



Disciplined Concurrent Models of Computation for Parallel Software

Edward A. Lee

*Robert S. Pepper Distinguished Professor and
UC Berkeley*

*Invited Keynote Talk
2008 Summer Institute
The Concurrency Challenge: Can We Make Parallel Programming Popular?
Sponsors: University of Washington and Microsoft Research*

*Blaine, WA
August 3 to August 7, 2008*



A Tired Debate...

- Shared memory
 - Threads, semaphores, mutexes, monitors...
- Message Passing
 - Synchronous, asynchronous, buffered, ...

This debate is a red herring!

*The choice of shared memory vs. message
passing should be left up to the architects and
compilers, and not up to the programmers.*

Lee, Berkeley 2



Shared memory is not an acceptable programmer's model

Nontrivial software written with threads, semaphores, and mutexes are incomprehensible to humans.

Lee, Berkeley 3



Consider a Simple Example

“The *Observer pattern* defines a one-to-many dependency between a subject object and any number of observer objects so that when the subject object changes state, all its observer objects are notified and updated automatically.”

Design Patterns, Eric Gamma, Richard Helm, Ralph Johnson, John Vlissides (Addison-Wesley Publishing Co., 1995. ISBN: 0201633612):

Lee, Berkeley 4



Observer Pattern in Java

```
public void addListener(listener) {  
    myListeners.add(listener);  
}  
  
public void setValue(newValue) {  
    myValue = newValue;  
  
    for (listener : myListeners) {  
        listener.valueChanged(newValue)  
    }  
}
```

Will this work in a
multithreaded context?

Thanks to Mark S. Miller for the details
of this example.

Lee, Berkeley 5



Observer Pattern With Mutual Exclusion

```
public synchronized void addListener(listener) {  
    myListeners.add(listener);  
}  
  
public synchronized void setValue(newValue) {  
    myValue = newValue;  
  
    for (listener : myListeners) {  
        listener.valueChanged(newValue)  
    }  
}
```

JavaSoft recommends against this.
What's wrong with it?

Lee, Berkeley 6



Mutexes are Minefields

```
public synchronized void addListener(Listener) {  
    myListeners.add(Listener);  
}
```

```
public synchronized void setValue(newValue) {  
    myValue = newValue;  
  
    for (Listener : myListeners) {  
        Listener.valueChanged(newValue)  
    }  
}
```

valueChanged() may attempt to acquire a lock on some other object and stall. If the holder of that lock calls addListener(), deadlock!

Lee, Berkeley 7





Should have used tools...

Tsk, tsk...

Should have used static analysis tools to detect the deadlock potential...

But detection is only half the battle...

How to fix it?

Lee, Berkeley 9



Simple Observer Pattern Becomes Not So Simple

```
public synchronized void addListener(Listener) {
    myListeners.add(Listener);
}

public synchronized void setValue(newValue) {
    synchronized(this) {
        myValue = newValue;
        listeners = myListeners.clone();
    }
    for (listener : listeners) {
        listener.valueChanged(newValue)
    }
}
```

while holding lock, make copy of listeners to avoid race conditions

notify each listener outside of synchronized block to avoid deadlock

**This still isn't right.
What's wrong with it?**

Lee, Berkeley 10



Simple Observer Pattern: How to Make It Right?

```
public synchronized void addListener(Listener) {
    myListeners.add(Listener);
}

public synchronized void setValue(newValue) {
    synchronized(this) {
        myValue = newValue;
        listeners = myListeners.clone();
    }
    for (Listener : listeners) {
        Listener.valueChanged(newValue)
    }
}
```

Suppose two threads call setValue(). One of them will set the value last, leaving that value in the object, but listeners may be notified in the opposite order. The listeners may be alerted to the value changes in the wrong order!

