
- No need for cache coherency

Basics of Scala

- Martin Odersky, EPFL, 2003
- JVM language – similar look and feel, semantics, object model, concurrency model to Java
- Adds anonymous functions, pattern matching, type inference
- Data structures are immutable by default
- Several parallelism models – threads, actors, parallel collections

Why Scala?

- Pragmatic, high-productivity
- Strong static typing, but dynamic feel through type inference
- Existing Java infrastructure and libraries
- Motivates and facilitates a functional approach, but doesn't enforce it
- A familiar language for many programmers

DFScala Model

- Data-driven
- Coarse-grained
- Nodes are functions
- Graph is dynamic
- Graph is statically checked for type correctness
- Graphs can be nested
- Avoid pitfalls of parallelism, make dataflow simpler

DFScala Library

- Dataflow nodes are Scala functions
- Type-safety using inference
- Nodes are runnable when the function's arguments are available
- Java thread pool
- Shared-state scheduler using software transactional memory

```
def fib(n : Int) : Int =  
  if (n <= 2)  
    1  
  else  
    fib(n-1) + fib(n-2)
```

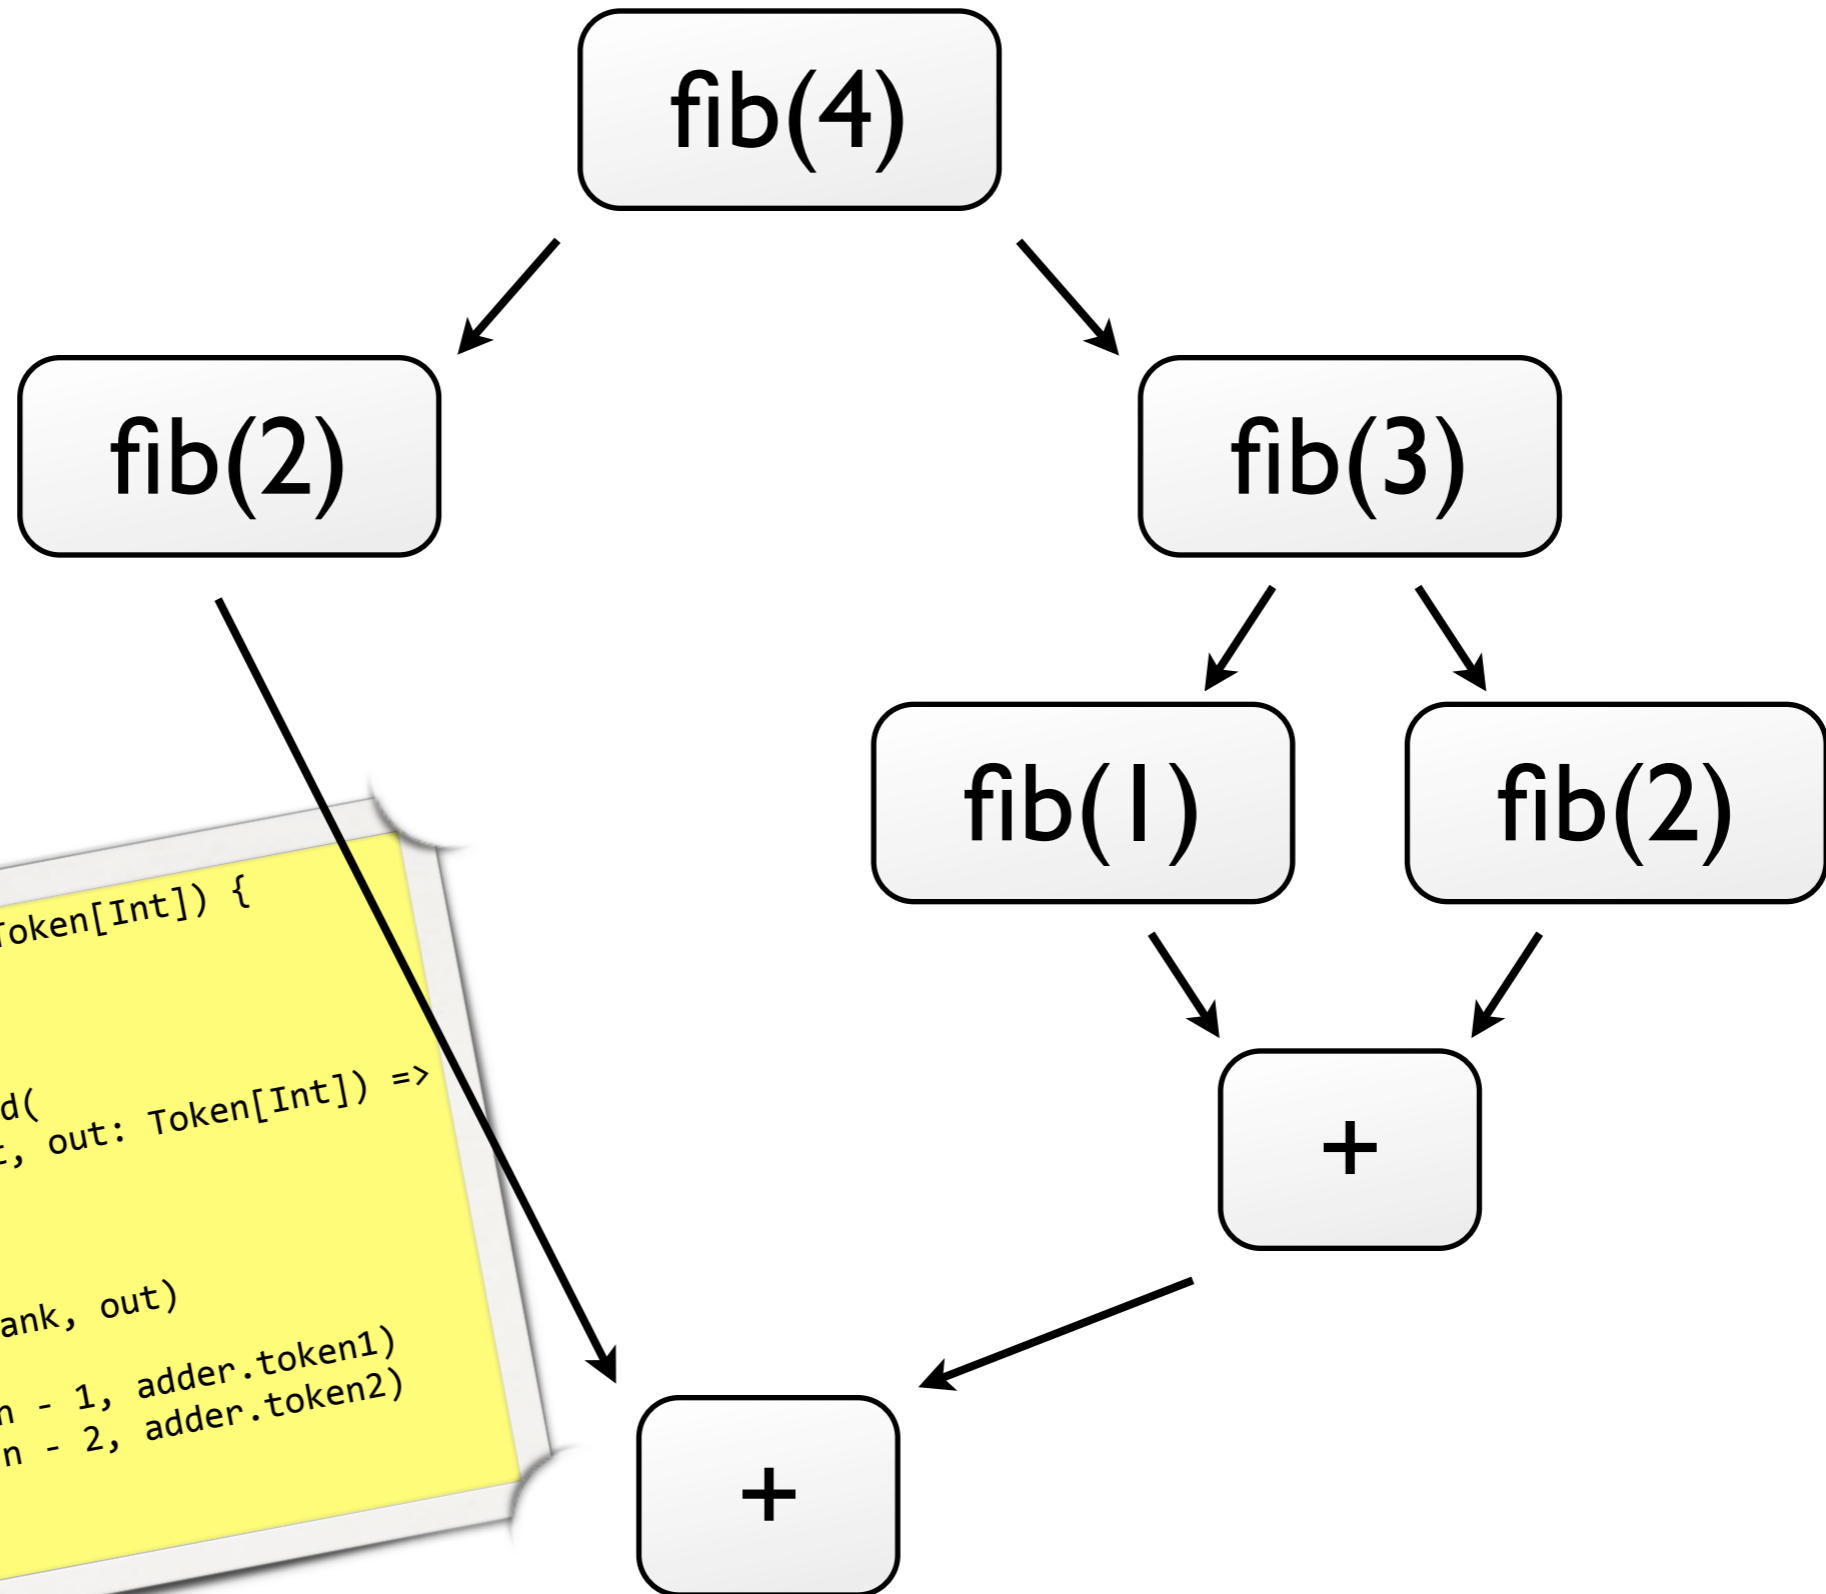


```
def fib(n : Int, out: Token[Int]) {  
  if (n <= 2)  
    out(1)  
  else {  
    val adder = thread(  
      (a: Int, b: Int, out: Token[Int]) =>  
        out(a + b)  
    )  
  
    adder(Blank, Blank, out)  
  
    thread(fib _, n - 1, adder.token1)  
    thread(fib _, n - 2, adder.token2)  
  }  
}
```

