

# Concurrent Models of Computation

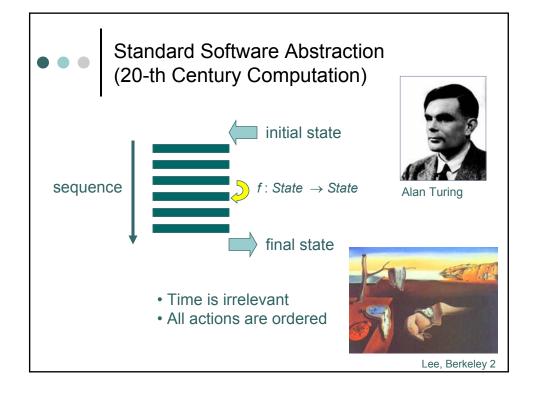
#### Edward A. Lee

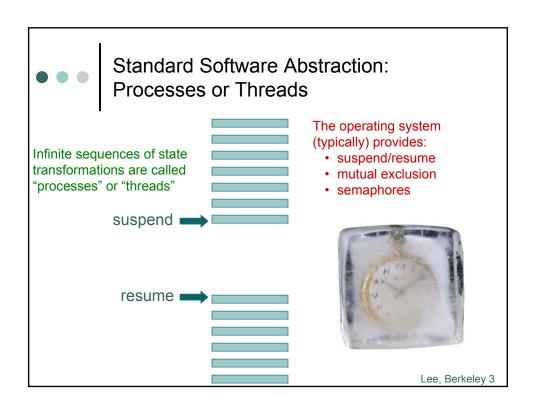
Professor, UC Berkeley Ptolemy Project

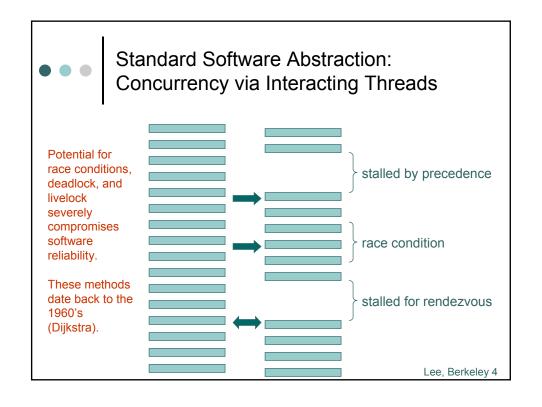
CHESS: Center for Hybrid and Embedded Software Systems

HP Workshop on Advanced Software Technologies July 20-22

HP Labs, Palo Alto, CA







#### A Stake in the Ground

Nontrivial concurrent programs based on processes, threads, semaphores, and mutexes are incomprehensible to humans.

- No amount of process improvement is going to change this.
  - the human brain doesn't work this way.
- Formal methods may help
  - scalability?
  - understandability?
- Better concurrency abstractions will help more





## Ptolemy Project Code Review A Typical Story

- Code review discovers that a method needs to be synchronized to ensure that multiple threads do not reverse each other's actions.
- No problems had been detected in 4 years of using the code.
- Three days after making the change, users started reporting deadlocks caused by the new mutex.
- Analysis and correction of the deadlock is hard.
- But code review successfully identified the flaw.

```
public synchronized void addChangeListener(ChangeListener listener) (
   NamedObj container = (NamedObj) getContainer();
   if (container != null) (
        container.addChangeListener(listener);
   ) else (
   if (_changeListeners == null) (
        __changeListeners = new LinkedList();
        __changeListeners.add(0, listener);
    ) else if (!_changeListeners.contains(listener)) (
        __changeListeners.add(0, listener);
    )
}
Lee, Berkeley 7
```

• • •

# Code Review Doesn't Always Work Another Typical Story

```
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, design reviewed to yellow, then code reviewed to green in 2000, causes a deadlock during a demo on April 26, 2004.

