



Using the Principles of Synchronous Languages in Discrete-event and Continuous-time Models

Edward A. Lee

Robert S. Pepper Distinguished Professor
Chair of EECS
UC Berkeley

With special thanks to Xiaojun Liu, Eleftherios Matsikoudis, Haiyang Zheng, and Ye (Rachel) Zhou

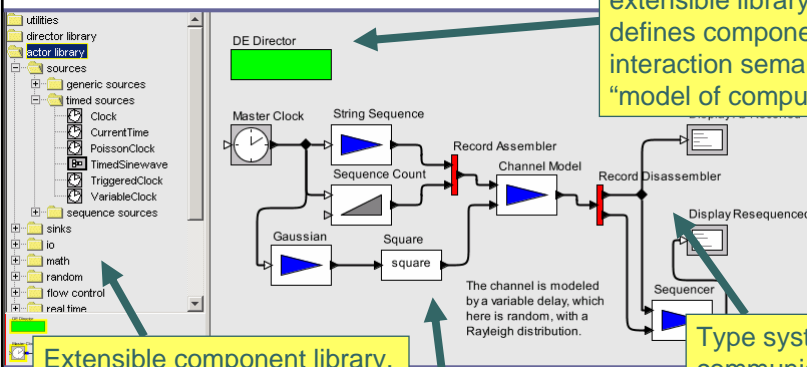
Workshop: Between Control and Software (in honor of Paul Caspi)
September 28, 2007
Grenoble, France



Ptolemy II: Our Laboratory for Studying Concurrent Models of Computation

Concurrency management supporting dynamic model structure.

Director from an extensible library defines component interaction semantics or "model of computation."



Extensible component library.

Visual editor for defining models

Type system for communicated data

Lee, Berkeley 2



Some Models of Computation Implemented in Ptolemy II

- o CSP – concurrent threads with rendezvous
- o CT – continuous-time modeling
- o DE – discrete-event systems
- o DDE – distributed discrete events
- o DDF – dynamic dataflow
- o DPN – distributed process networks
- o DT – discrete time (cycle driven)
- o FSM – finite state machines
- o Giotto – synchronous periodic
- o GR – 2-D and 3-D graphics
- o PN – process networks
- o SDF – synchronous dataflow
- o SR – synchronous/reactive
- o TM – timed multitasking

This talk will start with SR...

... and then show how CT and DE can be built on it, and FSM can be brought in to model hybrid systems.


Lee, Berkeley 3




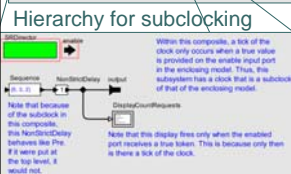
Synchronous/Reactive (SR) Models in Ptolemy II

SR Director

This model illustrates the use of SR primitive actors to make a Countdown actor. This (composite) actor outputs a true on the ready port when it is ready to count. In the same tick of the clock, the Sequence actor provides it with a starting number. It then counts down to zero on each subsequent tick of the clock, emitting true on ready when it again reaches zero.



- o Synchronous coordination language
- o With a few specialized actors
 - When
 - Default
- o Clocking
 - Global clock with subclocks
 - Structured multiclock
 - Structured nondeterminism



Generic actors defined in

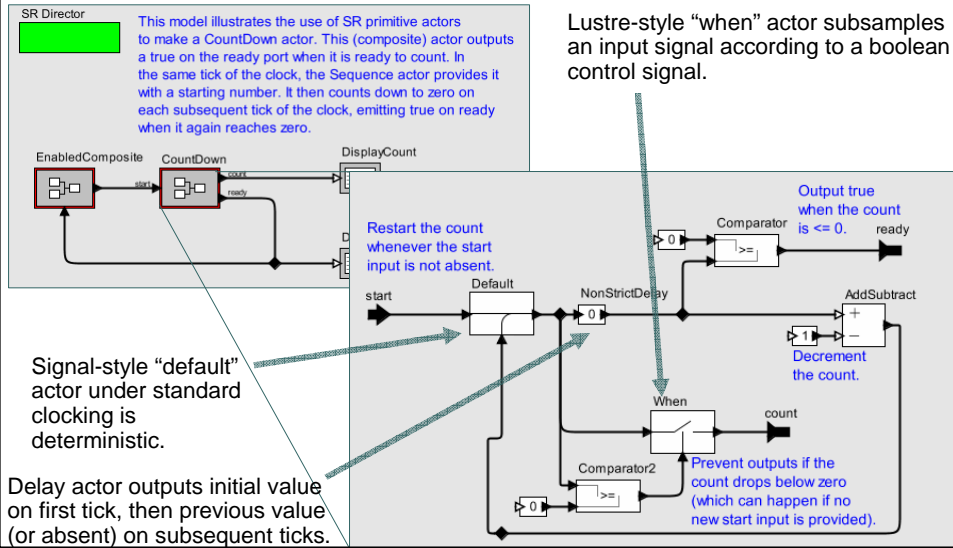
```
Java,  
C,  
MATLAB,  
Cal,  
Python, etc.
```

```
1 // Base SR actor that takes one tick every clock tick and returns the  
2 // value of the counter when the counter reaches the given start  
3 // value. An actor that takes a start value and a tick every clock tick.  
4 // Description: This actor counts down from the start value to zero.  
5 // Inputs: start (int), tick (boolean)  
6 // Outputs: ready (boolean)  
7  
8 // Note: This actor is a subclock of the enclosing model.  
9  
10 // Note: This actor is a subclock of the enclosing model.  
11  
12 // Note: This actor is a subclock of the enclosing model.  
13  
14 // Note: This actor is a subclock of the enclosing model.  
15  
16 // Note: This actor is a subclock of the enclosing model.  
17  
18 // Note: This actor is a subclock of the enclosing model.  
19  
20 // Note: This actor is a subclock of the enclosing model.  
21  
22 // Note: This actor is a subclock of the enclosing model.  
23  
24 // Note: This actor is a subclock of the enclosing model.  
25  
26 // Note: This actor is a subclock of the enclosing model.  
27  
28 // Note: This actor is a subclock of the enclosing model.  
29  
30 // Note: This actor is a subclock of the enclosing model.  
31  
32 // Note: This actor is a subclock of the enclosing model.  
33  
34 // Note: This actor is a subclock of the enclosing model.  
35  
36 // Note: This actor is a subclock of the enclosing model.  
37  
38 // Note: This actor is a subclock of the enclosing model.  
39  
40 // Note: This actor is a subclock of the enclosing model.  
41  
42 // Note: This actor is a subclock of the enclosing model.  
43  
44 // Note: This actor is a subclock of the enclosing model.  
45  
46 // Note: This actor is a subclock of the enclosing model.  
47  
48 // Note: This actor is a subclock of the enclosing model.  
49  
50 // Note: This actor is a subclock of the enclosing model.  
51  
52 // Note: This actor is a subclock of the enclosing model.  
53  
54 // Note: This actor is a subclock of the enclosing model.  
55  
56 // Note: This actor is a subclock of the enclosing model.  
57  
58 // Note: This actor is a subclock of the enclosing model.  
59  
60 // Note: This actor is a subclock of the enclosing model.  
61  
62 // Note: This actor is a subclock of the enclosing model.  
63  
64 // Note: This actor is a subclock of the enclosing model.  
65  
66 // Note: This actor is a subclock of the enclosing model.  
67  
68 // Note: This actor is a subclock of the enclosing model.  
69  
70 // Note: This actor is a subclock of the enclosing model.  
71  
72 // Note: This actor is a subclock of the enclosing model.  
73  
74 // Note: This actor is a subclock of the enclosing model.  
75  
76 // Note: This actor is a subclock of the enclosing model.  
77  
78 // Note: This actor is a subclock of the enclosing model.  
79  
80 // Note: This actor is a subclock of the enclosing model.  
81  
82 // Note: This actor is a subclock of the enclosing model.  
83  
84 // Note: This actor is a subclock of the enclosing model.  
85  
86 // Note: This actor is a subclock of the enclosing model.  
87  
88 // Note: This actor is a subclock of the enclosing model.  
89  
90 // Note: This actor is a subclock of the enclosing model.  
91  
92 // Note: This actor is a subclock of the enclosing model.  
93  
94 // Note: This actor is a subclock of the enclosing model.  
95  
96 // Note: This actor is a subclock of the enclosing model.  
97  
98 // Note: This actor is a subclock of the enclosing model.  
99  
100 // Note: This actor is a subclock of the enclosing model.
```