Lee, Berkeley 11



If the simplest design patterns yield such problems, what about non-trivial designs?

```
/**
 * CrossRefList is a list that maintains pointers to other CrossRefLists.
 * ...
 * @author Geroncio Galicia, Contributor: Edward A. Lee
 * @version $Id: CrossRefList.java,v 1.78 2004/04/29 14:50:00 eal Exp $
 * @since Ptolemy II 0.2
 * @Pt.ProposedRating Green (eal)
 * @Pt.AcceptedRating Green (bart)
 */
public final class CrossRefList implements Serializable {
    ...
    protected class CrossRef implements Serializable{
        ...
        // NOTE: It is essential that this method not be
        // synchronized, since it is called by _farContainer(),
        // which is. Having it synchronized can lead to
        // deadlock. Fortunately, it is an atomic action,
        // so it need not be synchronized.
        private Object _nearContainer() {
            return _container;
        }

        private synchronized Object _farContainer() {
            if (_far != null) return _far._nearContainer();
            else return null;
        }
        ...
    }
}
```

Code that had been in use for four years, central to Ptolemy II, with an extensive test suite with 100% code coverage, design reviewed to yellow, then code reviewed to green in 2000, causes a deadlock during a demo on April 26, 2004.

Lee, Berkeley 12



Perhaps Concurrency is Just Hard...

Sutter and Larus observe:

“humans are quickly overwhelmed by concurrency and find it much more difficult to reason about concurrent than sequential code. Even careful people miss possible interleavings among even simple collections of partially ordered operations.”

H. Sutter and J. Larus. Software and the concurrency revolution. ACM Queue, 3(7), 2005.

Lee, Berkeley 13



If concurrency were intrinsically hard, we would not function well in the physical world



It is not concurrency that is hard...

Lee, Berkeley 14



...It is shared memory that is Hard!

Threads are sequential processes that share memory. From the perspective of any thread, the entire state of the universe can change between any two atomic actions (itself an ill-defined concept).

Imagine if the physical world did that...

Lee, Berkeley 15



Concurrent programs using shared memory are incomprehensible because concurrency in the physical world does not work that way.

We have no experience!

Lee, Berkeley 16



Succinct Problem Statement

Threads are wildly nondeterministic.

The programmer's job is to prune away the nondeterminism by imposing constraints on execution order (e.g., mutexes, transactions) and limiting shared data accesses (e.g., OO design).

Lee, Berkeley 17



We have incrementally improved threads

- Object Oriented programming
- Coding rules (Acquire locks in the same order...)
- Libraries (Stapl, Java 5.0, ...)
- PGAS (UPC, Co-array FORTRAN, Titanium, ...)
- Patterns (MapReduce, ...)
- Transactions (Databases, ...)
- Formal verification (Blast, thread checkers, ...)
- Enhanced languages (Split-C, Cilk, Guava, ...)
- Enhanced mechanisms (Promises, futures, ...)

**But is it enough to refine a mechanism
with flawed foundations?**

Lee, Berkeley 18



E.g.: Verification:
Thread checkers?

Consider what it would take for static analysis tools to detect the out-of-order notification bug in our implementation of the listener pattern...

We want to tolerate a race on calls to `setValue()` (application nondeterminism), but not on notifications to listeners. To specify this for a verifier, we have to have already solved the problem!

Lee, Berkeley 19



For a brief optimistic instant, *transactions* looked like they might save us...

“TM is not as easy as it looks (even to explain)”

Michael L. Scott, invited keynote, (EC)2
Workshop, Princeton, NJ, July 2008

Lee, Berkeley 20



Do we have a sound foundation for concurrent programming?

If the foundation is bad, then we either tolerate *brittle designs* that are difficult to make work, or we have to rebuild from the foundations.



