



# Habano-Scala: Async-Finish Programming in Scala

Scala Days  
April 17, 2012

Shams Imam and Vivek Sarkar  
Rice University



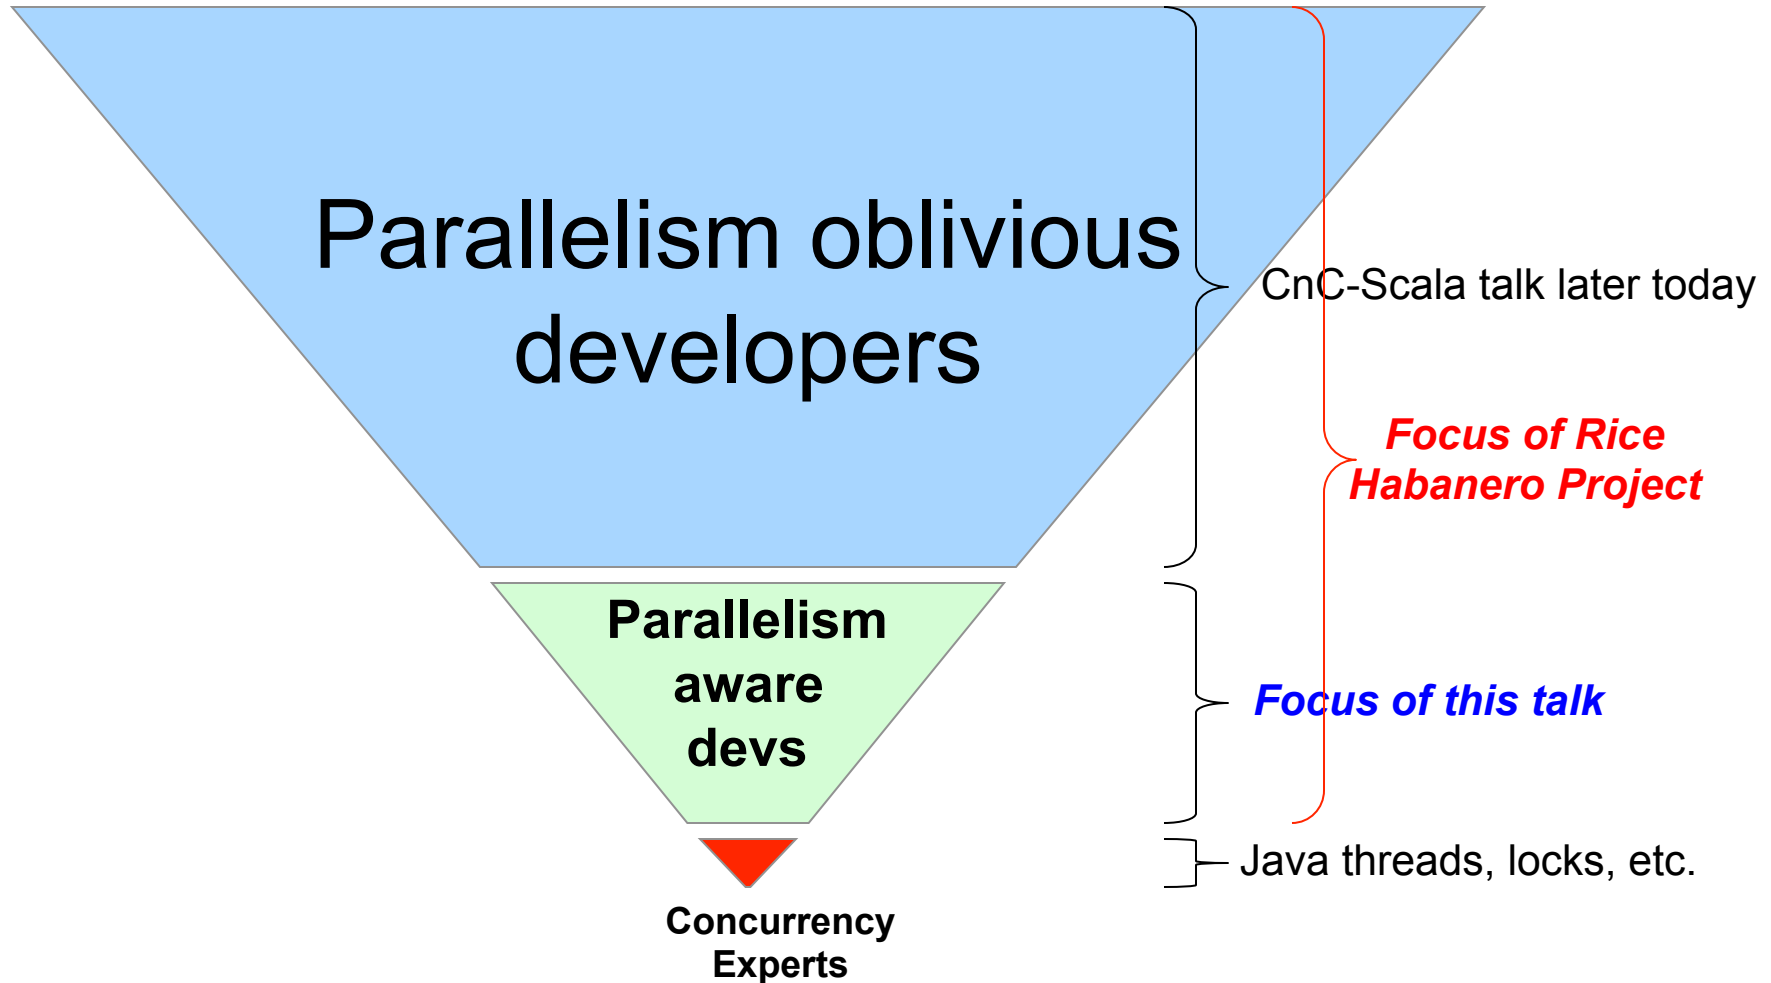
# Introduction



- Multi-core processors
  - Software Concurrency Revolution
  - renewed interest in parallel programming models
- Goal: increase productivity of parallel programming by both simplifying and generalizing current parallel programming models
  - Simplification → increase classes of developers who can write parallel programs
  - Generalization → increase classes of applications that can be supported by a common model



# Inverted Pyramid of Parallel Programming Skills





# Habanero-Scala



- Scala integration of Habanero-Java features
- Habanero-Java
  - developed at Rice University
  - derived from Java-based version of X10 language (v1.5) in 2007
  - targeted at parallelism-aware developers, not necessarily concurrency experts
  - used in sophomore-level undergraduate course on “Fundamentals of Parallel Programming” at Rice