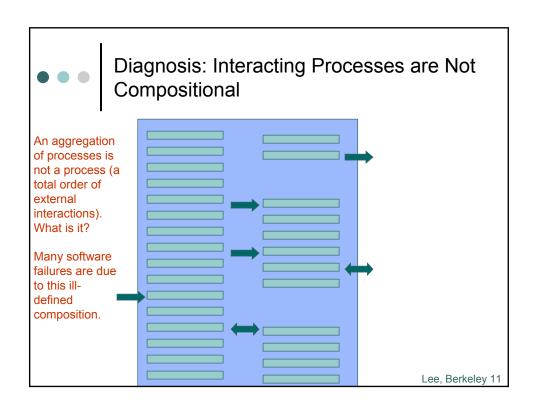
#### And Doubts Remain

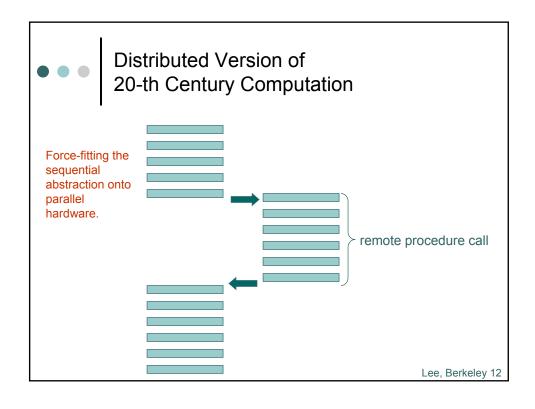
```
CrossRefList is a list that maintains pointers to other CrossRefLists.
@author Geroncio Galicia, Contributor: Edward A. Lee
@version $Id: CrossRefList.java,v 1.78 2004/04/29 14:50:00 eal Exp $
@since Ptolemy II 0.2
@Pt.ProposedRating Green (eal)
@Pt.AcceptedRating Green (bart)
                                                                       Safety of this code
                                                                       depends on policies
public final class CrossRefList implements Serializable {
                                                                       maintained by entirely
    protected class CrossRef implements Serializable{
                                                                       unconnected classes.
                                                                       The language and
        private synchronized void _dissociate() {
             unlink(); // Remove this.
// NOTE: Deadlock risk here! If _far is waiting
                                                                       synchronization
                                                                       mechanisms provide no
             // on a lock to this CrossRef, then we will get
                                                                       way to talk about these