Note that this whole enterprise is held up by threads

Lee, Berkeley 21

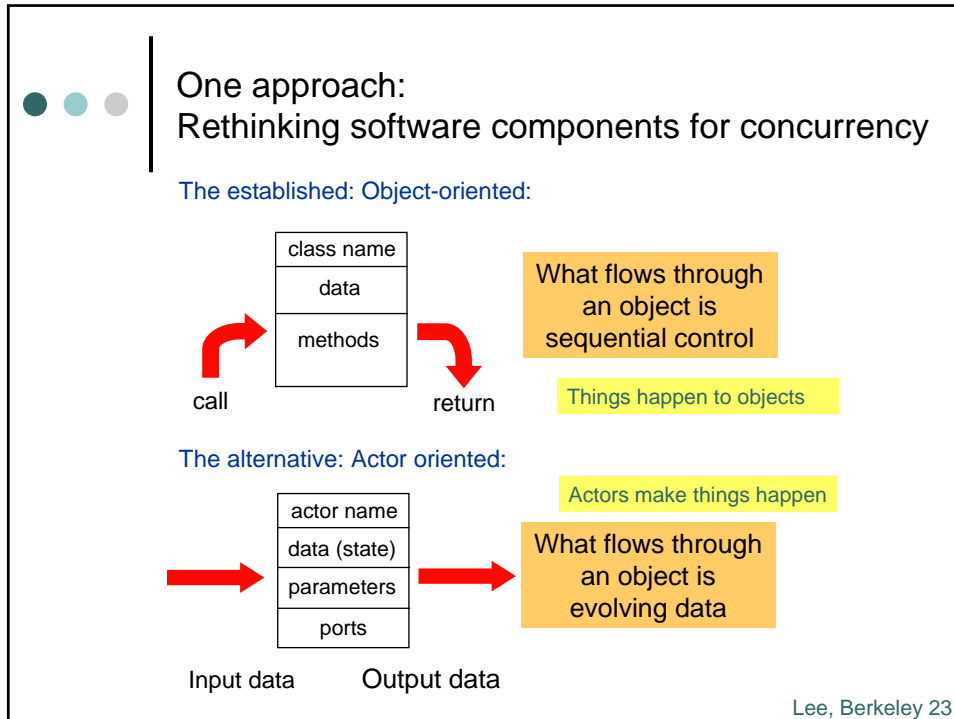


So, the answer must be message passing, right?

Not quite...

More discipline is needed than what is provided by today's message passing libraries.

Lee, Berkeley 22



- ## Examples of Actor-Oriented Systems
- Unix pipes
 - Dataflow systems
 - CORBA event service (distributed push-pull)
 - ROOM and UML-2 (dataflow, Rational, IBM)
 - VHDL, Verilog (discrete events, Cadence, Synopsys, ...)
 - LabVIEW (structured dataflow, National Instruments)
 - Modelica (continuous-time, constraint-based, Linkoping)
 - OPNET (discrete events, Opnet Technologies)
 - SDL (process networks)
 - Occam (rendezvous)
 - Simulink (Continuous-time, The MathWorks)
 - SPW (synchronous dataflow, Cadence, CoWare)
 - ...
- The semantics of these differ considerably, but all provide more understandable ways of expressing concurrency.*
- Lee, Berkeley 24



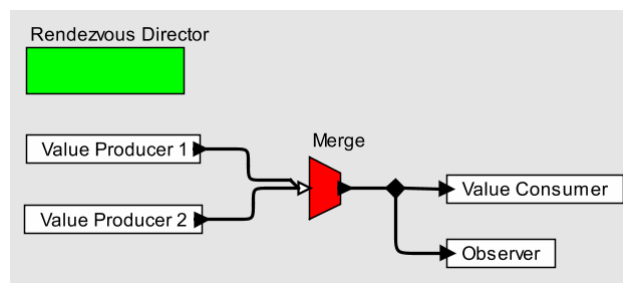
Recall the Observer Pattern

“The *Observer pattern* defines a one-to-many dependency between a subject object and any number of observer objects so that when the subject object changes state, all its observer objects are notified and updated automatically.”

Lee, Berkeley 25



Observer Pattern using an Actor-Oriented Language with Rendezvous Semantics



Each actor is a process, communication is via rendezvous, and the Merge explicitly represents nondeterministic multi-way rendezvous.

This is realized here in a *coordination language with a visual syntax*.

Lee, Berkeley 26



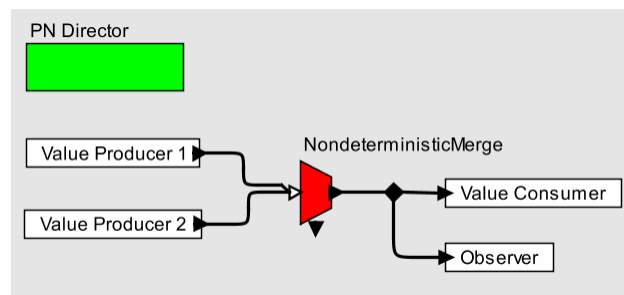
Now that we've made a trivial design pattern trivial, we can work on more interesting aspects of the design.