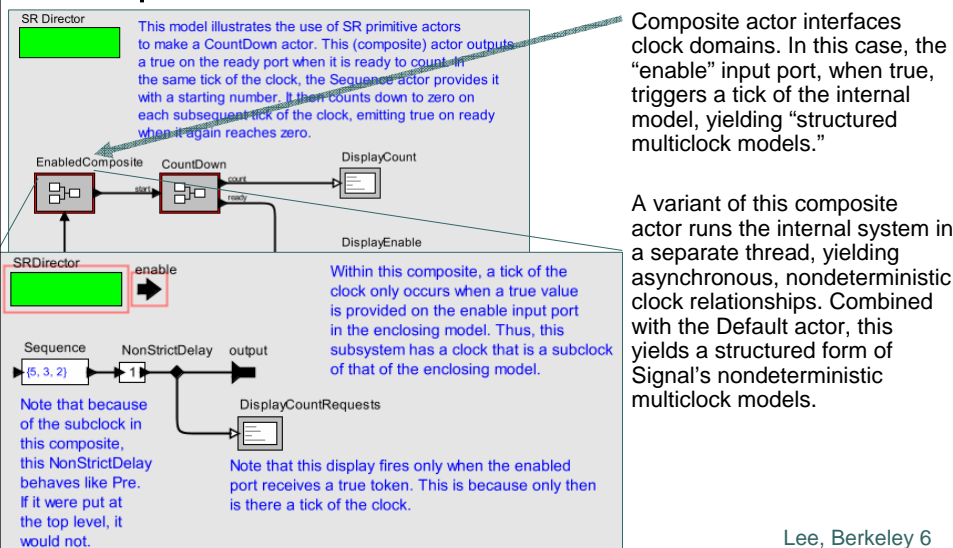
Lee, Berkeley 4



Synchronous/Reactive (SR) Models in Ptolemy II



Synchronous/Reactive (SR) Models in Ptolemy II



Lee, Berkeley 6



How Does This Work? Execution of Ptolemy II Actors


Flow of control:

- Initialization
- Execution
- Finalization

Lee, Berkeley 7



How Does This Work? Execution of Ptolemy II Actors



Flow of control:


- **Initialization**
- Execution
- Finalization

Initialize actors

Lee, Berkeley 8




How Does This Work? Execution of Ptolemy II Actors



Flow of control:

- Initialization
- Execution
- Finalization




Only the `postfire()` method can change the state of the actor.

Lee, Berkeley 9



How Does This Work? Execution of Ptolemy II Actors



Flow of control:

- Initialization
- Execution
- Finalization

Lee, Berkeley 10



Definition of the NonStrictDelay Actor (Sketch)



```
public class NonStrictDelay extends TypedAtomicActor {
    protected Token _previousToken;
    public Parameter initialValue;

    public void initialize() {
        _previousToken = initialValue.getToken();
    }

    public boolean prefire() {
        return true;
    }


    public void fire() {
        if (_previousToken != null) {
            if (_previousToken == AbsentToken.ABSENT) {
                output.sendClear(0);
            } else {
                output.send(0, _previousToken);
            }
        } else {
            output.sendClear(0);
        }
    }

    public boolean postfire() {
        if (input.isKnown(0)) {
            if (input.hasToken(0)) {
                _previousToken = input.get(0);
            } else {
                _previousToken = AbsentToken.ABSENT;
            }
        }
        return true;
    }
}
```

Lee, Berkeley 11



Definition of the NonStrictDelay Actor (Sketch)



initialization

```
public class NonStrictDelay extends TypedAtomicActor {
    protected Token _previousToken;
    public Parameter initialValue;

    public void initialize() {
        _previousToken = initialValue.getToken();
    }

    public boolean prefire() {
        return true;
    }

    public void fire() {
        if (_previousToken != null) {
            if (_previousToken == AbsentToken.ABSENT) {
                output.sendClear(0);
            } else {
                output.send(0, _previousToken);
            }
        } else {
            output.sendClear(0);
        }
    }

    public boolean postfire() {
        if (input.isKnown(0)) {
            if (input.hasToken(0)) {
                _previousToken = input.get(0);
            } else {
                _previousToken = AbsentToken.ABSENT;
            }
        }
        return true;
    }
}
```