             // deadlock. However, this will only happen if // we have two threads simultaneously modifying a
                                                                       systemwide properties.
             // model. At the moment (4/29/04), we have no // mechanism for doing that without first
             // acquiring write permission the workspace().
             // Two threads cannot simultaneously hold that
             // write access.
             if (_far != null) _far._unlink(); // Remove far
                                                                                    Lee, Berkeley 9
```

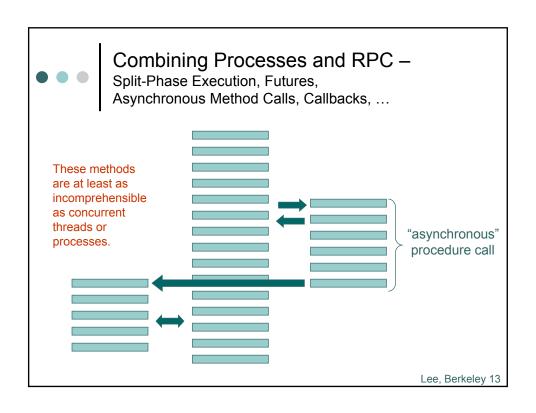
# What it Feels Like to Use the *synchronized* Keyword in Java

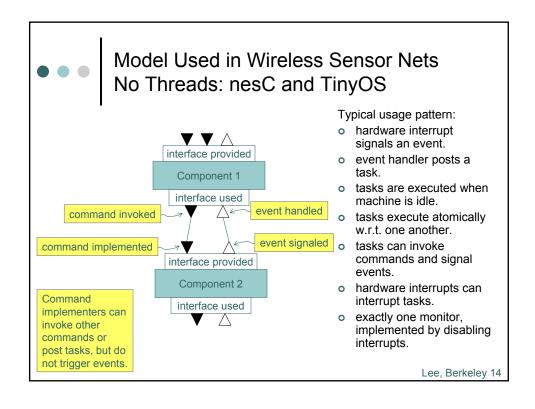


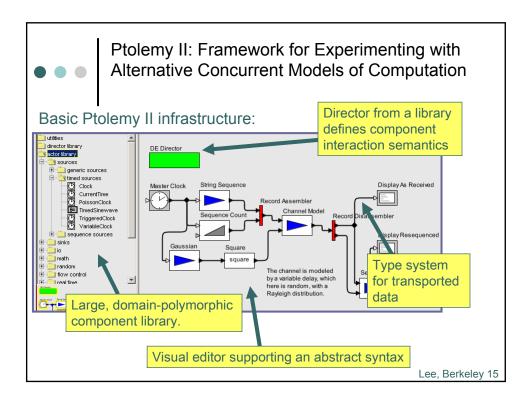
Image "borrowed" from an Lomega advertisement for Y2K software and disk drives, *Scientific America*, September 199

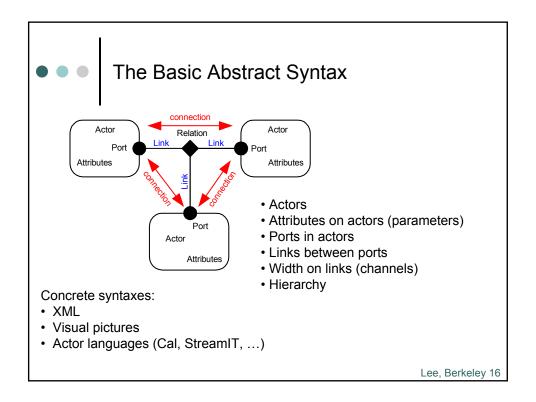


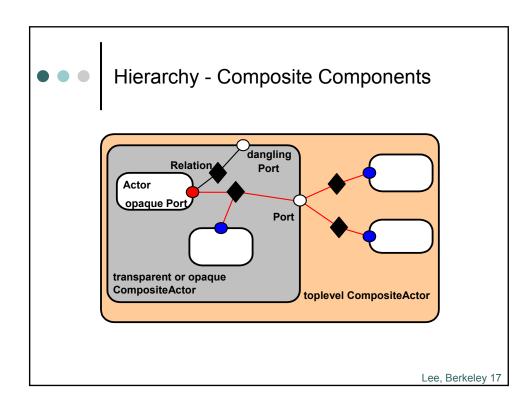


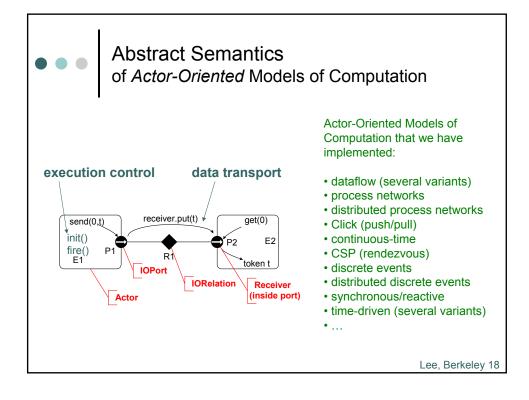


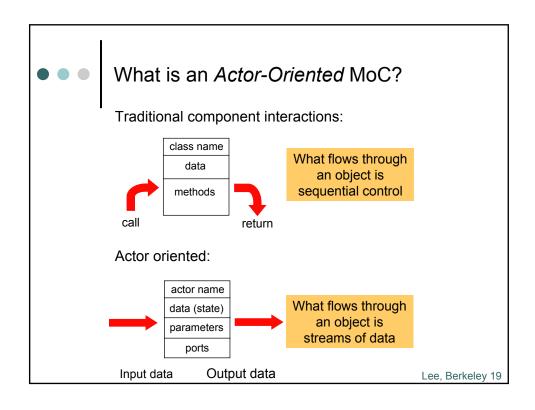




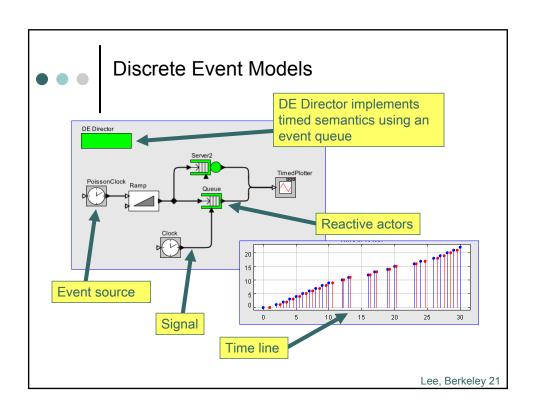


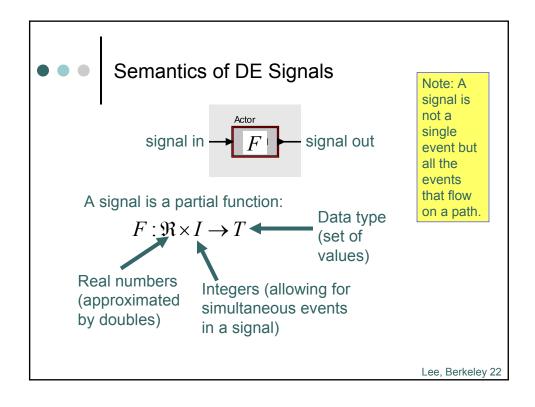


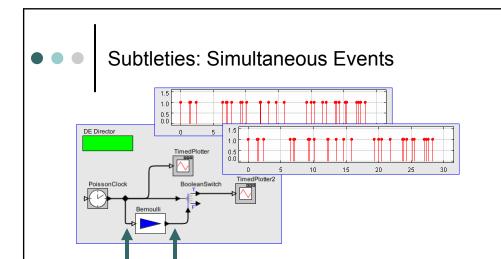




#### Models of Computation Implemented in Ptolemy II CI – Push/pull component interaction Click – Push/pull with method invocation CSP – concurrent threads with rendezvous CT – continuous-time modeling DE – discrete-event systems Most of DDE – distributed discrete events these are FSM – finite state