    - <https://wiki.rice.edu/confluence/display/PARPROG/COMP322>
    - Or search for “comp322 wiki”



# Goals for this talk



- Task parallelism
  1. Dynamic task creation & termination
    - async, finish, forall, foreach
  2. Mutual exclusion: isolated
  3. Coordination
    - futures, data-driven futures
  4. Collective and P2P synchronization:
    - phaser, next
  5. Locality control for tasks and data: places
- Actor extensions and unification with task parallelism



# Async and Finish



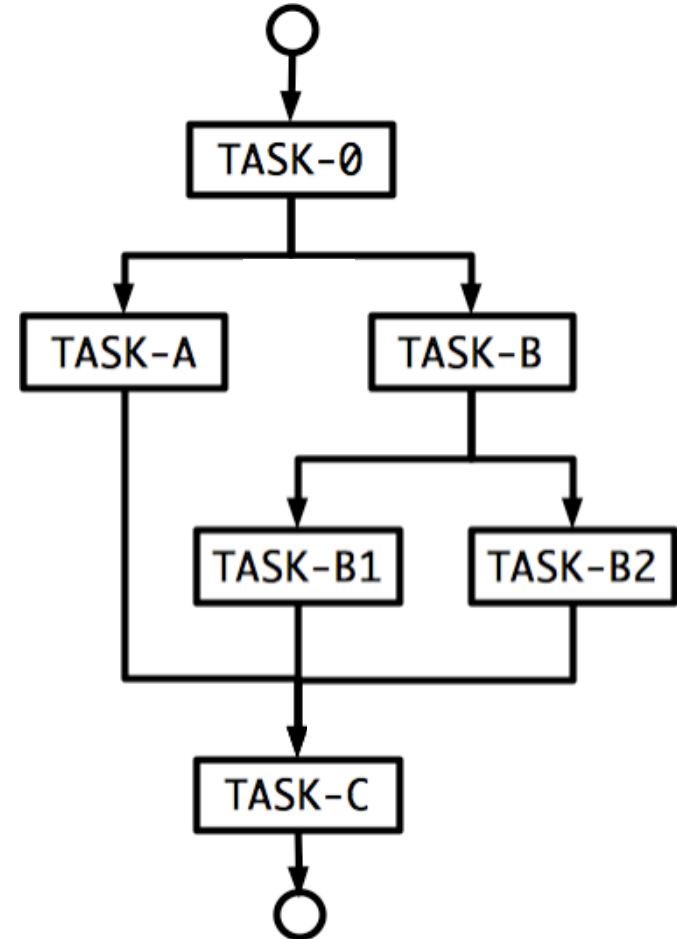
- `async { <stmt> }`
  - creates a new child task that executes `<stmt>`
  - parent task proceeds to operation following the `async`
  - `asyncSeq(<cond>) { S } ≡ if (<cond>) S else async { S }`
- `finish { <stmt> }`
  - execute `<stmt>`, but wait until *all* (transitively) spawned `asyncs` in `<stmt>`'s scope have terminated
  - Implicit finish between start and end of main program
- Async-Finish programs **cannot** create a deadlock cycle