E.g., suppose we don't care how long notification of the observer is deferred, as long as the observer is notified of all changes in the right order?

Lee, Berkeley 27



Observer Pattern using an Actor-Oriented Language with Kahn Semantics (Extended with Nondeterministic Merge)



Each actor is a process, communication is via streams, and the NondeterministicMerge explicitly merges streams nondeterministically.

Again a coordination language with a visual syntax.

Lee, Berkeley 28

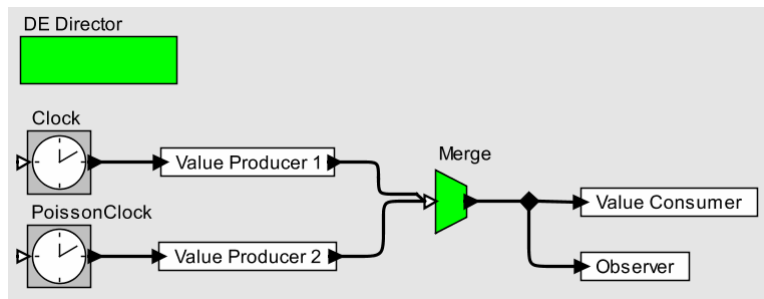


Suppose further that we want to explicitly specify the timing?

Lee, Berkeley 29



Observer Pattern using an Actor-Oriented Language with Discrete Event Semantics



Messages have a (semantic) time, and actors react to messages chronologically. Merge now becomes deterministic.

Again a coordination language with a visual syntax.

Lee, Berkeley 30



Isn't this just message passing?

Each of these realizations of the listener pattern can be implemented with a message passing library (or with threads, for that matter).

But a message passing library allows too much flexibility to yield comprehensible designs.

Its capabilities need to be used judiciously...

Lee, Berkeley 31



Consider for example Determinism

Most programs specify a particular computation, expecting the same input to yield the same output.

Occasionally, programs require nondeterminism, where the input/output relation is not a function. Multiple outputs are possible for the same input.

Regrettably, without considerable sophistication, message passing libraries often yield inadvertently nondeterminate programs.

Lee, Berkeley 32



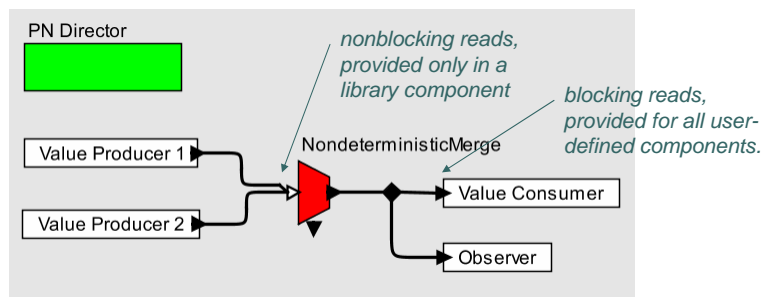
Consider Streams

- In 1974, Gilles Kahn showed that prefix-monotonic functions on streams composed deterministically.
- In 1977, Kahn and MacQueen showed that prefix monotonic functions could be implemented with blocking reads.
- Unix pipes use such blocking reads, and achieve determinate composition, but have limited expressiveness.
- Message passing libraries, however, are more flexible, and unless the programmer has studied Kahn, he is likely to mess up...

Lee, Berkeley 33



A *disciplined* use of streams follows Kahn-MacQueen semantics, except where *explicitly* requested by the programmer.



This becomes a *disciplined model of computation*, trivially easy for the programmer to understand, with excellent analogies in the physical world.

Lee, Berkeley 34



A few disciplined concurrent models of computation

- Kahn process networks
- Dataflow
- Synchronous/reactive systems
- Rendezvous
- Discrete-events
- ...

Each of these has many variants with sometimes subtle differences due to differing constraints imposed on the programmer.