machines actor DT – discrete time (cycle driven) oriented. Giotto – synchronous periodic GR – 2-D and 3-D graphics PN – process networks DPN – distributed process networks SDF – synchronous dataflow SR – synchronous/reactive TM – timed multitasking Lee, Berkeley 20



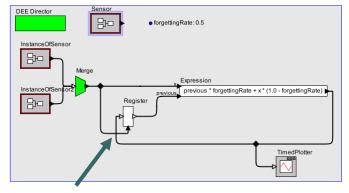




By default, an actor produces events with the same time as the input event. But in this example, we expect (and need) for the BooleanSwitch to "see" the output of the Bernoulli in the same "firing" where it sees the event from the PoissonClock. Events with identical time stamps are also ordered, and reactions to such events follow data precedence order.

Lee, Berkeley 23

## • • Subtleties: Feedback



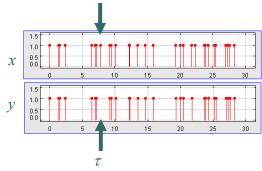
Data precedence analysis has to take into account the non-strictness of this actor (that an output can be produced despite the lack of an input).

### Discrete-Event Semantics

Cantor metric:

$$d(x,y) = 1/2^{\tau}$$

where  $\tau$  is the earliest time where x and y differ.



Lee, Berkeley 25

## • • Causality

 $\boldsymbol{x}$ 

x'

Causal:

$$d(y, y') \le d(x, x')$$

Strictly causal:

$$d(y,y') < d(x,x')$$

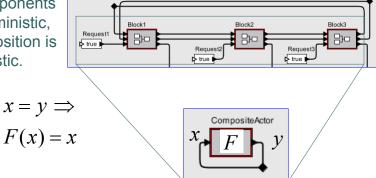
Delta causal:

$$\exists \delta < 1,$$
  
 $d(y, y') \le \delta d(x, x')$ 

A delta-causal component is a "contraction map."

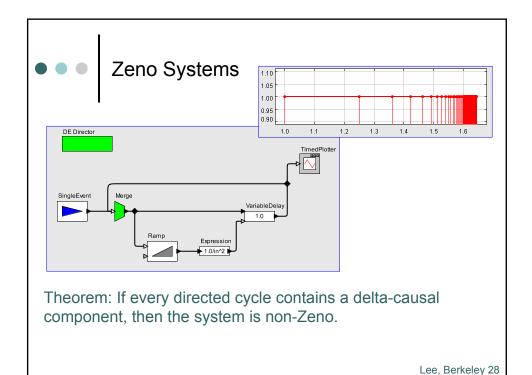
## Semantics of Composition

If the components are deterministic, the composition is deterministic.