# Async-Finish Example



```
1. // imports
2. object SeqApp extends App {
3.     println("Task 0")
4.
5.     println("Task A")
6.
7.     println("Task B")
8.     println("Task B1")
9.     println("Task B2")
10.    println("Task C")
11. }
```

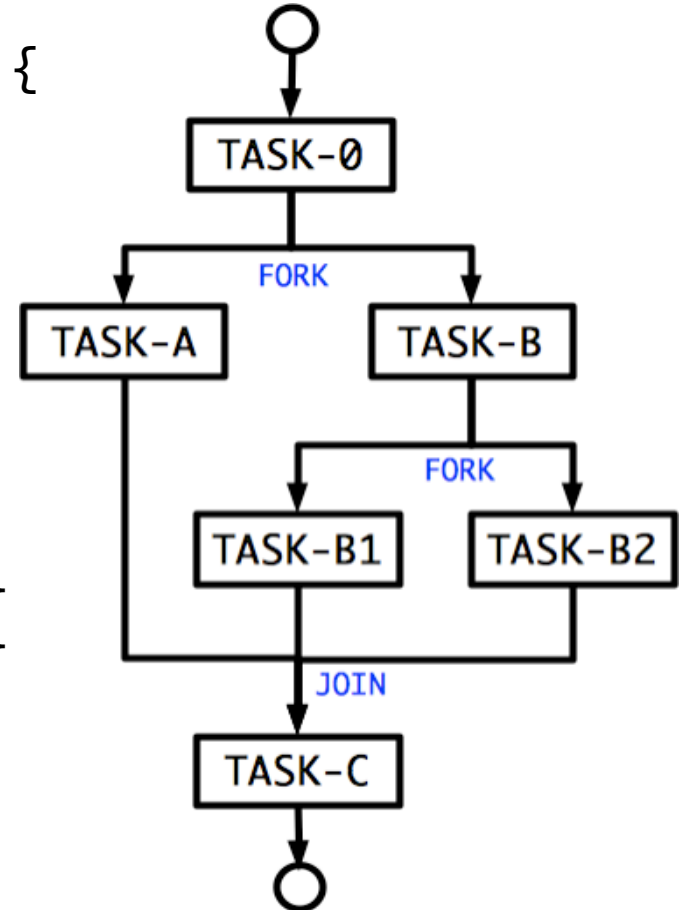




# Async-Finish Example (contd)



```
1. // imports
2. object ParApp extends HabaneroApp {
3.     println("Task 0")
4.     finish {
5.         async {
6.             println("Task A")
7.         }
8.         async {
9.             println("Task B")
10.            async { println("Task B1") }
11.            async { println("Task B2") }
12.        }
13.    }
14.    println("Task C")
15. }
```







# Forall and Foreach



- `forall(start, end) { f(i) } ≡  
 finish { for(i <- start until end) async { f(i) } }`
- `foreach(start, end) { f(i) } ≡  
 for(i <- start until end) async { f(i) }`
- `scala.collection.Iterables` support `asyncForall` and `asyncForeach` as extension methods
- E.g.

```
1 to 20 asyncForeach {  
  i =>  
    println(" i = " + i)  
}
```



# Synchronized access - isolated



- `isolated { <stmt> }`
  - Two tasks executing isolated statements with interfering accesses must perform the isolated statement in mutual exclusion
  - isolated statements can be nested (redundant)
  - support weak isolation, i.e. atomicity is guaranteed only with respect to other statements also executing inside isolated scopes



# Parallel DFS - example



```
1. object DepthFirstSearchApp extends HabaneroApp {
2.     ...
3.     finish { root.parent = root; root.compute() }
4. }

5. class Node() {
6.     def tryLabelling(p: Node): Boolean = {
7.         isolated {
8.             if (parent eq null)
9.                 parent = p
10.        }
11.        (parent eq p)
12.    }

13.    def compute(): Unit = {
14.        neighbors foreach { child =>
15.            if (child.tryLabelling(this)) async {
16.                child.compute()
17.            } } } }
```



# Futures – Tasks with Return Values



- `asyncFuture[T] { <stmt> }`
  - creates a new child task that executes `<stmt>`
  - parent task proceeds to operation following the `async`
  - return value of `<stmt>` must be of type `T`
  - `asyncFuture` expression returns a reference to a *container* of type `habanero.Future[T]`
  - `aFuture.get()` blocks if value is unavailable
    - `aFuture.get()` only waits for specified `async`
- Assignment of future references to final variables **guarantees** deadlock freedom with `get()` operations
- In addition, no data races are possible on future return values