Lee, Berkeley 12



Definition of the NonStrictDelay Actor (Sketch)



prefire: can the actor fire?

```
public class NonStrictDelay extends TypedAtomicActor {
    protected Token _previousToken;
    public Parameter initialValue;

    public void initialize() {
        _previousToken = initialValue.getToken();
    }

    public boolean prefire() {
        return true;
    }


    public void fire() {
        if (_previousToken == AbsentToken.ABSENT) {
            output.sendClear(0);
        } else {
            output.send(0, _previousToken);
        }
    }

    public boolean postfire() {
        if (input.isKnown(0)) {
            if (input.hasToken(0)) {
                _previousToken = input.get(0);
            } else {
                _previousToken = AbsentToken.ABSENT;
            }
        }
        return true;
    }
}
```

Lee, Berkeley 13



Definition of the NonStrictDelay Actor (Sketch)



fire: produce outputs (in this case, the output does not depend on the input).

```
public class NonStrictDelay extends TypedAtomicActor {
    protected Token _previousToken;
    public Parameter initialValue;

    public void initialize() {
        _previousToken = initialValue.getToken();
    }

    public boolean prefire() {
        return true;
    }

    public void fire() {
        if (_previousToken != null) {
            if (_previousToken == AbsentToken.ABSENT) {
                output.sendClear(0);
            } else {
                output.send(0, _previousToken);
            }
        } else {
            output.sendClear(0);
        }
    }

    public boolean postfire() {
        if (input.isKnown(0)) {
            if (input.hasToken(0)) {
                _previousToken = input.get(0);
            } else {
                _previousToken = AbsentToken.ABSENT;
            }
        }
        return true;
    }
}
```

Lee, Berkeley 14



Definition of the NonStrictDelay Actor (Sketch)



```

public class NonStrictDelay extends TypedAtomicActor {
    protected Token _previousToken;
    public Parameter initialValue;

    public void initialize() {
        _previousToken = initialValue.getToken();
    }

    public boolean prefire() {
        return true;
    }

    public void fire() {
        if (input.isKnown(0)) {
            if (input.hasToken(0)) {
                _previousToken = input.get(0);
            } else {
                _previousToken = AbsentToken.ABSENT;
            }
        }
    }

    public boolean postfire() {
        if (input.isKnown(0)) {
            if (input.hasToken(0)) {
                _previousToken = input.get(0);
            } else {
                _previousToken = AbsentToken.ABSENT;
            }
        }
        return true;
    }
}

```

postfire:
record
state
changes

Lee, Berkeley 15



Standard Synchronous Semantics

Let V be a family of values (a data type, or alphabet). Let

$$V_{\perp} = V \cup \{\varepsilon, \perp\}$$

be the set of values plus “absent” (ε) and “unknown” (\perp). Define a flat partial order $<$ where $\perp < \varepsilon$ and $\perp < v$ for all $v \in V$. At each tick, every actor realizes a monotonic firing function (in this order). The signal values at the tick are the least fixed point of the composition of these firing functions.

Lee, Berkeley 16



Ptolemy II SR Domain has a Constructive Version of the Synchronous Semantics

Let

$$V_\epsilon = V \cup \{\epsilon\}$$

(No “unknown” and no partial order). Let \mathbb{N} be the non-negative integers. Let s be a signal, given as a partial function:

$$s: \mathbb{N} \rightarrow V_\epsilon$$

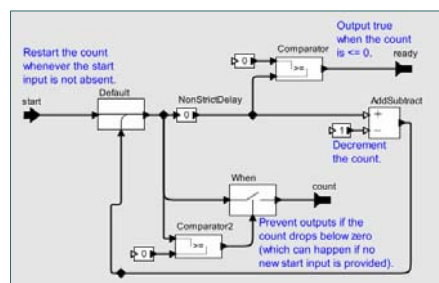
defined on an initial segment of \mathbb{N} . An actor is a function mapping input signals into output signals. This function is required to be monotonic in a prefix order. The signals in a model are the least fixed point of the composition of these actor functions.

Lee, Berkeley 17



Execution of an SR Model (Conceptually)

- Start with all signals empty (i.e. defined on the empty initial segment).
- Initialize all actors.
- Invoke the following on all actors until either all signals are defined on the initial segment $\{0\}$ or no progress can be made:
if (prefire()) { fire(); }
- If not all signals are defined on $\{0\}$, declare a causality loop.
- Invoke postfire() for all actors.
- Repeat to define signals on the initial segment $\{0, 1\}$.
- Etc.



The correctness of this is guaranteed by the fixed point semantics. Efficiency, of course, depends on being smart about the order in which actors are invoked.

Lee, Berkeley 18



Metric Time in SR

- By default, “time” does not advance when executing an SR model in Ptolemy II (“current time” remains at 0.0, a real number).
- Optionally, the SR Director can increment time by a fixed amount on each clock tick.
- More interestingly, SR can be embedded within timed MoCs that model the environment and govern the passage of time.

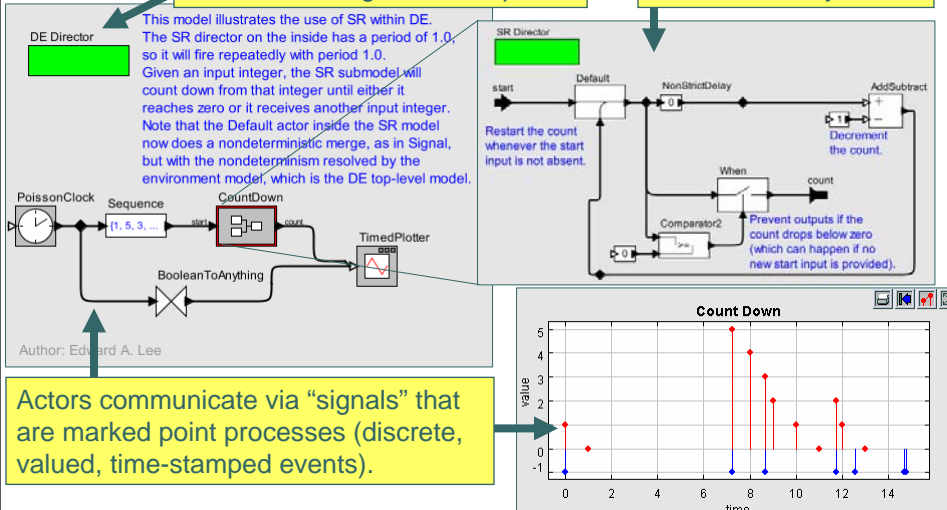
Lee, Berkeley 19



Discrete Events (DE): A Timed Concurrent Model of Computation

DE Director implements timed semantics using an event queue

SR subsystem implements structured nondeterminacy.