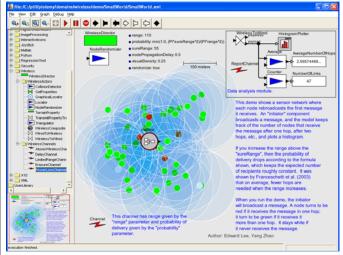


Banach fixed point theorem:

- · Contraction map has a unique fixed point
- · Execution procedure for finding that fixed point
- · Successive approximations to the fixed point

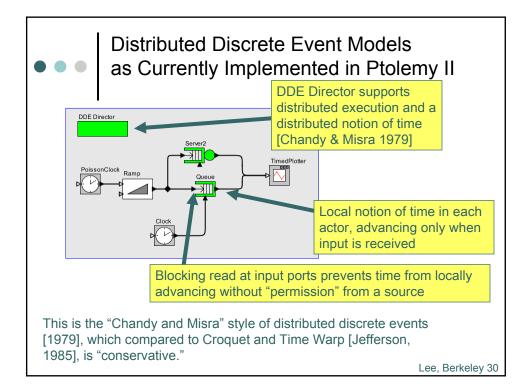


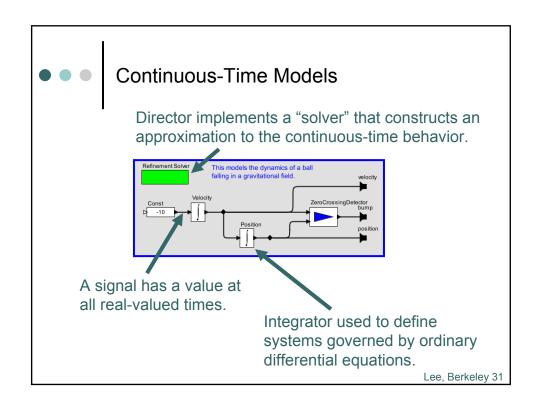
# Extension of Discrete-Event Modeling for Wireless Sensor Nets

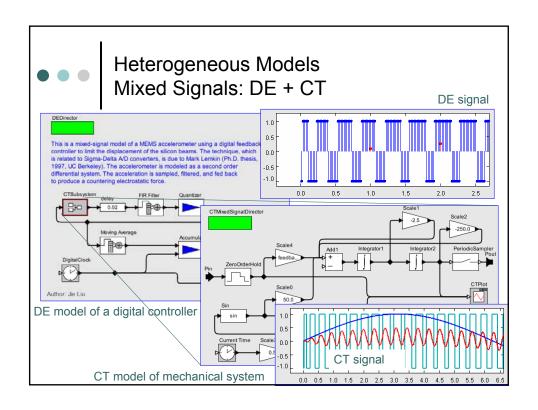


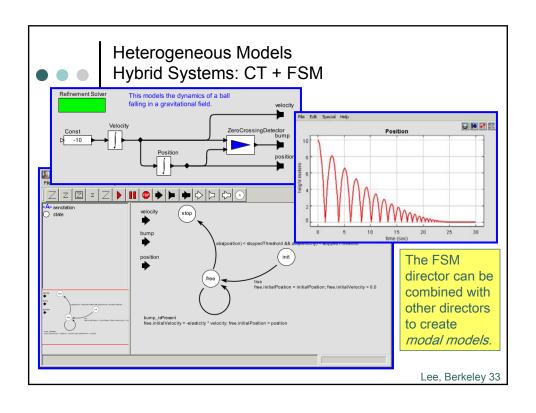
VisualSense extends the Ptolemy II discreteevent domain with communication between actors representing sensor nodes being mediated by a *channel*, which is another actor.