# Futures – example



```
1.  def fib(n: Int): Int = {
2.      if (n < 2) {
3.          n
4.      } else {
5.          val x = asyncFuture {
6.              fib(n - 1)
7.          }
8.          val y = asyncFuture {
9.              fib(n - 2)
10.         }
11.
12.         x.get() + y.get()
13.     }
```



# Data-Driven Futures (DDFs)



- separation of classical “futures” into data (DDF) and control (`asyncAwait`) parts
- Operations:
  - `ddf[T]()`: new instance using factory method
  - `put(someValue)`: only a single `put()` is allowed on the DDF
  - `asyncAwait()`: declare data/control dependency in an `async`
  - `get()`: returns the value associated with the DDF
- Accesses to values inside the DDF are guaranteed to be race-free and deterministic



# DDF – Fib example



```
1. finish {
2.     val res = ddf[Int]()
3.     async {
4.         fib(N, res)
5.     }
6. }
7. println("fib(" + N + ") = " + res.get())

8. def fib(n: Int, v: DataDrivenFuture[Int]): Unit = {
9.     if (n < 2) {
10.        v.put(n)
11.    } else {
12.        val (res1, res2) = (ddf[Int](), ddf[Int]())
13.        async {
14.            fib(n - 1, res1)
15.        }
16.        async {
17.            fib(n - 2, res2)
18.        }
19.        asyncAwait(res1, res2) {
20.            v.put(res1.get() + res2.get())
21.        } } }
```



# Phasers




- Support Collective and Point-to-Point synchronization
- Tasks can register in
  - *signal-only/wait-only mode* for producer/consumer synchronization
  - *signal-wait* mode for barrier synchronization
- *next* operation is guaranteed to be deadlock-free
- HJ programs with phasers, finish, async, async-await (but not isolated) are guaranteed to be deterministic if they are data-race-free





# Phasers – Iterative Averaging example



```
1. finish {
2.   val myPhasers = Array.tabulate[Phaser](n + 2)(i => phaser())
3.   for (index <- 1 to n) {
4.     val (me, left, right) = (index, index - 1, index + 1)
5.     val leftPhaser = myPhasers(left).inMode(PhaserMode.WAIT)
6.     val selfPhaser = myPhasers(me).inMode(PhaserMode.SIG)
7.     val rightPhaser = myPhasers(right).inMode(PhaserMode.WAIT)
8.     asyncPhased(leftPhaser, selfPhaser, rightPhaser) {
9.       for (iter <- 0 until N) {
10.        val loopVal = 0.5 * (dataArray(left) + dataArray(right))
11.        // Allow others to proceed and modify dataArray
12.        next
13.        // update the 'owning' element
14.        dataArray(me) = loopVal
15.        // notify others that value has been updated
16.        next
17.      } } } }
```



# Phaser Accumulators



- A parallel reduction construct which separates reduction computations into the parts of
  - sending data,
  - performing the computation itself, and
  - retrieving the result
- Support two logical operations:
  - `send(value)`: to send a value for accumulation in the current phase
  - `result()`: to receive the accumulated value from the previous phase



# Sum Reduction example



```
1.  finish {
2.    val ph = phaser()
3.    val sumAccum = intAccumulator(Operator.SUM, ph)

4.    for (i <- 1 to 30) asyncPhased(ph.inMode(PhaserMode.SIG)) {
5.      sumAccum.send(i)
6.      sumAccum.send(i + 30)
7.      sumAccum.send(i + 60)
8.    }

9.    asyncPhased(ph.inMode(PhaserMode.WAIT)) {
10.     // wait for the tasks from for to complete
11.     next
12.     val resVal: Int = sumAccum.result()
13.     println("Sum(1..90) = " + resVal)
14.   }
15. }
```



# Places



- Logical location where tasks are run
  - enables locality control and load balancing among worker threads
- `async(<some-place>) { <stmt> }` launches an `async` at the specified place
- Current place can be obtained by invoking `here()`
- Set of places are ordered and `aPlace.next()` and `aPlace.prev()` may be used to cycle through them
- System property, `-Dhs.places p:w`, allows the user to specify how many places (`p`) and workers per place (`w`) the runtime should be initialized with.



# Actors and Async/Finish Tasks



- Actors interact seamlessly with async and finish compliant constructs in Habanero-Scala
- Simplifies termination detection
  - wrap actors in a finish scope
- Parallelize message processing inside actors
- Two actor implementations:
  - compliant with Standard Scala actors, extend from HabaneroActor instead of Actor
  - a more efficient implementation, that extends from the HabaneroReactor class