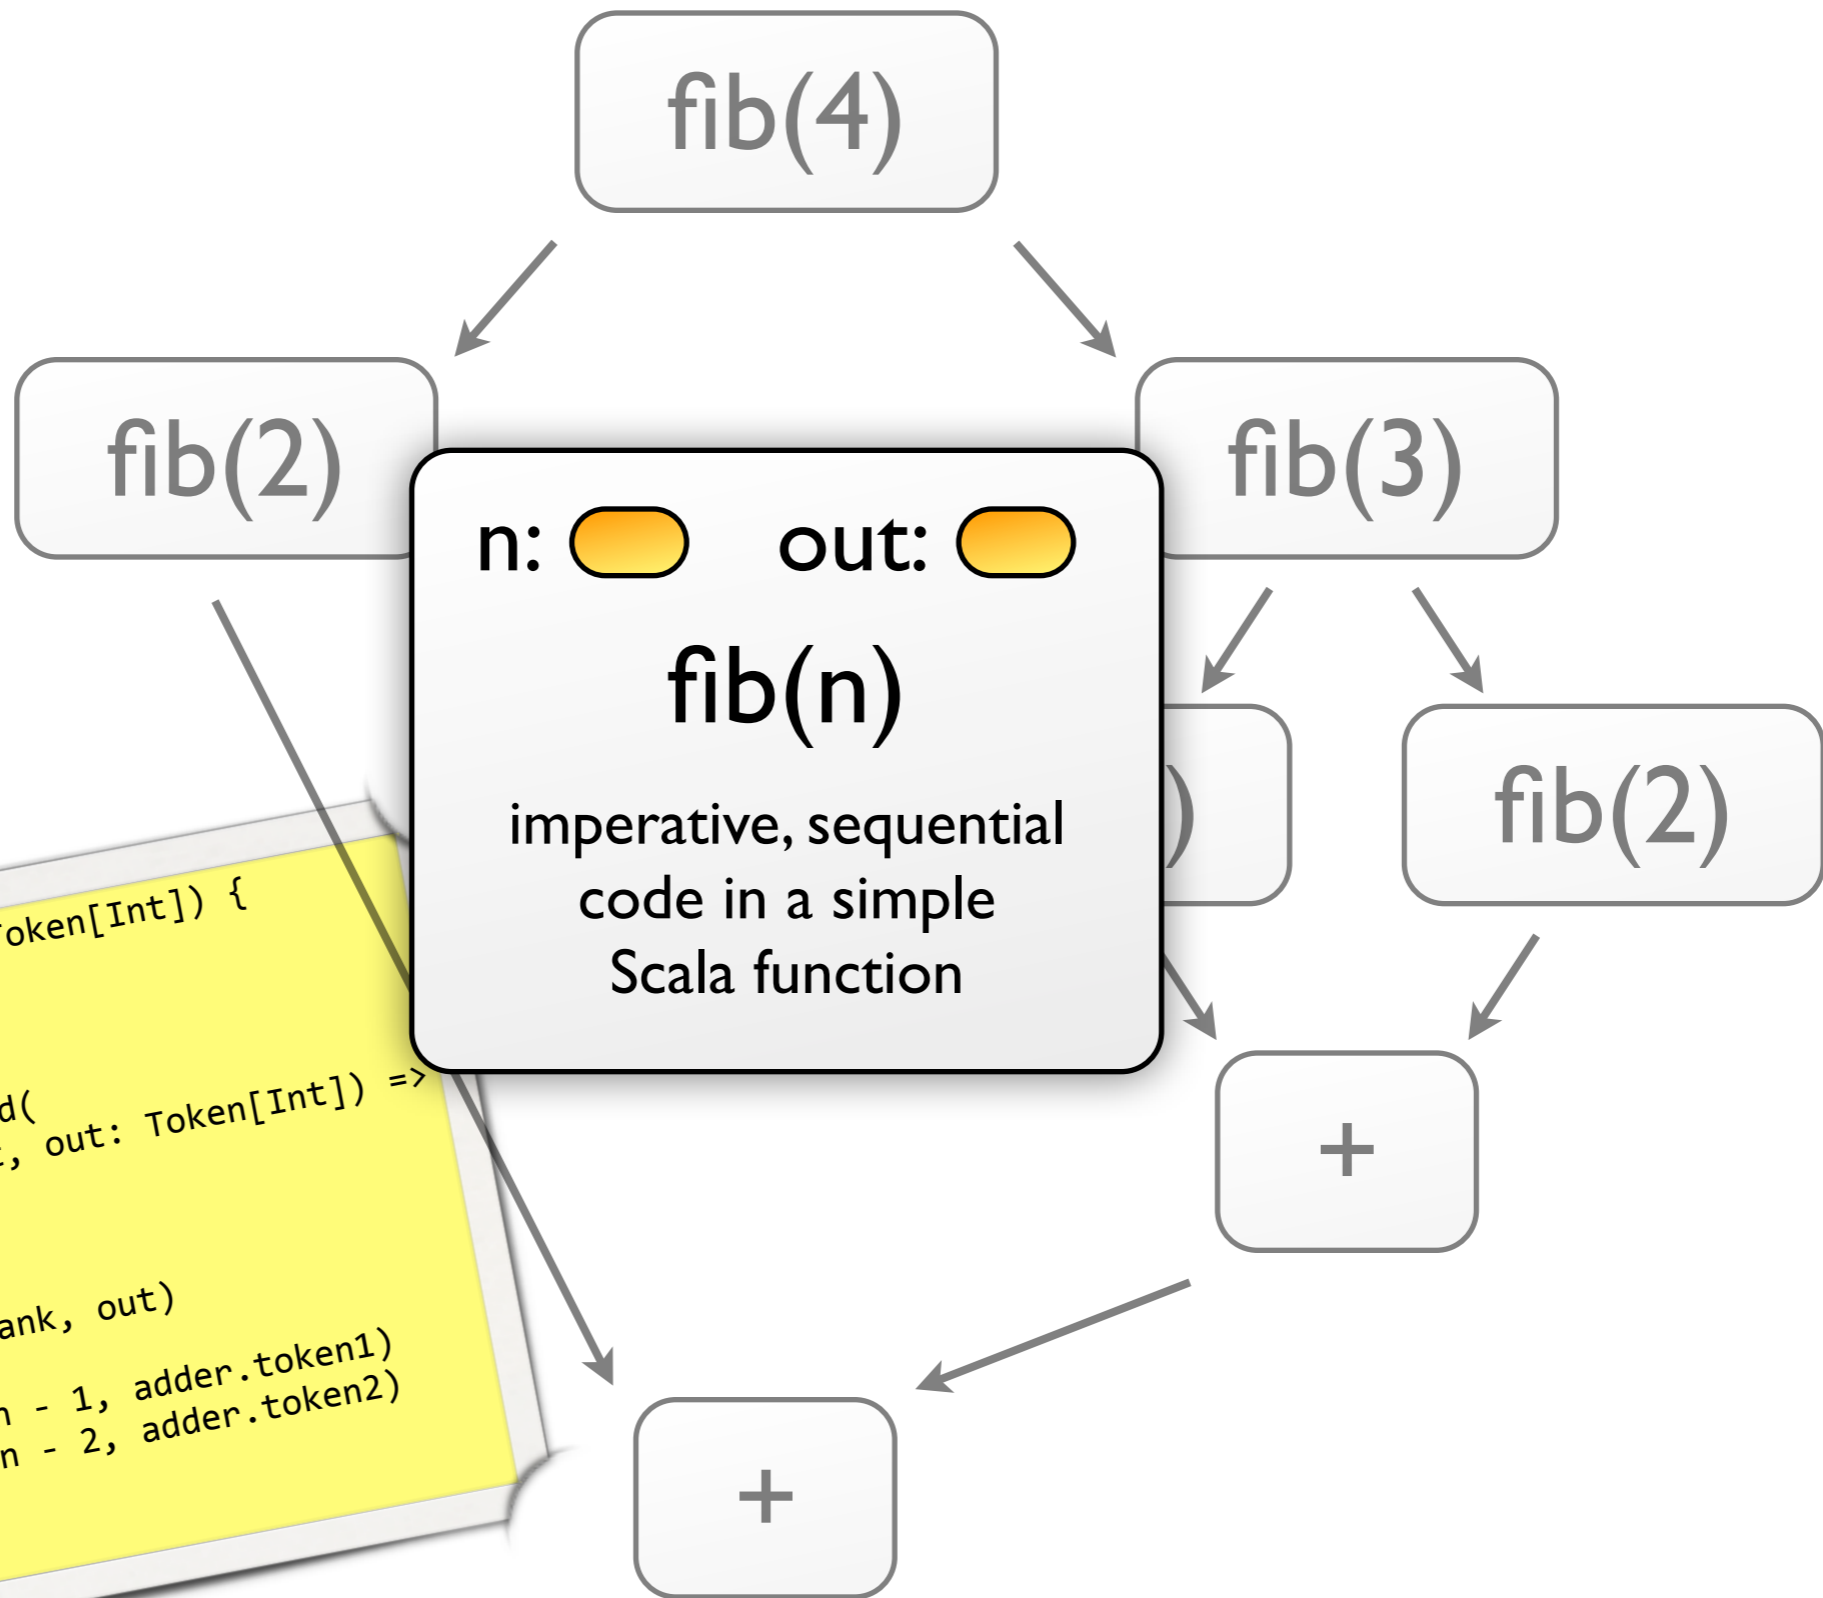
```
def fib(n : Int, out: Token[Int]) {  
  if (n <= 2)  
    out(1)  
  else {  
    val adder = thread(  
      (a: Int, b: Int, out: Token[Int]) =>  
        out(a + b)  
    )  
  
    adder(Blank, Blank, out)  
  
    thread(fib _, n - 1, adder.token1)  
    thread(fib _, n - 2, adder.token2)  
  }  
}
```

```
def fib(n : Int, out: Token[Int]) {  
  if (n <= 2)  
    out(1)  
  else {  
    val adder = thread(  
      (a: Int, b: Int, out: Token[Int]) =>  
        out(a + b)  
    )  
  
    adder(Blank, Blank, out)  
  
    thread(fib _, n - 1, adder.token1)  
    thread(fib _, n - 2, adder.token2)  
  }  
}
```