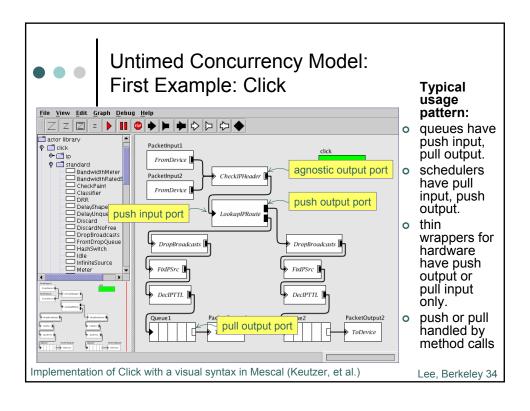
The example at the left shows a grid of nodes that relay messages from an *initiator* (center) via a channel that models a low (but nonzero) probability of long range links being viable.



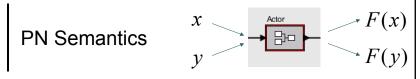








#### **Untimed Concurrency Model:** Second Example: Process Networks This model, whose structure is due to Kahn and MacQueen, calculates integers whose prime factors are only 2, 3, and 5, with no redundancies It uses the OrderedMerge actor, which takes two monotonically increasing input sequences and merges them into one monotonically increasing output sequence. This BooleanSwitch is used SampleDelay to starve the model after · {5} **•** a power of 5 greater than Comparato 1000000 is produced. This results in deterministically Limit on powers of 5 stopping the execution. **▶** 1000000 **▶** actor == thread OrderedMerge SampleDelay3 {3} ▶ signal == stream OrderedMerge2 SampleDelay2 In the PN domain, each actor executes reads block in its own Java thread. That thread iteratively reads inputs, performs computation, and produces outputs Display writes don't Kahn, MacQueen, 1977 The output is an ordered sequence of integers of the form $2^n * 3^m * 5^k$ , where n, m and k are non-negative integers Lee, Berkeley 35



- o A signal is a sequence of values
- o Define a prefix order:

$$x \sqsubset y$$

means that x is a prefix of y.

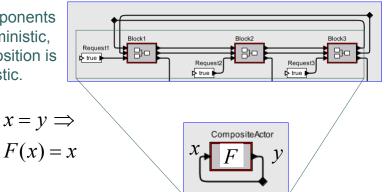
Actors are monotonic functions:

$$x \sqsubseteq y \Rightarrow F(x) \sqsubseteq F(y)$$

 Stronger condition: Actors are continuous functions (intuitively: they don't wait forever to produce outputs).

#### PN Semantics of Composition (Kahn, '74)

If the components are deterministic, the composition is deterministic.

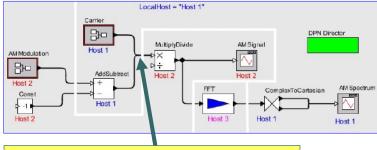


Knaster-Tarski fixed point theorem:

- · Continuous function has a unique least fixed point
- Execution procedure for finding that fixed point
- · Successive approximations to the fixed point

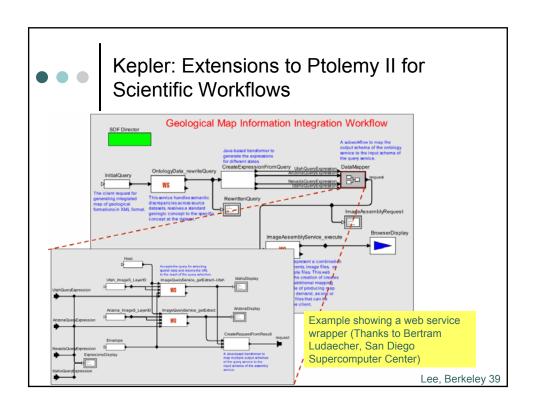
Lee, Berkeley 37

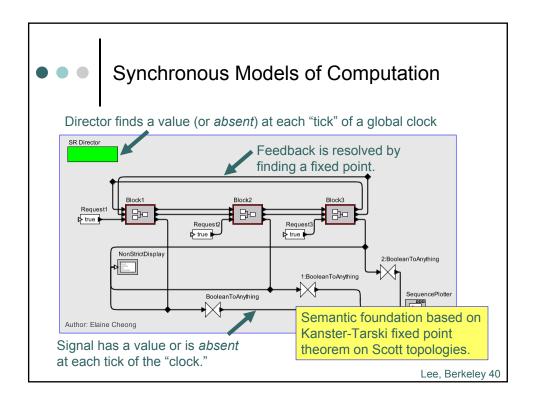
#### Distributed Process Networks



Transport mechanism between hosts is provided by the director. Transparently provides guaranteed delivery and ordered messages.