Our Applications of DE

- Modeling and simulation of
 - Communication networks (mostly wireless)
 - Hardware architectures
 - Systems of systems
 - The environment of a system
- Design and software synthesis for
 - Sensor networks
 - Distributed real-time software
 - Hardware/software systems

Lee, Berkeley 21



First Attempt at a Model for Signals

Let \mathbb{R}_+ be the non-negative real numbers. Let V be an arbitrary family of values (a data type, or alphabet). Let

$$V_\varepsilon = V \cup \{\varepsilon\}$$

be the set of values plus “absent.” Let s be a signal, given as a partial function:

$$s: \mathbb{R}_+ \rightarrow V_\varepsilon$$

defined on an initial segment of \mathbb{R}_+


Note similarity to SR semantics. Only the domain of signals has changed.

Lee, Berkeley 22



First Attempt at a Model for Signals

$$s: \mathbb{R}_+ \rightarrow V_\varepsilon$$



Initial segment $I \subseteq \mathbb{R}_+$ where the signal is defined.


Absent: $s(\tau) = \varepsilon$ for almost all $\tau \in I$.

This model is not rich enough because it does not allow a signal to have multiple events at the same time.

Lee, Berkeley 23



Example Motivating the Need for Simultaneous Events Within a Signal



Newton's Cradle:

- Steel balls on strings
- Collisions are events
- Momentum of the middle ball has three values at the time of collision.

This example has continuous dynamics as well
(I will return to this)

Other examples:

- Batch arrivals at a queue.
- Software sequences abstracted as instantaneous.
- Transient states.

Lee, Berkeley 24



A Better Model for Signals: Super-Dense Time

Let $T = \mathbb{R}_+ \times \mathbb{N}$ be a set of “tags” where \mathbb{N} is the natural numbers, and give a signal s as a partial function:

$$s: T \rightarrow V_\varepsilon$$

defined on an initial segment of T , assuming a lexical ordering on T :


$$(t_1, n_1) \leq (t_2, n_2) \iff t_1 < t_2, \text{ or } t_1 = t_2 \text{ and } n_1 \leq n_2.$$

This allows signals to have a sequence of values at any real time t .

Lee, Berkeley 25



Super Dense Time



Initial segment $I \subseteq \mathbb{R}_+ \times \mathbb{N}$ where the signal is defined

Absent: $s(\tau) = \varepsilon$ for almost all $\tau \in I$.

Lee, Berkeley 26



Events and Firings

$$s: T \rightarrow V_\varepsilon$$

- A *tag* is a time-index pair, $\tau = (t, n) \in T = \mathbb{R}_+ \times \mathbb{N}$.
- An *event* is a tag-value pair, $e = (\tau, v) \in T \times V$.
- $s(\tau)$ is an event if $s(\tau) \neq \varepsilon$.


Operationally, events are processed by presenting all input events at a tag to an actor and then *firing* it.

However, this is not always possible!

Lee, Berkeley 27



A Feedback Design Pattern



In this model, a sensor produces measurements that are combined with previous measurements using an exponential forgetting function.

The feedback loop makes it impossible to present the Register actor with all its inputs at any tag before firing it.

Lee, Berkeley 28



Solving Feedback Loops

Possible solutions:

- Find algebraic solution
- All actors have time delay
- Some actors have time delay, and every directed loop must have an actor with time delay.
- All actors have delta delay
- Some actors have delta delay and every directed loop must have an actor with delta delay.

Although each of these solutions is used, all are problematic.

The root of the problem is simultaneous events.

Lee, Berkeley 29




Consider “Find Algebraic Solution”

This solution is used by Simulink, but is ill posed.


Consider:

$$y(t) = x^2(t) + u(t)$$
$$x(t) = K y(t)$$



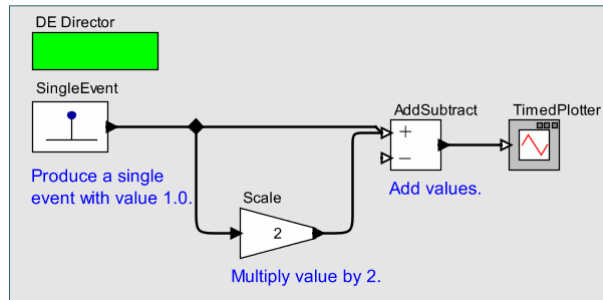
This has two solutions:

$$y(t) = 1.072, \quad x(t) = 0.268, \text{ or}$$
$$y(t) = 14.9282, \quad x(t) = 3.7321$$





Consider “All Actors Have Time Delay”



If all actors have time delay, this produces either:

- Event with value 1 followed by event with value 2, or
- Event with value 1 followed by event with value 3.

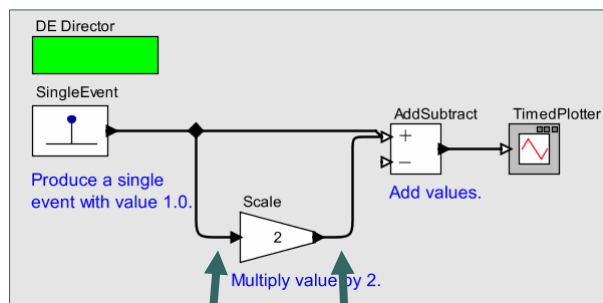
(the latter if signal values are persistent).

Neither of these is likely what we want.

Lee, Berkeley 31



Consider “All Actors Have Delta Delay”




With delta delays, if an input event is $((t, n), v)$, the corresponding output event is $((t, n+1), v')$. Every actor is assumed to give a delta delay.

This style of solution is used in VHDL.