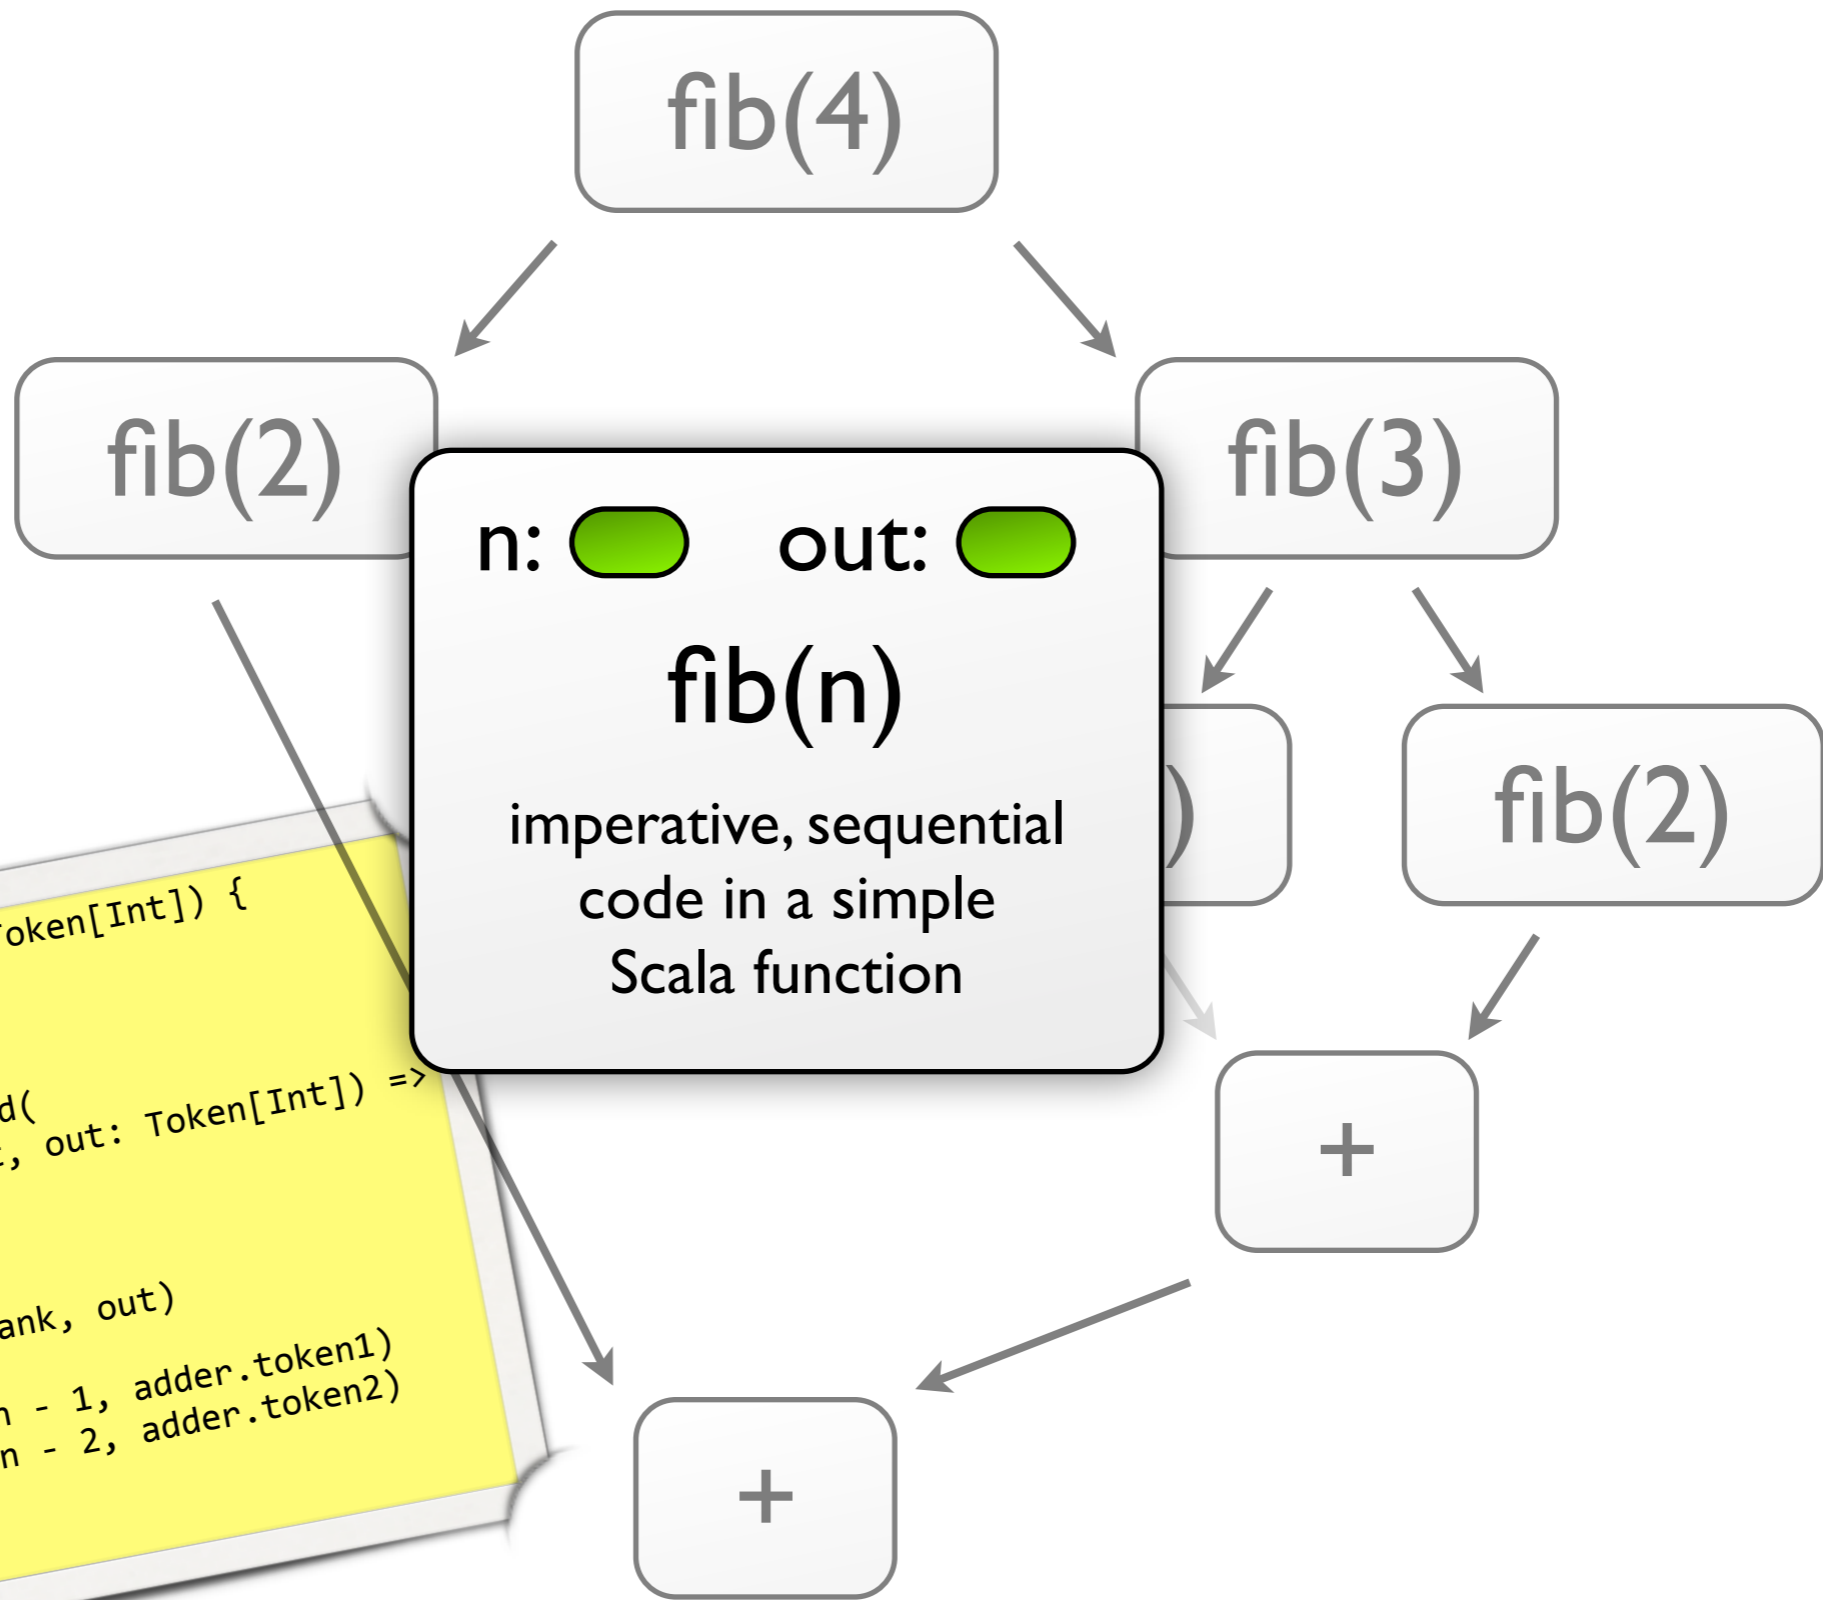
```
def fib(n : Int, out: Token[Int]) {  
  if (n <= 2)  
    out(1)  
  else {  
    val adder = thread(  
      (a: Int, b: Int, out: Token[Int]) =>  
        out(a + b)  
    )  
  
    adder(Blank, Blank, out)  
  
    thread(fib _, n - 1, adder.token1)  
    thread(fib _, n - 2, adder.token2)  
  }  
}
```



```
def fib(n : Int, out: Token[Int]) {  
  if (n <= 2)  
    out(1)  
  else {  
    val adder = thread(  
      (a: Int, b: Int, out: Token[Int]) =>  
        out(a + b)  
    )  
    adder(Blank, Blank, out)  
    thread(fib _, n - 1, adder.token1)  
    thread(fib _, n - 2, adder.token2)  
  }  
}
```



```
def fib(n : Int, out: Token[Int]) {  
  if (n <= 2)  
    out(1)  
  else {  
    val adder = thread(  
      (a: Int, b: Int, out: Token[Int]) =>  
        out(a + b)  
    )  
    adder(Blank, Blank, out)  
    thread(fib _, n - 1, adder.token1)  
    thread(fib _, n - 2, adder.token2)  
  }  
}
```



```

def fib(n : Int, out: Token[Int]) {
  if (n <= 2)
    out(1)
  else {
    val adder = thread(
      (a: Int, b: Int, out: Token[Int]) =>
        out(a + b)
    )
    adder(Blank, Blank, out)
    thread(fib _, n - 1, adder.token1)
    thread(fib _, n - 2, adder.token2)
  }
}
  