Lee, Berkeley 35



A few variants of dataflow, for example

- Dynamic dataflow [Arvind, 1981]
- Structured dataflow [Matwin & Pietrzykowski 1985]
- K-bounded loops [Culler, 1986]
- Synchronous dataflow [Lee & Messerschmitt, 1986]
- Structured dataflow and LabVIEW [Kodosky, 1986]
- PGM: Processing Graph Method [Kaplan, 1987]
- Synchronous languages [Lustre, Signal, 1980's]
- Well-behaved dataflow [Gao, 1992]
- Boolean dataflow [Buck and Lee, 1993]
- Multidimensional SDF [Lee, 1993]
- Cyclo-static dataflow [Lauwereins, 1994]
- Integer dataflow [Buck, 1994]
- Bounded dynamic dataflow [Lee and Parks, 1995]
- Heterochronous dataflow [Girault, Lee, & Lee, 1997]
- ...

Lee, Berkeley 36



Isn't this just message passing again?

Dataflow models can be built with message passing libraries (and with threads). But should the programmer be asked to handle the considerable subtleties?

Few programmers will get it right...

Lee, Berkeley 37



Some Subtleties


- Termination, deadlock, and livelock (halting)
- Bounding the buffers.
- Fairness
- Parallelism
- Data structures and shared data
- Determinism
- Syntax

Lee, Berkeley 38




Dennis-Style Dataflow

Synchronous Dataflow



Dynamic Dataflow



Firing rules:
the number of
tokens
required to fire
an actor.

Communication between actors is via potentially unbounded streams of *tokens*. Each actor has *firing rules*, which specify how availability of input tokens enables a computation. Constraints on the firing rules can yield very nice properties.

Without these nice properties, scheduling policy is a very subtle question.

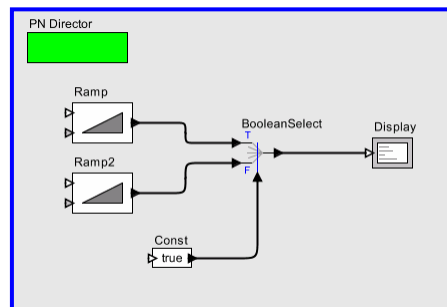
A *signal* or *stream* is a (potentially infinite) sequence of communicated data tokens..

Lee, Berkeley 39



Question 1: Is "Fair" Scheduling a Good Idea?

In the following model, what happens if every actor is given an equal opportunity to run?

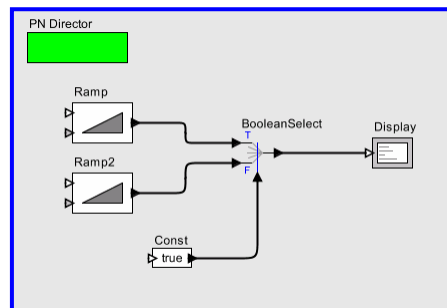


Lee, Berkeley 40



Question 2: Is “Data-Driven” Execution a Good Idea?

In the following model, if actors are allowed to run when they have input data on connected inputs, what will happen?

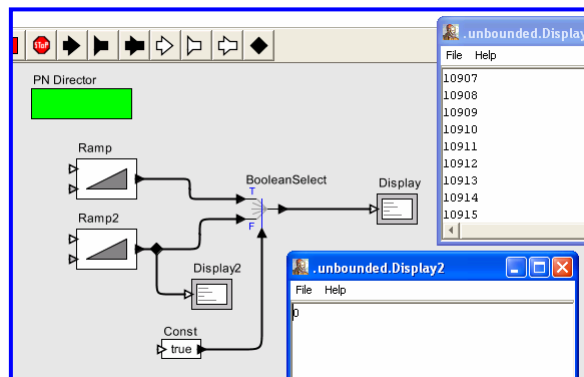


Lee, Berkeley 41



Question 3: When are Outputs Required?

Is the execution shown for the following model the “right” execution?




Lee, Berkeley 42



Question 4: Is “Demand-Driven” Execution a Good Idea?



In the following model, if actors are allowed to run when another actor requires their outputs, what will happen?



Lee, Berkeley 43





Question 5: What is the “Correct” Execution of This Program?



Lee, Berkeley 44



Question 6: What is the Correct Behavior of this Program?



Lee, Berkeley 45



Naïve Schedulers Fail

- Fair
- Demand driven
- Data driven
- Most mixtures of demand and data driven

If programmers are building such programs with message passing libraries or threads, what will keep them from repeating these mistakes that have been made by top experts in the field?