Lee, Berkeley 32



Consider “All Actors Have Delta Delay”




- If all actors have a delta delay, this produces either:
- Event with value 1 followed by event with value 2, or
 - Event with value 1 followed by event with value 3 (the latter if signal values are persistent, as in VHDL).
- Again, neither of these is likely what we want.*

Lee, Berkeley 33



More Fundamental Problem: Delta Delay Semantics is Not Compositional



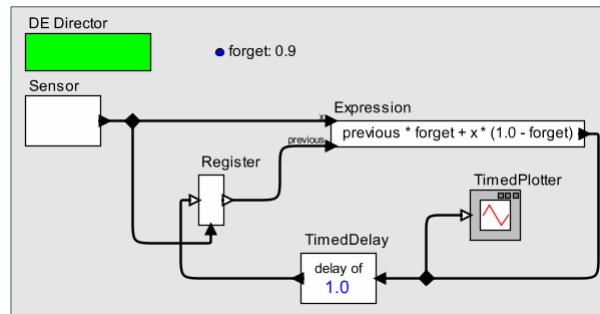
The top composition of two actors will have a two delta delays, whereas the bottom abstraction has only a single delta delay.

Under delta delay semantics, a composition of two actors cannot have the semantics of a single actor.

Lee, Berkeley 34



Consider “Some actors have time delay, and every directed loop must have an actor with time delay.”

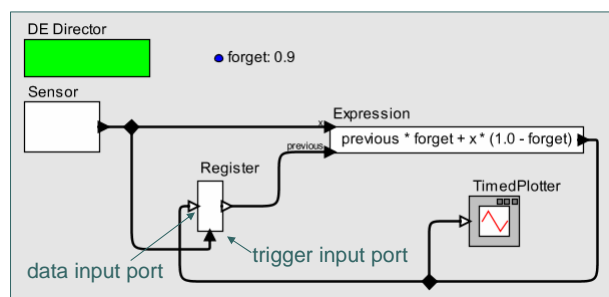


Any non-zero time delay imposes an upper bound on the rate at which sensor data can be accepted. Exceeding this rate will produce erroneous results.

Lee, Berkeley 35



Consider “Some actors have delta delay, and every directed loop must have an actor with delta delay.”



The output of the Register actor must be at least one index later than the data input, hence this actor has at least a delta delay.


To schedule this, could break the feedback loop at actors with delta delay, then do a topological sort.

Lee, Berkeley 36



Naïve Topological Sort is not Compositional

Breaking loops where an actor has a delta delay and performing a topological sort is not a compositional solution:

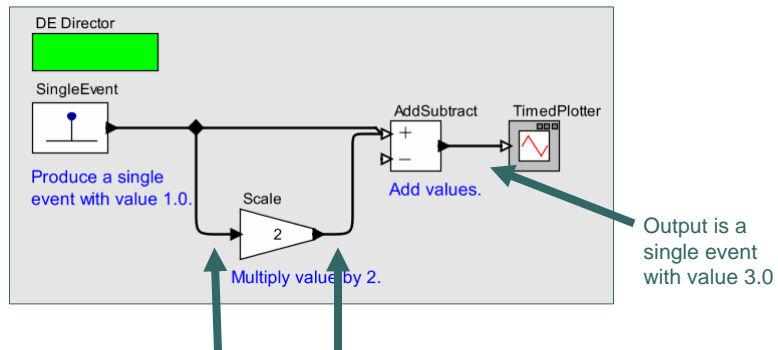


Does this composite actor have a delta delay or not?

Lee, Berkeley 37



Our Answer: No Required Delay, and Feedback Loops Have (Unique) Least Fixed Points Semantics



Given an input event $((t, n), v)$, the corresponding output event is $((t, n), v')$. The actor has no delay.

The semantics is similar to SR, except that time may advance by a variable amount.

Lee, Berkeley 38



Advancing Time

- A signal is a partial function

$$s : T \rightarrow V_{\epsilon}$$

defined on an initial segment of

$$T = \mathbb{R}_+ \times \mathbb{N}$$

- But how to increment the initial segment on which the signal is defined? It won't work to just proceed to the next one, as we did with SR.


Lee, Berkeley 39



Execution of a DE Model (Conceptually)

- Start with all signals empty.
- Initialize all actors (some will post tags on the event queue).
- Take the smallest tag (t, n) from the event queue.
- Invoke the following on all actors that have input events until either all signals are defined on the initial segment $S = [(0,0), (t,n)]$ or no progress can be made:


```
if (prefire()) { fire(); }
```
- If not all signals are defined on S , declare a causality loop.
- Invoke postfire() for all actors (some will post tags on the event queue).
- Repeat with the next smallest tag on the event queue.



This is exactly the execution policy of SR, except that rather than just choosing the next tag in the tag set, we use a sorted event queue to choose an interval over which to increment the initial segment.

Lee, Berkeley 40



Subtle Difference Between SR and DE


- In SR, every actor is fired at every tick of its clock, as determined by a clock calculus and/or structured subclocks.
- In DE, an actor is fired at a tag only if it has input events at that tag or it has previously posted an event on the event queue with that tag.

In DE semantics, event counts may matter. If every actor were to be fired at every tick, then adding an actor in one part of a model could change the behavior in another part of the model in unexpected ways.

Lee, Berkeley 41



Summary: Existence and Uniqueness of the Least Fixed Point Solution.



- Signal: $s: \mathbb{R}_+ \times \mathbb{N} \rightarrow V_e$
- Set of signals: S
- Tuples of signals: $s \in S^N$
- Actor: $F: S^N \rightarrow S^M$


A unique least fixed point, $s \in S^N$ such that $F(s) = s$, exists and be constructively found if S^N is a CPO and F is (Scott) continuous.

Under our execution policy, actors are usually (Scott) continuous.

Lee, Berkeley 42



But: Need to Worry About Liveness: Deadlocked Systems



Existence and uniqueness of a solution is not enough.

The least fixed point of this system consists of empty signals. It is deadlocked!


This is the same as SR causality loops.

Lee, Berkeley 43




Another Liveness Concern is *not* Present in SR: *Zeno Systems*

DE systems may have an infinite number of events in a finite amount of time. These “Zeno systems” can prevent time from advancing.



This model illustrates a Zeno condition, where an infinite number of events occur before time 2.0, and hence the Clock actor is unable to ever produce its output at time 2.0.