```

n: out:

fib(n)

imperative, sequential
code in a simple
Scala function


Benefits

- Programmers only need to worry about functions and the usual function boundaries
- Use existing arguments instead of l-structures
- No mutable state
- Graph is statically checked for type safety
- Can be used with existing code and libraries


Nested Graphs

- DFSScala by default does not allow for any parallelism within nodes
- Each function forming a node is entirely sequential and runs from start to finish without pauses
- As an explicit exception, dataflow graphs can be nested within nodes

Nested Graphs



Nested Graphs



Compare to Akka

- Dataflow library for Scala
- Part of a framework for large, distributed, concurrent applications (think Erlang)
- Released after we started work
- Based on the Oz model
- New concepts – dataflow variables (I-structures)
- Concurrency within functions using continuations

Compare to Akka

```
def fib(n: Int): Int = {  
  val r, a, b = Promise[Int]()  
  
  flow { r << a() + b() }  
  flow { a << fib(n - 1) }  
  flow { b << fib(n - 2) }  
  
  return r()  
}
```

Compare to CNC-Scala

- Port of C++ Intel Concurrent Collections to Scala
- Also have a focus on ‘parallelism oblivious developers’ and productivity
- Parallelism is within collections
- Each collection element is an I-structure, dependencies between them form the dataflow graph
- Also concurrency in functions using continuations

Comparisons

	DFScala	Akka	CNC
Dataflow	✓	✓	✓
Scala	✓	✓	✓
Dynamic	✓	✓	✓
Task parallel	✓	✓	✓
Pure Scala	✓	✓	
Special dataflow variables		✓	✓