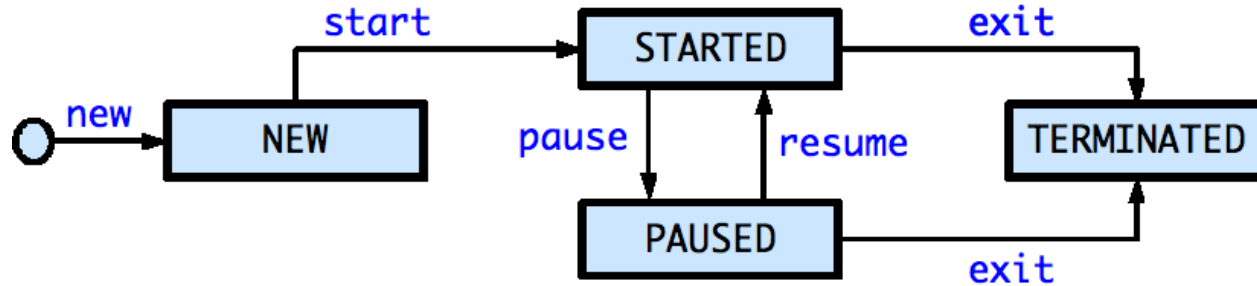
# Example of detecting Actor Termination using finish



```
1. object LightActorApp extends HabaneroApp {
2.   finish {
3.     val pong = new PongActor().start()
4.     val ping = new PingActor(msgs, pong).start()
5.     ping ! StartMessage()
6.   }
7.   println("Both actors terminated")
8. }
9. // class PingActor not displayed
10. class PongActor extends HabaneroActor {
11.   var pongCount = 0
12.   def act() {
13.     loop { react {
14.       case PingMessage =>
15.         sender ! PongMessage
16.         pongCount = pongCount + 1
17.       case StopMessage =>
18.         exit('stop)
19.   } } } }
```



# Pause/Resume extension for Actors



- paused state
  - actor will no longer process messages sent to it
- new operations:
  - `pause()`: move from started to paused state
  - `resume()`: move from paused to started state
- Pausing an actor prevents it from processing the next message until it is resumed



# Non-blocking receives



- Simulates synchronous communication **without** blocking

```
1. class ActorPerformingReceive extends HabaneroReactor {
2.   override def behavior() = {
3.     case msg: SomeMessage =>
4.       ...
5.       val theDdf = ddf[ValueType]()
6.       anotherActor ! new Message(theDdf)
7.       pause() // delay processing next message
8.       asyncAwait(theDdf) {
9.         val responseVal = theDdf.get()
10.        // process the current message
11.        ...
12.        resume() // enable next message processing
13.      }
14.      // return in paused state
15.      ...
16.    } }
```





# Stateless Actors



- Actors with no state, can actively process multiple messages without violating actor constraints

```
1. class StatelessActor() extends Habanero-Rea/A-ctor {
2.     ...
3.     override def behavior() = {
4.         case msg: SomeMessage =>
5.             async {
6.                 processMessage(msg)
7.             }
8.         if (enoughMessagesProcessed) {
9.             exit()
10.        }
11.        // return immediately to process next message
12.    } }
```