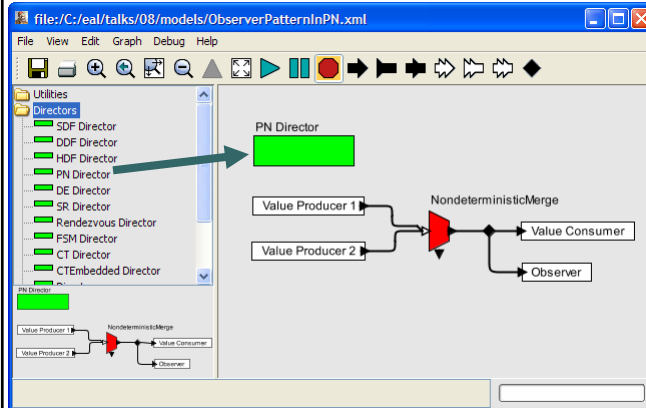
Lee, Berkeley 46



These problems have been solved!
Let's not make programmers re-solve them for every program.

Library of directors

Program using actor-oriented components and a PN MoC



In Ptolemy II, a programmer specifies a *director*, which provides much more structure than message-passing or thread library. It provides a concurrent *model of computation* (MoC).

Lee, Berkeley 47



The PN Director solves the above problems by implementing a “useful execution”

Define a **correct execution** to be any execution for which after any finite time every signal is a prefix of the signal given by the (Kahn) least-fixed-point semantics.

Define a **useful execution** to be a correct execution that satisfies the following criteria:

1. For every non-terminating model, after any finite time, a useful execution will extend at least one stream in finite (additional) time.
2. If a correct execution satisfying criterion (1) exists that executes with bounded buffers, then a useful execution will execute with bounded buffers.


Lee, Berkeley 48



Programmers should not have to figure out how to solve these problems!

Undecidability and Turing Completeness [Buck 93]

Given the following four actors and Boolean streams, you can construct a universal Turing machine:



Hence, the following questions are undecidable:

- Will a model deadlock (terminate)?
- Can a model be executed with bounded buffers?

Lee, Berkeley 49



Our solution: Parks' Strategy [Parks 95]

This “solves” the undecidable problems:

- Start with an arbitrary bound on the capacity of all buffers.
- Execute as much as possible.
- If deadlock occurs and at least one actor is blocked on a write, increase the capacity of at least one buffer to unblock at least one write.
- Continue executing, repeatedly checking for deadlock.


This delivers a useful execution (possibly taking infinite time to tell you whether a model deadlocks and how much buffer memory it requires).

Lee, Berkeley 50




More constrained MoCs yield better to static analysis. E.g.

Synchronous Dataflow (SDF) [Lee & Messerschmitt, 87]




Limit the expressiveness by constraining the number of tokens consumed and produced on each firing to be constant. Eliminates:



Lee, Berkeley 51



Balance Equations



Let q_A, q_B be the number of firings of actors A and B.

Let p_C, c_C be the number of token produced and consumed on a connection C.

Then the system is *in balance* if for all connections C


$$q_A p_C = q_B c_C$$

where A produces tokens on C and B consumes them.

Lee, Berkeley 52

● ● ● | Example

Consider this example:

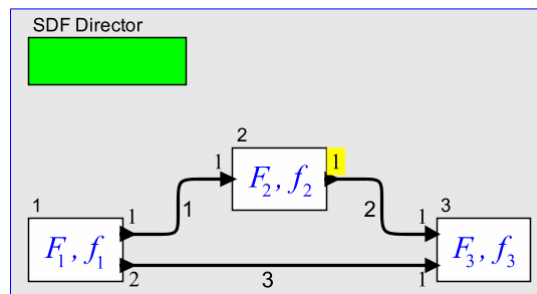


The balance equations imply that actor 3 must fire twice as often as the other two actors.

Lee, Berkeley 53

● ● ● | Inconsistent Models have no Non-Trivial Solution to the Balance Equations

$$\Gamma = \begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \\ 2 & 0 & -1 \end{bmatrix}$$




This production/consumption matrix has rank 3, so there are no nontrivial solutions to the balance equations.

Note that this model can execute forever, but it requires unbounded memory.

Lee, Berkeley 54



Deadlock



Some dataflow models cannot execute forever. In the above model, the feedback loop injects initial tokens, but not enough for the model to execute.