In this case, our execution policy fails to implement the Knaster-Tarski constructive procedure because some of the signals are not total.

Lee, Berkeley 44



Liveness

- A signal is *total* if it is defined for all tags in T .
- A model with no inputs is *live* if all signals are total.
- A model with inputs is *live* if all input signals are total implies all signals are total.


Liveness ensures freedom from deadlock and Zeno.

- *Whether a model is live is, in general, undecidable.*
- We have developed a useful sufficient condition based on *causality* that ensures liveness.

Lee, Berkeley 45



Causality Ensures Liveness of an Actor



A monotonic actor F is *causal* if for all sets of input signals S_i , the corresponding set of output signals $S_o = F(S_i)$ satisfy

$$\bigcap_{s \in S_i} \text{dom}(s) \subseteq \bigcap_{s \in S_o} \text{dom}(s).$$

An immediate consequence of this definition is that a causal actor is live. Thus, whether a composition of actors is causal will tell us whether it is live.

Causality does not imply continuity and continuity does not imply causality. Continuity ensures existence and uniqueness of a least fixed point, whereas causality ensures liveness.

Lee, Berkeley 46



Strict Causality Ensures Liveness of a Feedback Composition

A composition of causal actors without directed cycles is itself a causal actor. With cycles, we need:

- A monotonic actor F is *strictly causal* if for all sets of input signals S_i , the corresponding set of output signals $S_o = F(S_i)$ either consists only of total signals (defined over all T) or

$$\bigcap_{s \in S_i} \text{dom}(s) \subset \bigcap_{s \in S_o} \text{dom}(s).$$

(\subset denotes strict subset). If F is a strictly causal actor with one input and one output, then $F(s_{\perp}) \neq s_{\perp}$. F must “come up with something from nothing.”

Lee, Berkeley 47



Continuity, Liveness, and Causality


Theorem: Given a totally ordered tag set and a network of causal and continuous actors where in every dependency loop in the network there is at least one strictly causal actor, then the network is a causal and continuous actor.

This gives us sufficient, but not necessary condition for freedom deadlock and Zeno.

Lee, Berkeley 48



Recall Deadlocked System

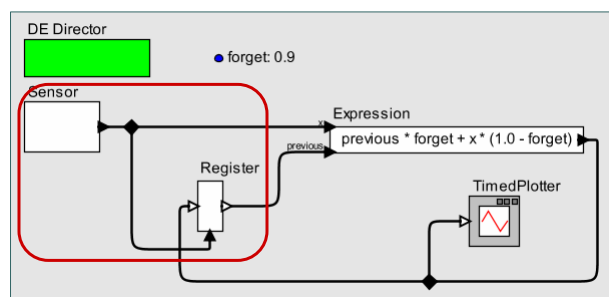


The feedback loop has no strictly causal actor.

Lee, Berkeley 49



Feedback Loop that is Not Deadlocked



This feedback loop also has no strictly causal actor, unless...

We aggregate the two actors as shown into one.


Lee, Berkeley 50



Causality Interfaces Make Scheduling of Execution and Analysis for Liveness Efficient

A causality interface exposes just enough information about an actor to make scheduling and liveness analysis efficient.

An algebra of interfaces enables inference of the causality interface of a composition.



Lee, Berkeley 51



Models of Computation Implemented in Ptolemy II

- o CSP – concurrent threads with rendezvous
- o CT – continuous-time modeling
- o DE – discrete-event systems
- o DDE – distributed discrete events
- o DDF – dynamic dataflow
- o DPN – distributed process networks
- o DT – discrete time (cycle driven)
- o FSM – finite state machines
- o Giotto – synchronous periodic
- o GR – 2-D and 3-D graphics
- o PN – process networks
- o SDF – synchronous dataflow
- o SR – synchronous/reactive
- o TM – timed multitasking

Done

But will also establish connections with Continuous Time (CT) and hybrid systems (CT + FSM)

SR is a special case of DE where the tag set is discrete.

Lee, Berkeley 52



Standard Model for Continuous-Time Signals

In ODEs, the usual formulation of the signals of interest is a function from the time line (a connected subset of the reals) to the reals:

$$p: \mathbb{R}_+ \rightarrow \mathbb{R}^n$$

$$\dot{p}: \mathbb{R}_+ \rightarrow \mathbb{R}^n$$

$$\ddot{p}: \mathbb{R}_+ \rightarrow \mathbb{R}^n$$

Such signals are continuous at $t \in \mathbb{R}_+$ if (e.g.):


$$\forall \epsilon > 0, \exists \delta > 0, \text{ s.t. } \forall \tau \in (t-\delta, t+\delta), \quad \|\dot{p}(t) - \dot{p}(\tau)\| < \epsilon$$

Lee, Berkeley 53



Piecewise Continuous Signals

In hybrid systems of interest, signals have discontinuities.



Piecewise continuous signals are continuous at all $t \in \mathbb{R}_+ \setminus D$ where $D \subset \mathbb{R}_+$ is a *discrete set*.¹


¹A set D with an order relation is a *discrete set* if there exists an order embedding to the integers.

Lee, Berkeley 54



Operational Semantics of Hybrid Systems

A computer execution of a hybrid system is constrained to provide values on a discrete set:

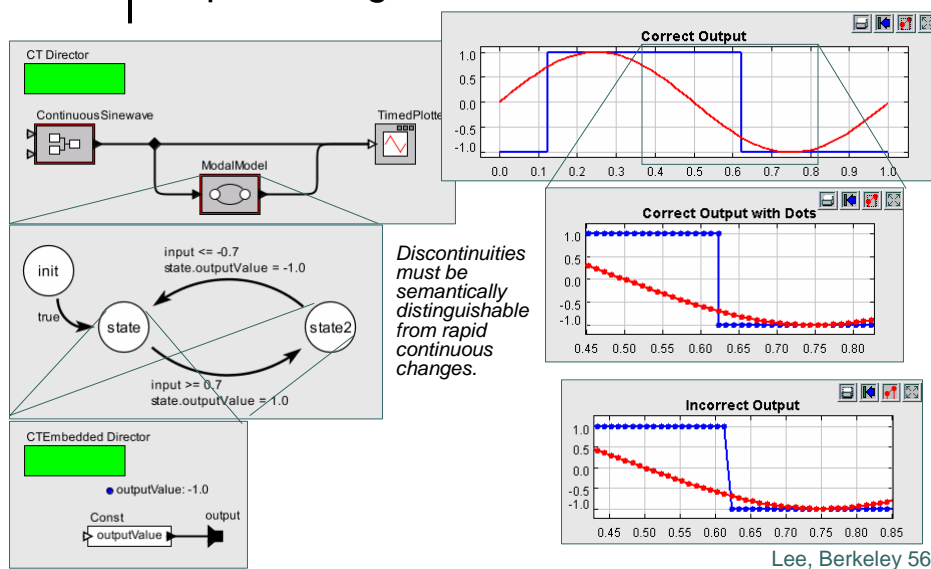


Given this constraint, choosing $T \subset \mathbb{R}$ as the domain of these functions is an unfortunate choice. It makes it impossible to unambiguously represent discontinuities.