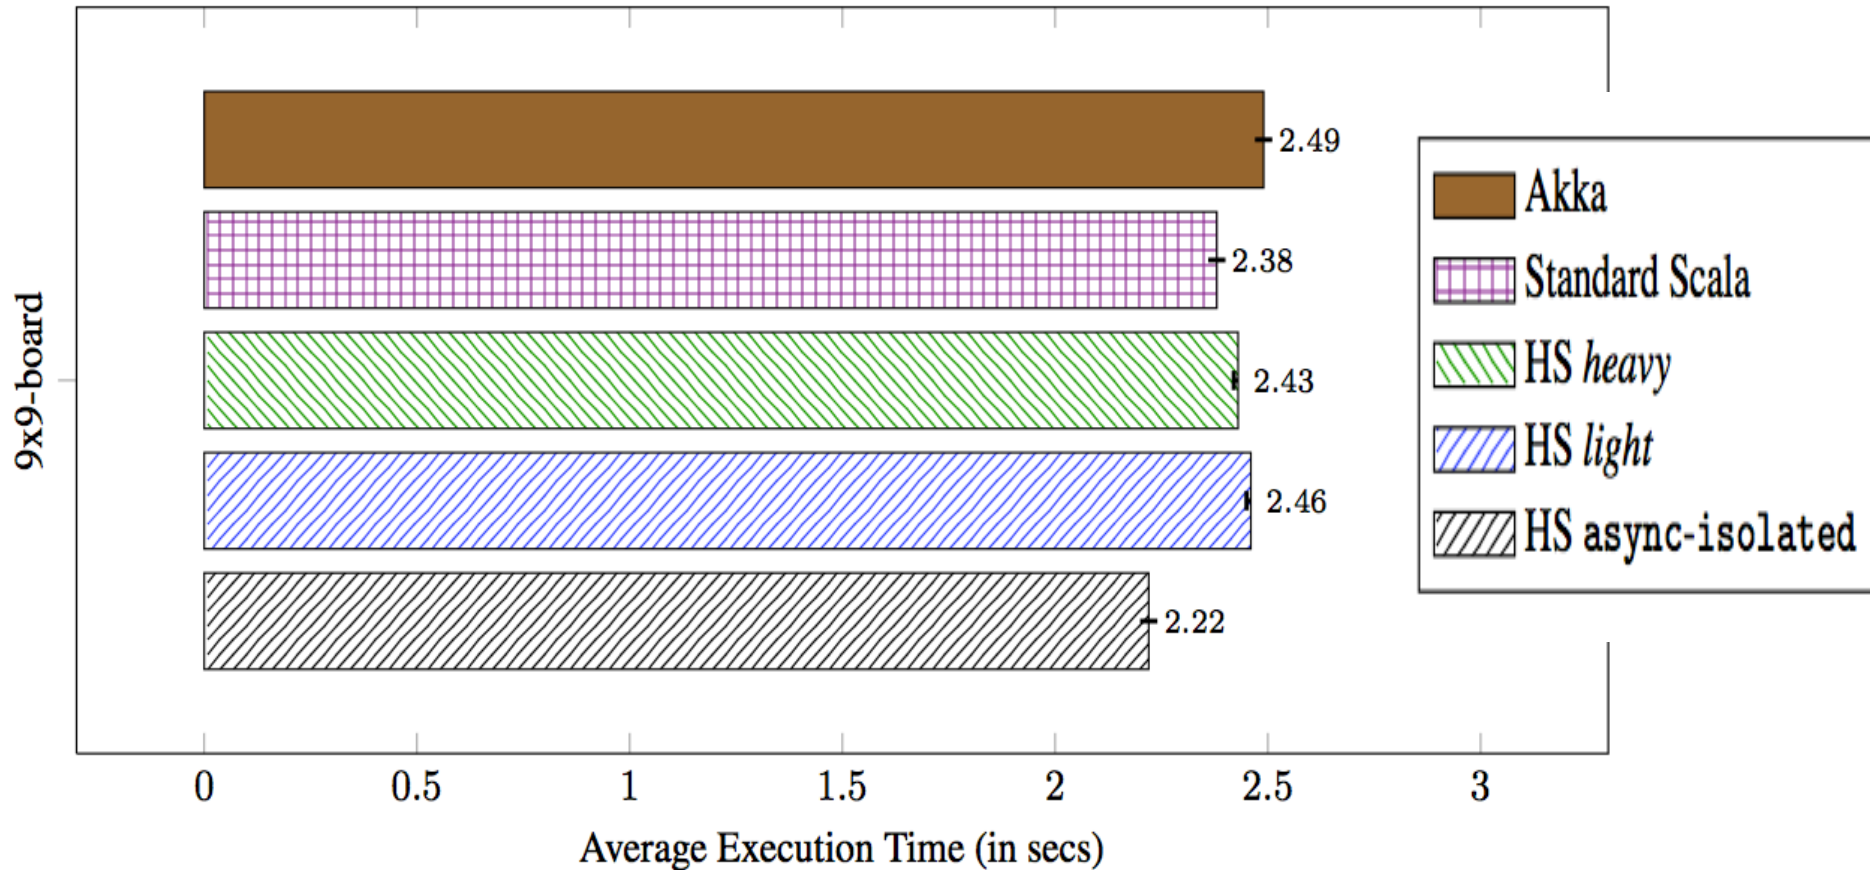
# Experimental Setup



- 12-core (two hex-cores) 2.8 GHz Intel Westmere SMP
- 48 GB memory, running Red Hat Linux (RHEL 6.0)
- Hotspot JDK 1.7
- Scala version 2.9.1-1
- Habanero-Scala 0.1.3
- Arithmetic mean of last thirty iterations from hundred iterations on ten separate JVM invocations