Lee, Berkeley 55



Decidable Models

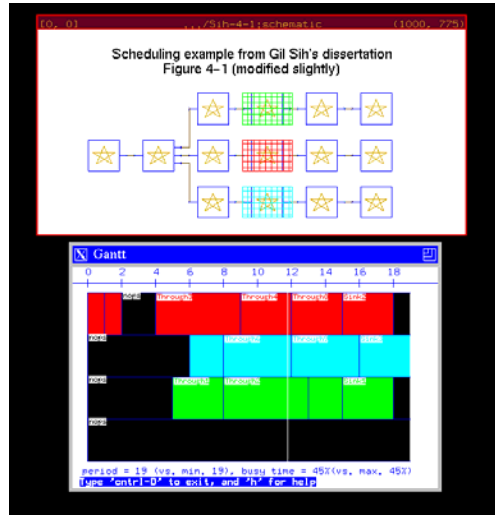
For SDF, boundedness and deadlock are decidable. Moreover, parallel scheduling can be done statically, and useful optimization problems can be solved. See for example:

1. Ha and Lee, "Compile-Time Scheduling and Assignment of Dataflow Program Graphs with Data-Dependent Iteration," *IEEE Trans. on Computers*, November, 1991.
2. Sih and Lee, "Declustering: A New Multiprocessor Scheduling Technique," *IEEE Trans. on Parallel and Distributed Systems*, June 1993.
3. Sih and Lee, "A Compile-Time Scheduling Heuristic for Interconnection-Constrained Heterogeneous Processor Architectures," *IEEE Trans. on Parallel and Distributed Systems*, February 1993.

Lee, Berkeley 56

Gabriel and Ptolemy Classic Leveraged SDF to Generate Parallel Code

SDF model, parallel schedule, and synthesized parallel code (1990)



```
codeblock(std) {
  : initialize address registers for coef and
  delayLineove #saddr(coef)+$val(coefLen)-1,r3
: insert here
  move #ref(delayLineStart),r5
: delayLine
  move #swal(stepSize),x1
  move #ref(operand),x0
  mpyr x0,x1,a
  move a,x0
  move x:(r3),b u:(r5)+,y0
}

codeblock(loop) {
  do #swal(loopVal), $label(endlloop)
  macr x0,y0,b
  move b,x:(r3)-
  move x:(r3),b u:(r5)+,y0
$label(endlloop)
}

codeblock(noloop) {
  macr x0,y0,b
  move b,x:(r3)-
  move x:(r3),b u:(r5)+,y0
}
```

It is an interesting (and rich) research problem to minimize interlocks and communication overhead in complex multirate applications.


Lee, Berkeley 57

Although this makes scheduling decidable, complex optimization problems remain.

Programmers should not have to solve these!

Optimization criteria that might be applied:


- Minimize buffer sizes.
- Minimize the number of actor activations.
- Minimize the size of the representation of the schedule (code size).
- Maximize the throughput.
- Minimize latency.



See Bhattacharyya, Murthy, and Lee, *Software Synthesis from Dataflow Graphs*, Kluwer Academic Press, 1996.

Lee, Berkeley 58

Example: Minimum Buffer Schedule for a 6-Actor Dataflow Model



```

ABABCABCABABCBCDEAFFFFFFBABCABCABABCDE
AFFFFFFBCABABCABCABABCDEAFFFFFFBCABABC
DEAFFFFFFBABCABCABABCBCDEAFFFFFFBABCABC
ABCDEAFFFFFFBCABABCABCABABCDEAFFFFFFBCA
FFFFFFBABCABCBCDEAFFFFFFBABCABCABABCBCDEAF
FFFFFFBABCABCABABCBCDEAFFFFFFBCABABCABCABABC
DEAFFFFFFBCABABCBCDEAFFFFFFBABCABCABABCABCA
BCDEAFFFFFFBABCABCABABCBCDEAFFFFFFBCAFFFFFFB
ABCABCABABCBCDEAFFFFFFBCABABCBCBCDEAFFFFFFBA
BCABCABABCBCBCDEAFFFFFFBABCABCABABCBCDEAFFF
FFBCABABCABCABABCBCDEAFFFFFFBCABABCBCBCDEAF
FFFFFFBABCABCABABCBCBCDEAFFFFFFFEBAFFFFFFBCABC
ABABCBCDEAFFFFFFBCABABCABCABABCBCDEAFFFFFFBCA
BABCABCBCDEAFFFFFFBABCABCABABCBCBCDEAFFFFFFB
ABCABCABABCBCDEAFFFFFFBCABABCABCABABCBCDEAF
FFFFFFBCABABCBCDEAFFFFFFFEFFFFFF
    
```