Lee, Berkeley 55



Discontinuities Are Not Just Rapid Changes






Solution is the Same: Superdense Time

$$p: \mathbb{R}_+ \times \mathbb{N} \rightarrow \mathbb{R}^n$$

$$\dot{p}: \mathbb{R}_+ \times \mathbb{N} \rightarrow \mathbb{R}^n$$

$$\ddot{p}: \mathbb{R}_+ \times \mathbb{N} \rightarrow \mathbb{R}^n$$




This makes it quite easy to construct models that combine continuous dynamics with discrete dynamics.

Lee, Berkeley 57




Ideal Solver Semantics

[Liu and Lee, HSCC 2003]



In the *ideal solver semantics*, an ODE governing the hybrid system has a unique solution for intervals $[t_i, t_{i+1})$, the interval between discrete time points. A discrete trace loses nothing by not representing values within these intervals.

This elaborates our DE models only by requiring that an ODE solver be consulted when advancing time.



Lee, Berkeley 58



Recall Subtle Difference Between SR and DE. CT is more like SR.

- In SR, every actor is fired at every tick of its clock, as determined by a clock calculus and/or structured subclocks.
- In DE, an actor is fired at a tag only if it has input events at that tag or it has previously posted an event on the event queue with that tag.
- In CT, every actor is fired at every tick of the clock, as determined by an ODE solver.

In CT semantics, a signal has a value at every tag. But the solver chooses to explicitly represent those values only at certain tags.

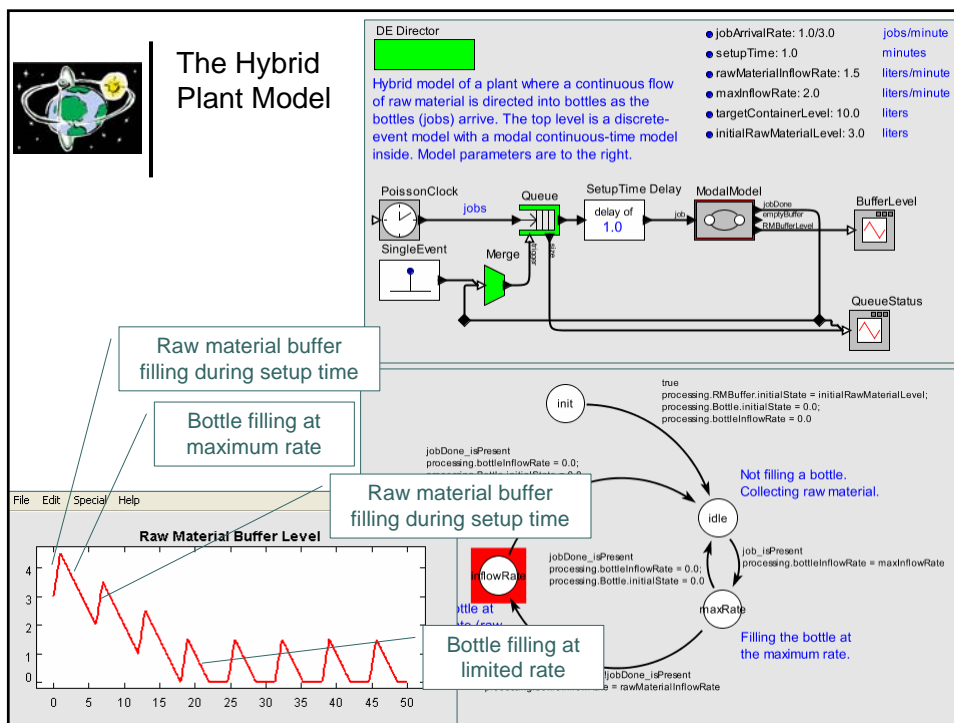
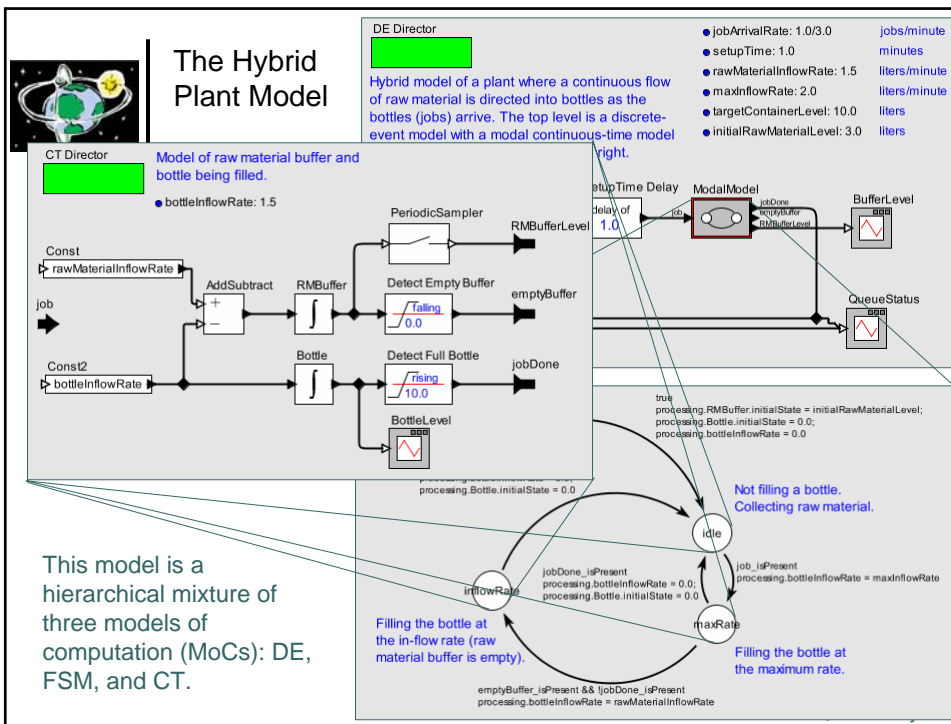
Lee, Berkeley 59



Key Contribution

- *With a common underlying fixed-point semantics, SR, DE, and CT can be arbitrarily combined hierarchically!*
- They can also be combined with FSMs using “modal models,” to get a rich hybrid systems modeling semantics.