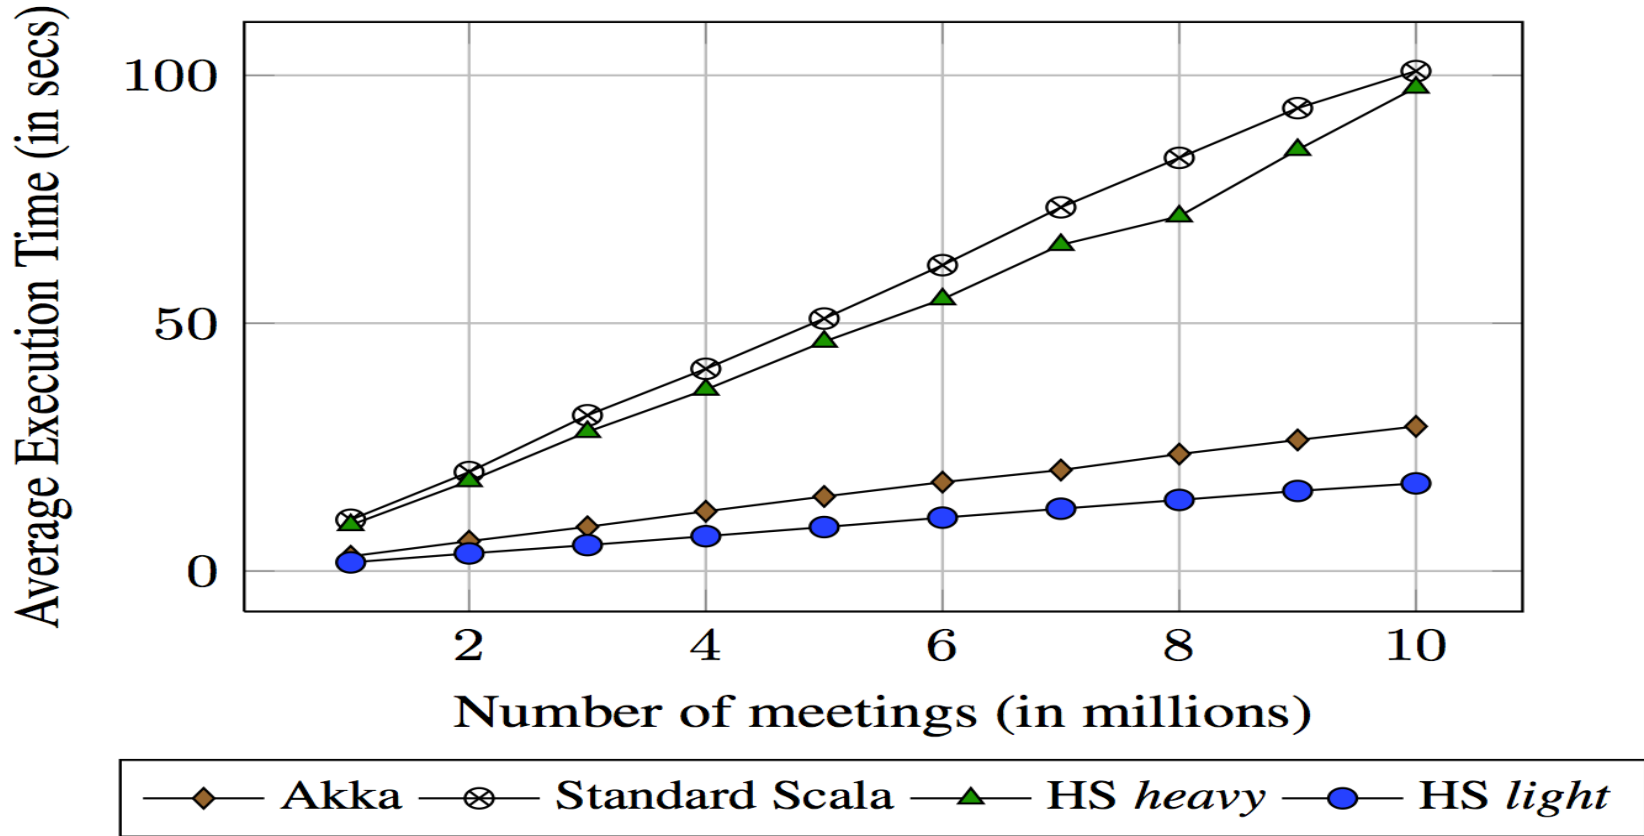
# Sudoku – Constraint Satisfaction



- Actors use Master-Worker style and perform similarly
- Async-Isolated version 7% faster than the actor solutions, about 10% faster than other HS solutions.




# Chameneos Benchmark



- Measures effects of contention – adding messages to mailbox
- HS *light* actors performs best