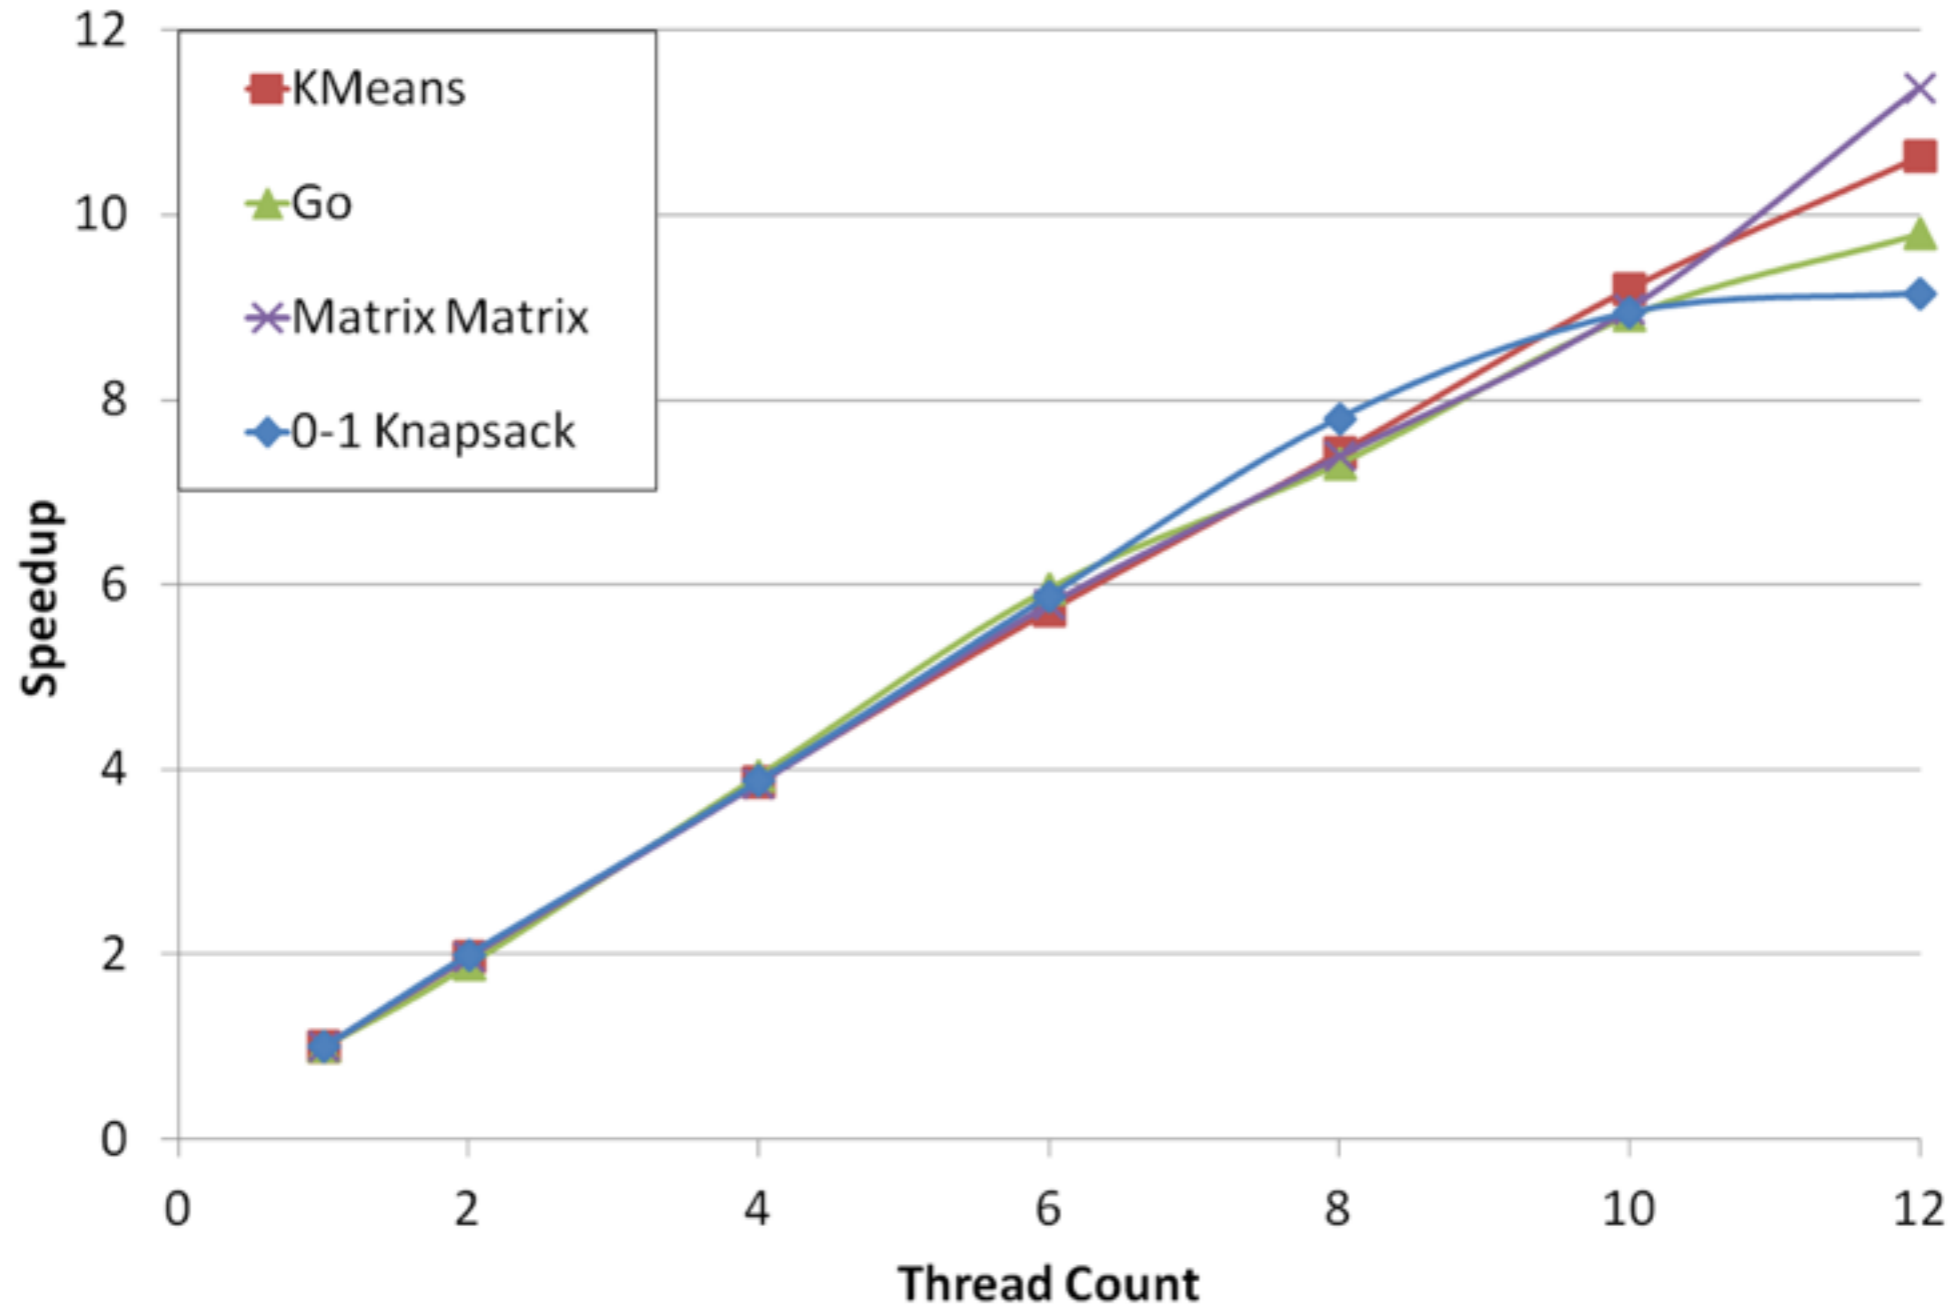
Summary

- Focus on programmability
- Coarse-grained
- Nodes are written as Scala functions
- Edges are normal function arguments
- Real software, plus tooling, available today