Created by Dominique Ragot, Thales Communications







# Languages Based on the Synchronous Model of Computation

- Lustre (and SCADE)
- Esterel
- Signal
- Statecharts (and UML state machines)
- Argos
- 0 ...

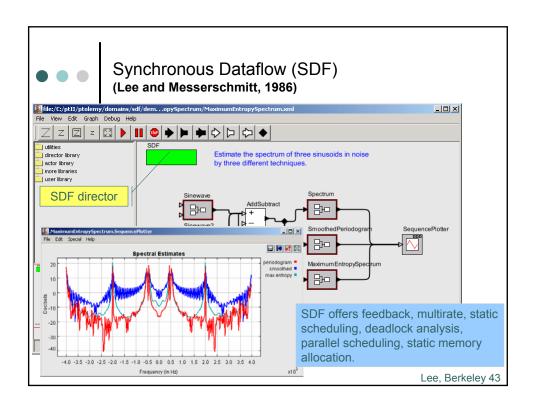
Lee, Berkeley 41

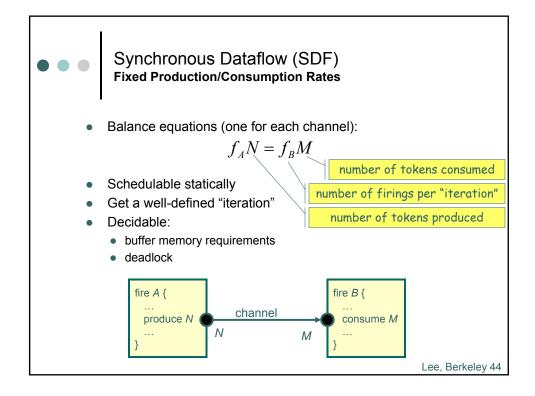
### $\bullet$

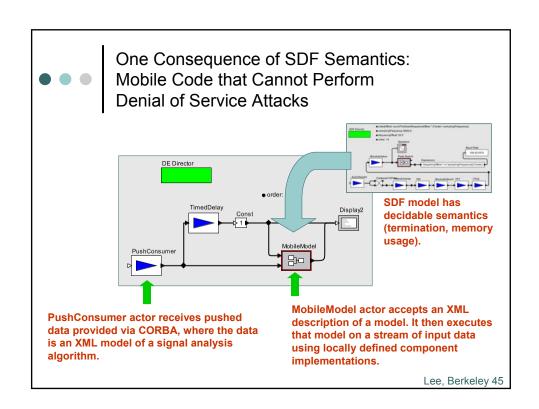
#### **Dataflow Models of Computation**

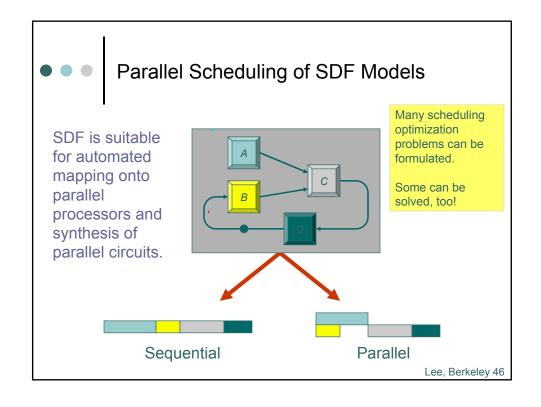
- Computation graphs [Karp & Miller 1966]
- Process networks [Kahn 1974]
- Static dataflow [Dennis 1974]
- o Dynamic dataflow [Arvind, 1981]
- K-bounded loops [Culler, 1986]
- Synchronous dataflow [Lee & Messerschmitt, 1986]
- Structured dataflow [Kodosky, 1986]
- o PGM: Processing Graph Method [Kaplan, 1987]
- Synchronous languages [Lustre, Signal, 1980's]
- Well-behaved dataflow [Gao, 1992]
- Boolean dataflow [Buck and Lee, 1993]
- o 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]
- o ...