Source: Shuvra Bhattacharyya

Lee, Berkeley 59

SDF, by itself, is too restrictive.

Extensions improve expressiveness.

- Heterochronous Dataflow [Girault, Lee, and Lee, 97]
- Structured Dataflow [Kodosky 86, Thies et al. 02]
- (the other) Synchronous Dataflow [Halbwachs et al. 91]
- Cyclostatic Dataflow [Lauwereins 94]
- Multidimensional SDF [Lee & Murthy 96]
- Parameterized Dataflow [Bhattacharya et al. 00]
- Teleport Messages [Thies et al. 05]

All of these remain decidable


And there are many other non-dataflow actor-oriented MoCs to bring into the mix!

Lee, Berkeley 60



Work to be done

- Develop language support for actor-oriented design (like what C++ did for object-oriented design).
- Generalize parallel schedulers to work for more expressive MoCs.
- Support mixing MoCs to enable to exploiting static analysis where possible (this is partially done in Ptolemy II, but much work is left).
- Develop and support design patterns that expose parallelism in actor-oriented designs (e.g. gather-scatter, MapReduce, etc.)



Lee, Berkeley 61



Conclusion:

Disciplined Concurrent Models of Computation

- Do not use nondeterministic programming models to accomplish deterministic ends.
- Use concurrency models that have analogies in the physical world (actors, not threads).
- Provide these in the form of models of computation (MoCs) with well-developed semantics and tools.
- Use specialized MoCs to exploit semantic properties (avoid excess generality).
- Leave the choice of shared memory or message passing to the compiler.



Lee, Berkeley 62



Example of an SDF Program

Thanks to Rodric Rabbah, IBM Watson Center, for suggesting this example.

- 3-D gravitational simulation of n bodies

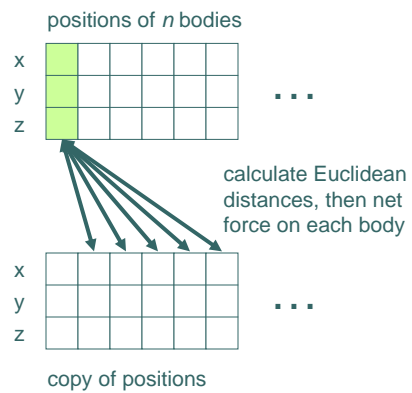



Lee, Berkeley 63



Example of an SDF Program

- 3-D gravitational simulation of n bodies




Lee, Berkeley 64



Example of an SDF Program

- 3-D gravitational simulation of n bodies




Lee, Berkeley 65



Example of an SDF Program

- 3-D gravitational simulation of n bodies

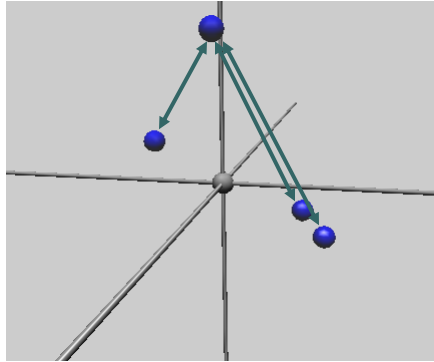


Lee, Berkeley 66



Example of an SDF Program

- 3-D gravitational simulation of n bodies



$$F(t) = ma(t)$$

$$a(t) = F(t)/m$$

$$v(t) = \int_0^t a(\tau) d\tau + v(0)$$

$$p(t) = \int_0^t v(\tau) d\tau + p(0)$$

A simple (naïve) approximation:

$$v(t + \Delta) = v(t) + \Delta a(t)$$

$$p(t + \Delta) = p(t) + \Delta v(t)$$

Lee, Berkeley 67