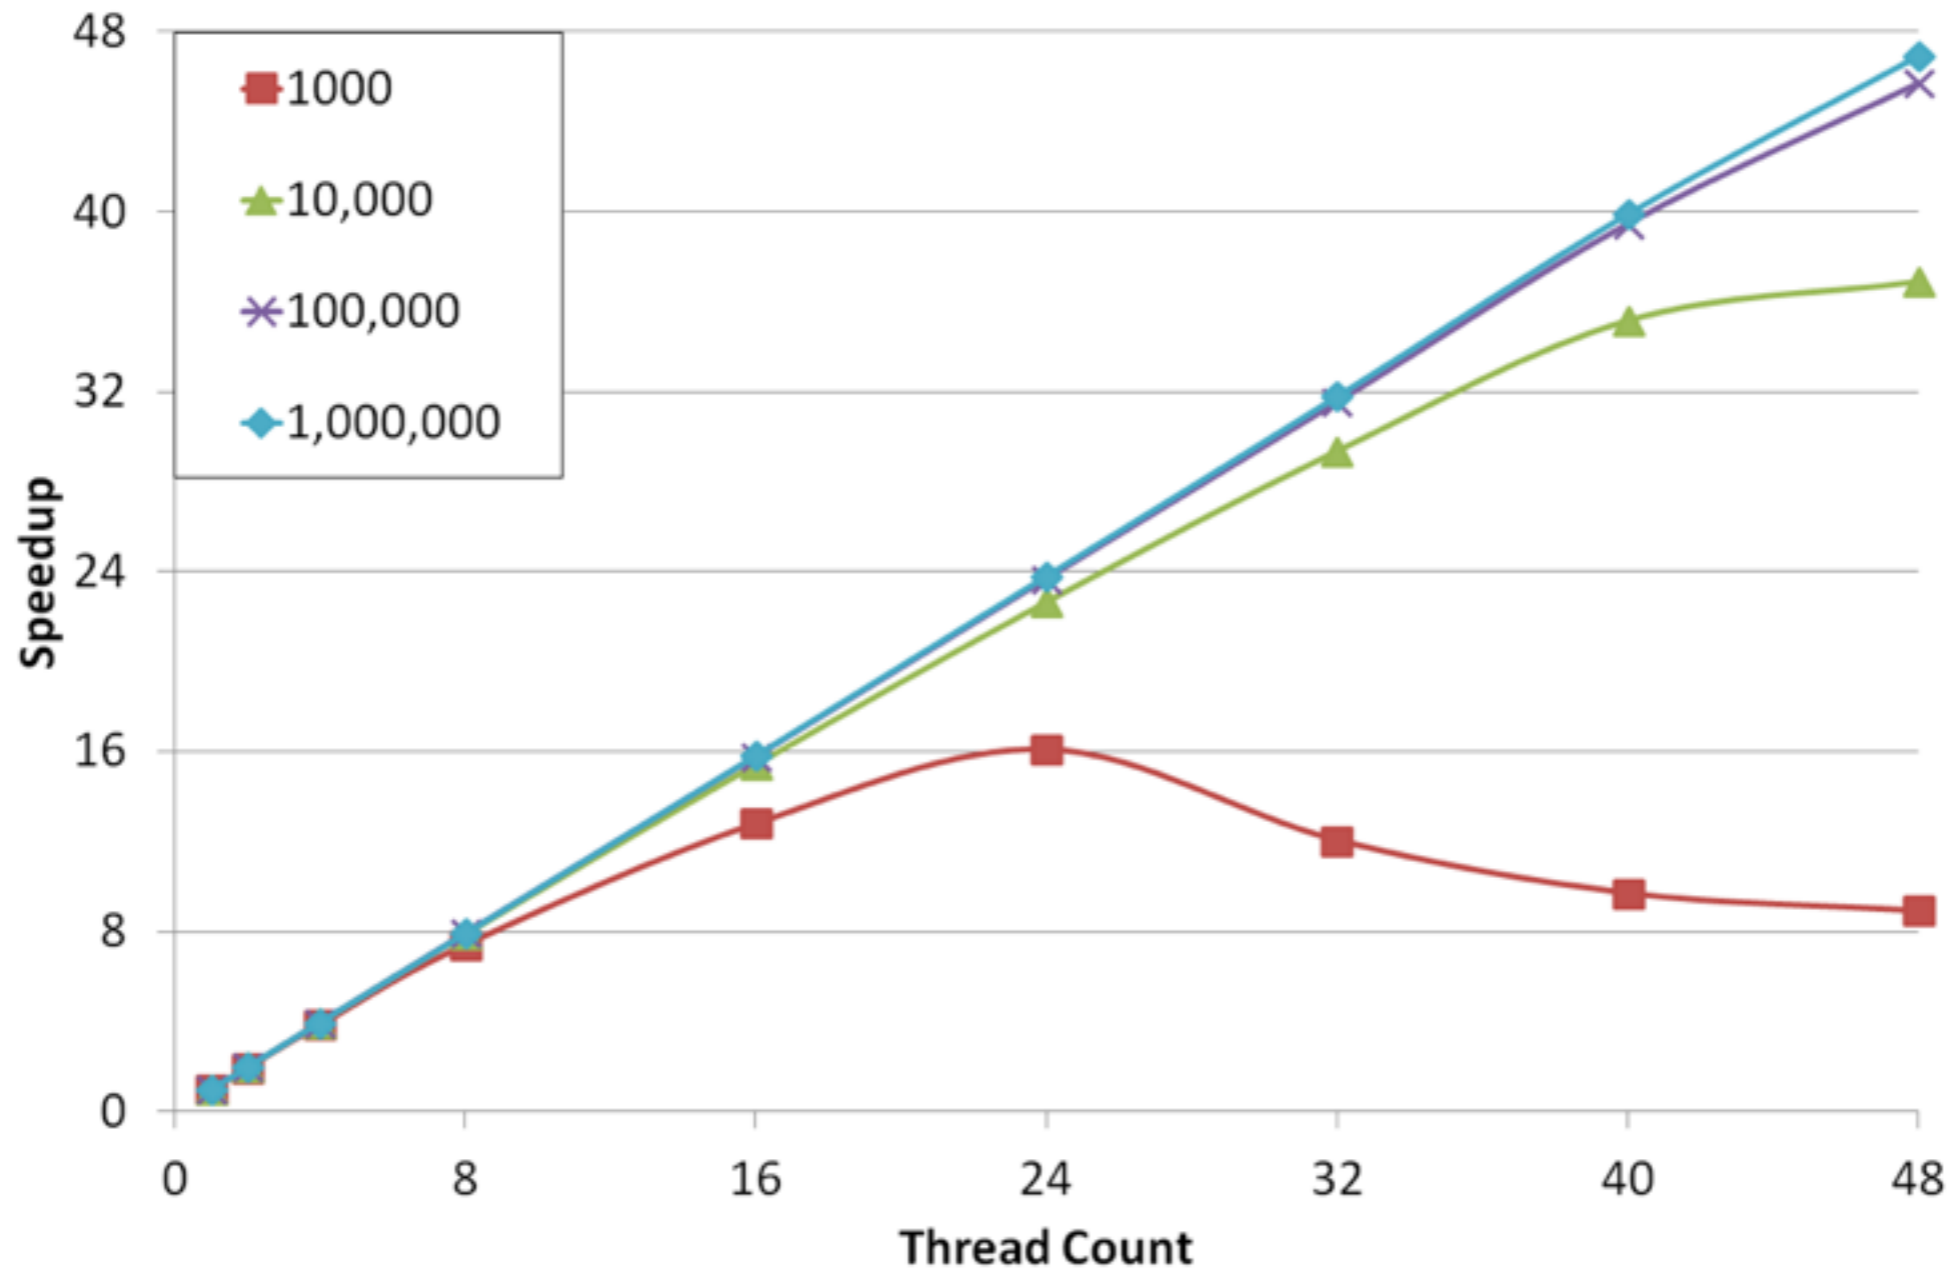
Benchmarks

- Matrix Multiplication
- KMeans
- 0-1 Knapsack
- Go
- Parallel collections

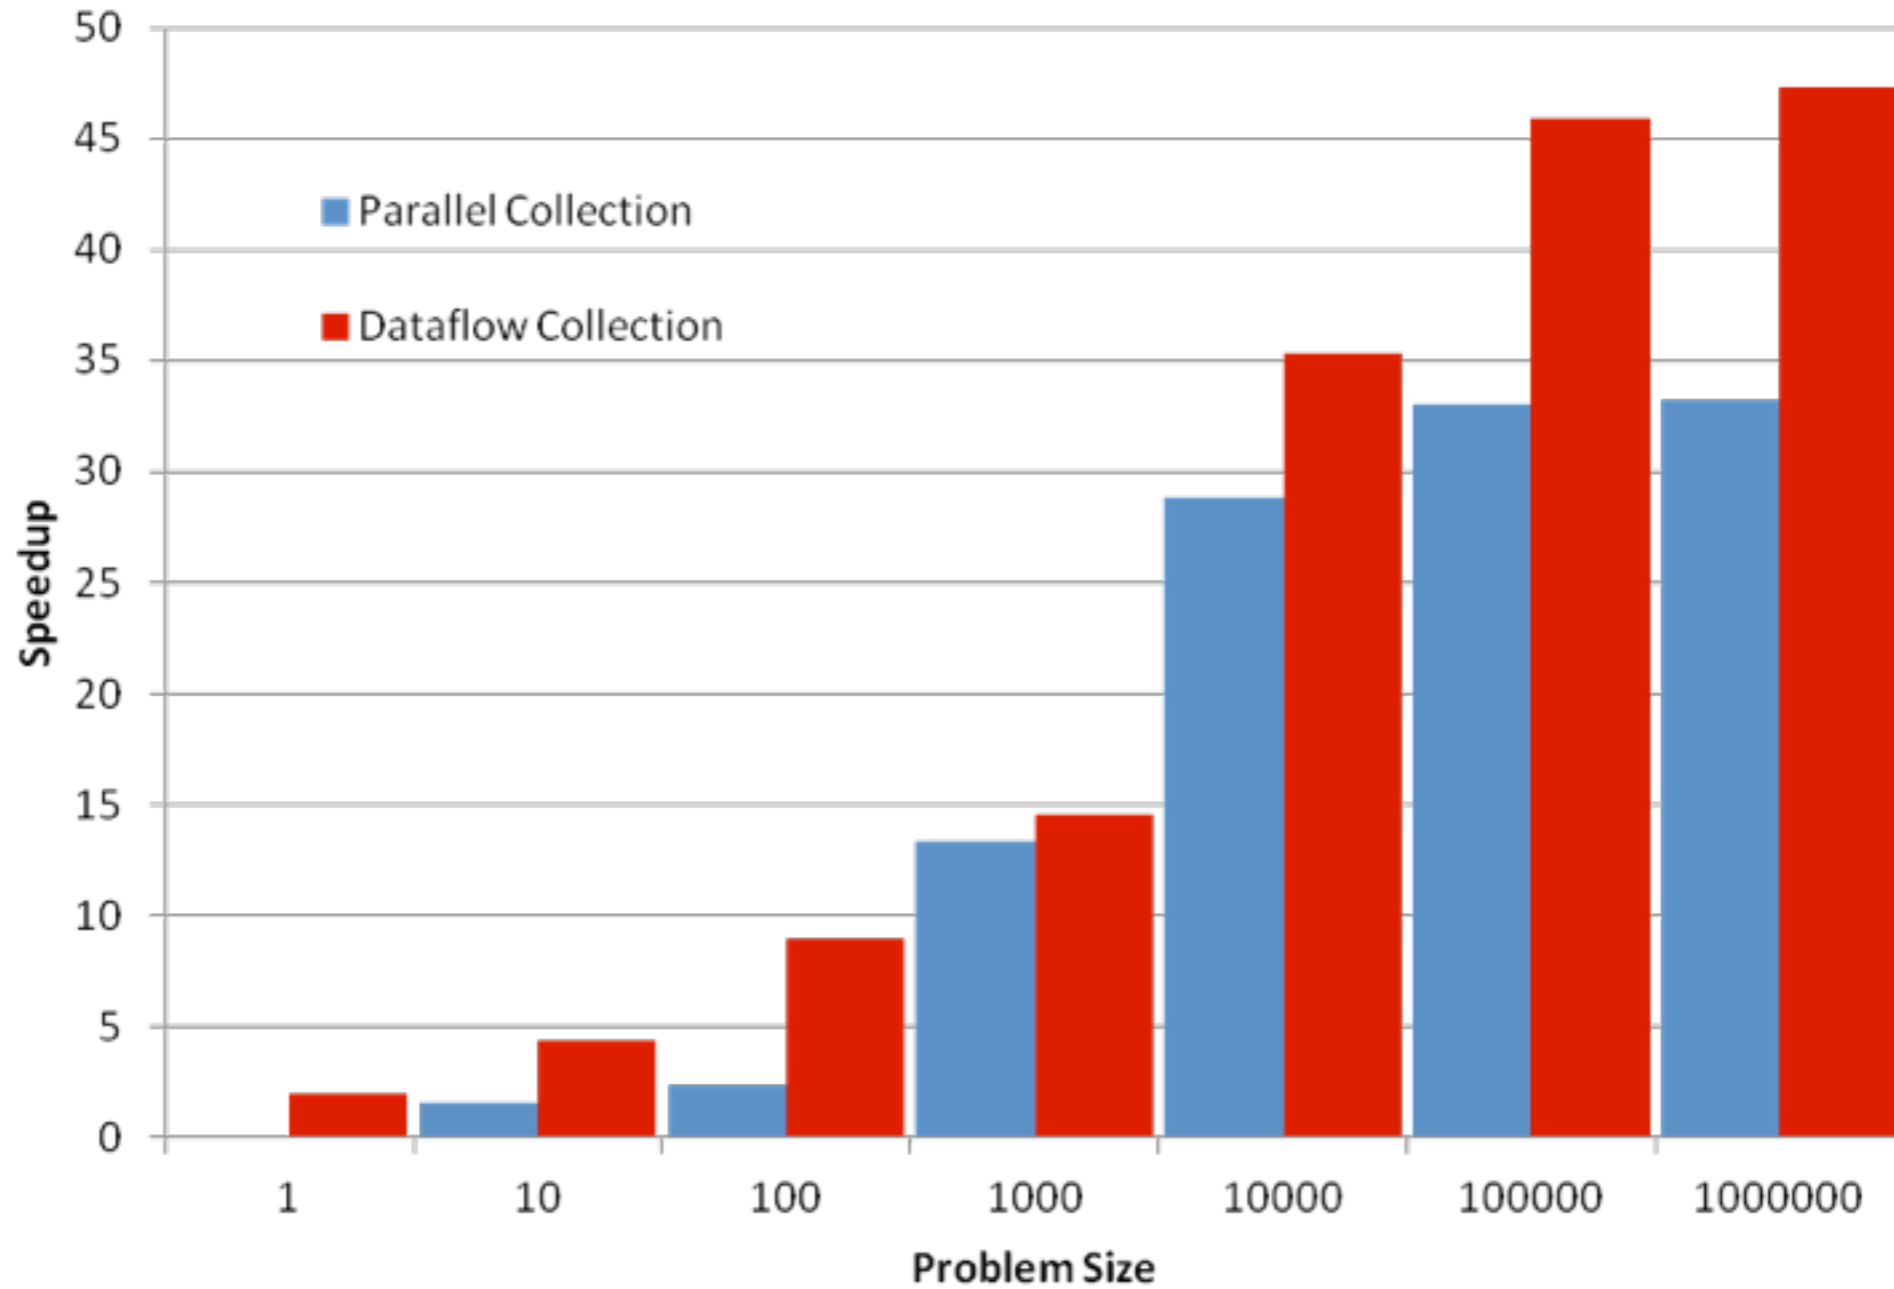




Benchmark Speedup, 2x 6-core AMD Opteron



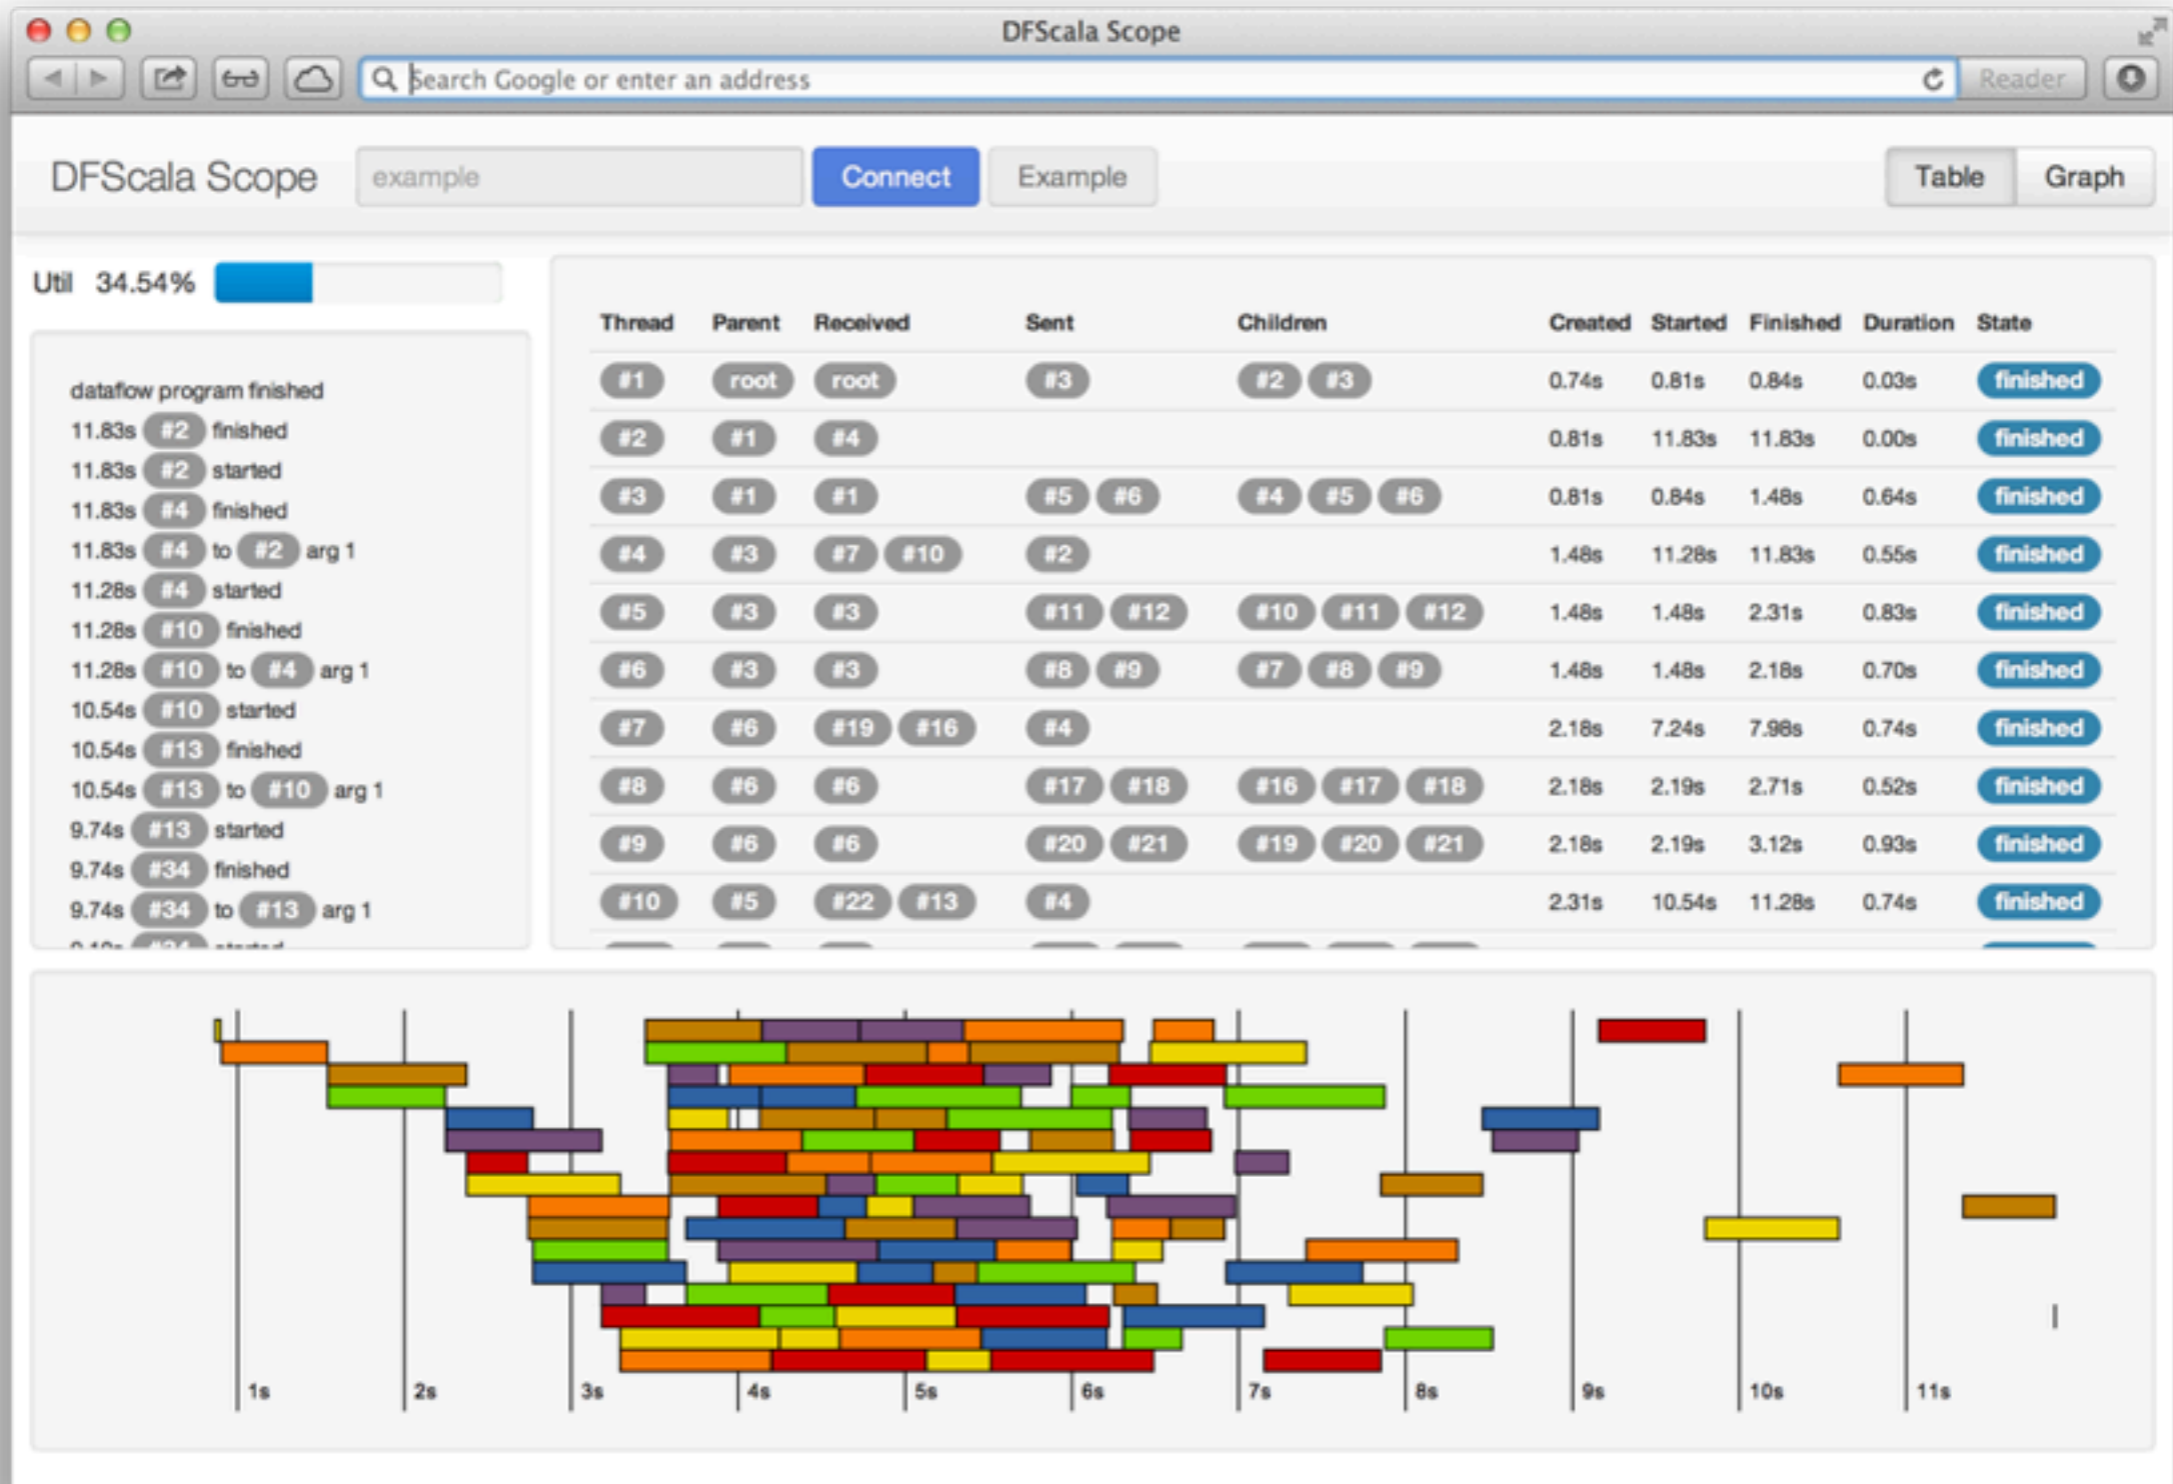
Collections Speedup, 4x 12-core AMD Opteron



**Collections Compared to Scala Parallel Collections,
4x 12-core AMD Opteron**

Tooling

- Pure Scala library
- Loggers for analysing runtime performance
- Graphical debugger



Available Today

<http://apt.cs.man.ac.uk/projects/TERAFLUX/DFScala/>

or Google for 'manchester dfscala'

- Currently used in research into combining dataflow with transactional memory
- Research into irregular parallelism

DFScala in a Nutshell

- Focus on programmability
- Coarse-grained
- Nodes are written as Scala functions
- Edges are normal function arguments
- Graphs are statically checked for type-safety
- Real software, plus tooling, available today
- Empirical results show it is scalable and can be applied to make a high performance collections library