Lee, Berkeley 60





Conclusions

A constructive fixed point semantics for synchronous/reactive models generalizes naturally to discrete-event and continuous-time models, enabling arbitrary combinations of the three modeling styles.

Lee, Berkeley 63



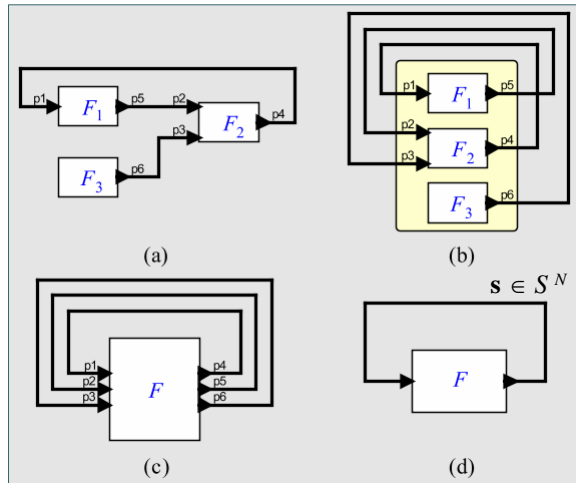
Further Reading

- [1] E. A. Lee and H. Zheng, "Leveraging Synchronous Language Principles for Heterogeneous Modeling and Design of Embedded Systems," in EMSOFT Salzburg, Austria: ACM, 2007.
- [2] X. Liu and E. A. Lee, "CPO Semantics of Timed Interactive Actor Networks," UC Berkeley, Berkeley, CA, Technical Report EECS-2006-67, May 18 2006.
- [3] X. Liu, E. Matsikoudis, and E. A. Lee, "Modeling Timed Concurrent Systems," in CONCUR 2006 - Concurrency Theory, Bonn, Germany, 2006.
- [4] A. Cataldo, E. A. Lee, X. Liu, E. Matsikoudis, and H. Zheng, "A Constructive Fixed-Point Theorem and the Feedback Semantics of Timed Systems," in Workshop on Discrete Event Systems (WODES), Ann Arbor, Michigan, 2006.
- [5] E. A. Lee, "Modeling Concurrent Real-time Processes Using Discrete Events," Annals of Software Engineering, vol. 7, pp. 25-45, March 4th 1998 1999.
- [6] E. A. Lee, H. Zheng, and Y. Zhou, "Causality Interfaces and Compositional Causality Analysis," in *Foundations of Interface Technologies (FIT), Satellite to CONCUR*, San Francisco, CA, 2005.

Lee, Berkeley 64



Observation: Any Composition is a Feedback Composition



- Signal: $s: \mathbb{R}_+ \times \mathbb{N} \rightarrow V_\epsilon$
- Set of signals: S
- Tuples of signals: $s \in S^N$
- Actor: $F: S^N \rightarrow S^M$

If every actor is a function, then the semantics of the overall system is the least $s \in S^N$ such that $F(s) = s$.

We have a least fixed point semantics.

Lee, Berkeley 65



Prefix Order**

- Recall that a signal s is a partial function: defined on an initial segment of T . Such a function can be given by its *graph*, $s \subset T \times V_\epsilon$.
- A signal s_1 is a *prefix* of a signal s_2 if $s_1 \subseteq s_2$. The prefix relation is a partial order on the set S of signals.
- Fact: S with the prefix order is a complete semilattice (and hence also a CPO).
- Generalizes easily to tuples of signals S^N .


Lee, Berkeley 66



Monotonic and Continuous Functions**

A function $F: S \rightarrow S$ is *monotonic* if it is order-preserving,

$$\forall s_1, s_2 \in S, \quad s_1 \subseteq s_2 \implies F(s_1) \subseteq F(s_2).$$



The same function is (Scott) *continuous* if for all directed sets $S' \subseteq S$, $F(S')$ is a directed set and

$$F(\bigvee S') = \bigvee F(S').$$

Here, $F(S')$ is defined in the natural way as $\{F(s) \mid s \in S'\}$, and $\bigvee X$ denotes the least upper bound of the set X .

Every continuous function is monotonic, and behaves as follows:
Extending the input (in time or tags) can only extend the output.

Lee, Berkeley 67




Knaster-Tarski Fixed-Point Theorem**

A classic fixed point theorem states that if F is continuous, then it has a least fixed point, and that least fixed point is

$$\bigvee \{F^n(\perp_S) \mid n \in \mathbb{N}\},$$

where \perp_S is the least element of S (the empty signal) and \mathbb{N} is the natural numbers.


- Start with empty signals.
- Iteratively apply function F .
- Converge to the unique solution.



Lee, Berkeley 68



Semantics of Merge



At time t , input sequences are interleaved. That is, if the inputs are s_1 and s_2 and

$$\begin{aligned} s_1(t, 0) &= v_1, \\ s_2(t, 0) &= w_1, \quad s_1(t, 1) = w_2 \end{aligned}$$


(otherwise absent) then the output s is

$$s(t, 0) = v_1, \quad s(t, 1) = w_1, \quad s(t, 2) = w_2.$$

Lee, Berkeley 69



Implementation of Merge



```
private List pendingEvents;
fire() {
  foreach input s {
    if (s is present) {
      pendingEvents.append(event from s);
    }
  }
  if (pendingEvents has events) {
    send to output (pendingEvents.first);
    pendingEvents.removeFirst();
  }
  if (pendingEvents has events) {
    post event at the next index on the event queue;
  }
}
```

Lee, Berkeley 70