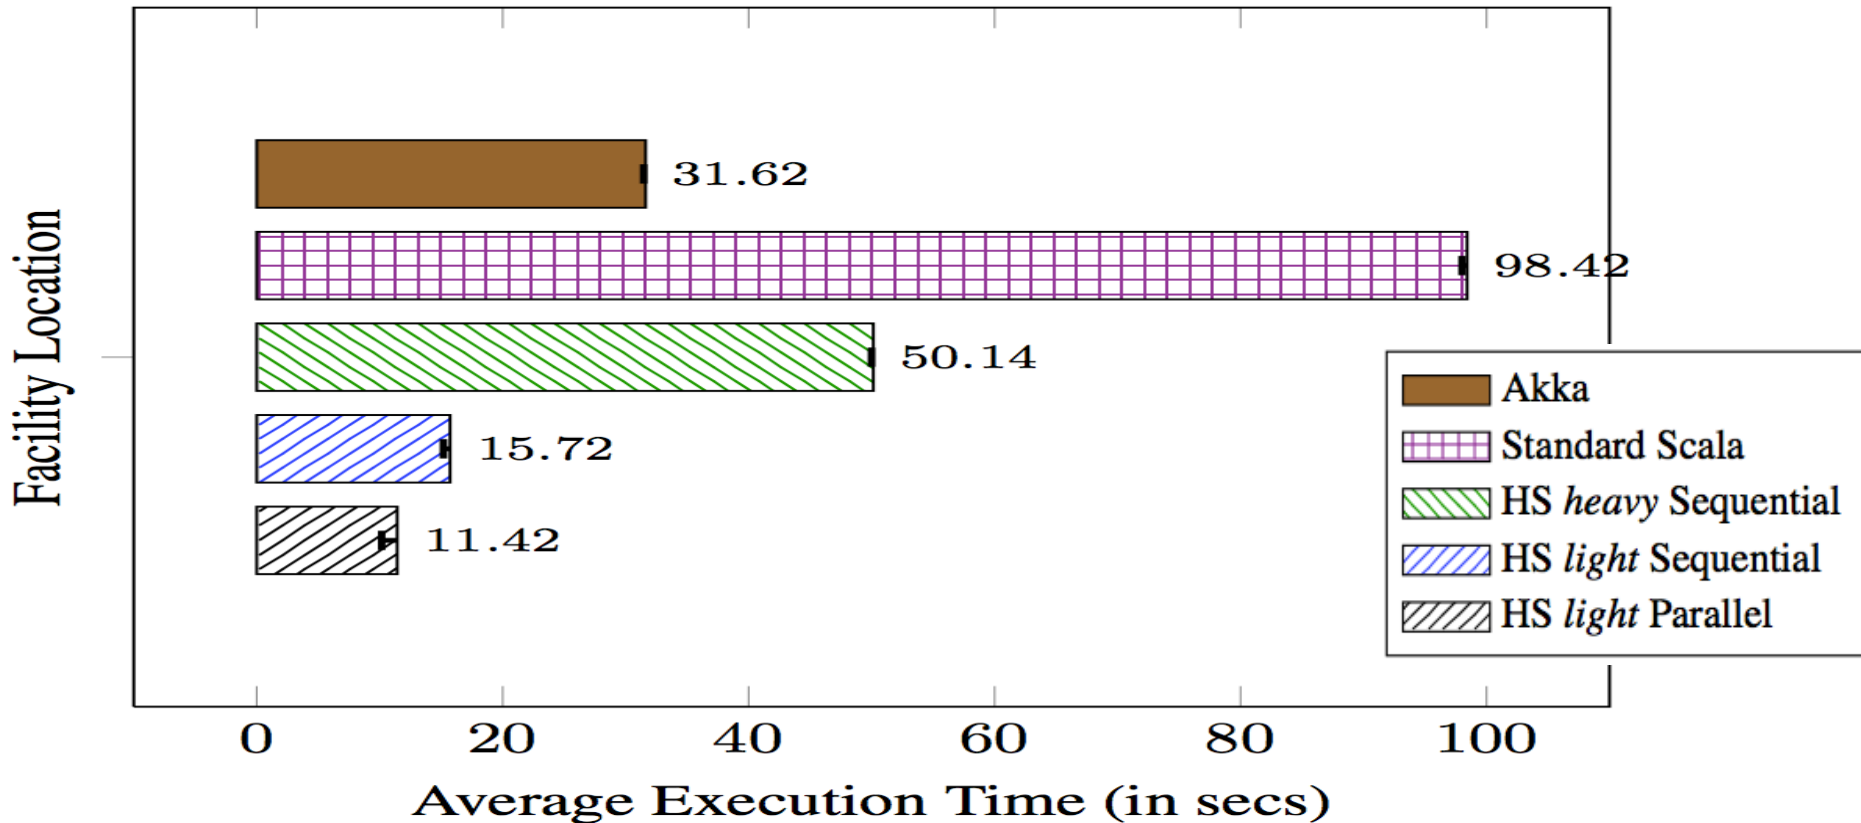
# Prime Sieve Benchmark



- Example of a dynamic pipeline, good fit for actors.
- HS places and thread binding benefits



# Hierarchical Facility Location



- Hybrid solution fastest, about 27% faster than any of the other Actor solutions



# Summary



- HS is a safe and powerful mid-level parallel language
  - programmers with a basic knowledge of Scala to get started quickly with expressing a wide range of parallel patterns
  - Deadlock freedom for programs using finish, async, futures, phasers, isolated
  - Data-race freedom for values accessed through futures and data-driven futures
  - Simplifies writing actor programs
- Runs on standard JRE's and delivers good performance on multicore SMPs
- Available for download at:

<http://habanero-scala.rice.edu/>



# Acknowledgments



- Habanero Group
  - Vincent Cavé
  - Dragos Sbirlea
  - Sagnak Tasirlar





# Thank you!