Many tools, software frameworks, and hardware architectures have been built to support one or more of these.

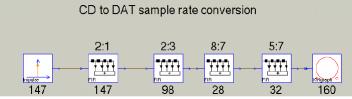








## Scheduling Tradeoffs (Bhattacharyya, Parks, Pino)

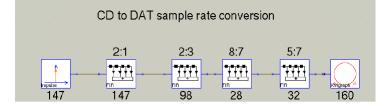


Scheduling strategy	Code	Data
Minimum buffer schedule, no looping	13735	32
Minimum buffer schedule, with looping	9400	32
Worst minimum code size schedule	170	1021
Best minimum code size schedule	170	264

Source: Shuvra Bhattacharyya

Lee, Berkeley 47

#### Minimum Buffer Schedule



Source: Shuvra Bhattacharyya

#### Selected Generalizations

- Multidimensional Synchronous Dataflow (1993)
  - Arcs carry multidimensional streams
  - One balance equation per dimension per arc
- Cyclo-Static Dataflow (Lauwereins, et al., 1994)
  - Periodically varying production/consumption rates
- Boolean & Integer Dataflow (1993/4)
  - Balance equations are solved symbolically
  - Permits data-dependent routing of tokens
  - Heuristic-based scheduling (undecidable)
- Dynamic Dataflow (1981-)
  - Firings scheduled at run time
  - Challenge: maintain bounded memory, deadlock freedom, liveness
  - Demand driven, data driven, and fair policies all fail
- Kahn Process Networks (1974-)
  - Replace discrete firings with process suspension
  - Challenge: maintain bounded memory, deadlock freedom, liveness
- Heterochronous Dataflow (1997)
  - Combines state machines with SDF graphs
  - · Very expressive, yet decidable

Lee, Berkeley 49



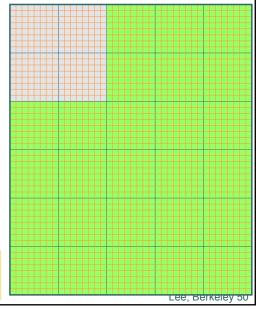
Multidimensional SDF (Lee, 1993)

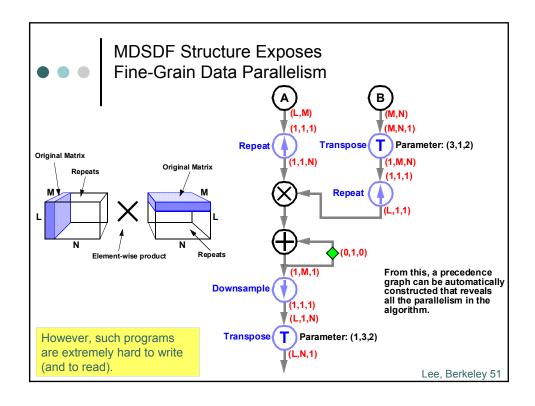
• Production and consumption of *N*-dimensional arrays of data:



- Balance equations and scheduling policies generalize.
- Much more data parallelism is exposed.

Similar (but dynamic) multidimensional streams have been implemented in Lucid.







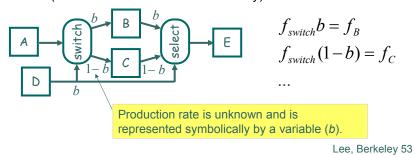
- Actors cycle through a regular production/consumption pattern.
- Balance equations become:

$$f_A \sum_{i=0}^{R-1} N_{i \bmod P} = f_B \sum_{i=0}^{R-1} M_{i \bmod Q}; \ R = lcm(P,Q)$$
 
$$\begin{array}{c} \text{cyclic production pattern} \\ \\ \text{channel} \\ \\ \dots \\ \\ N_0, \dots, N_{P-1} \\ \end{array} \begin{array}{c} \text{fire } B \left\{ \\ \dots \\ \text{consume } M \\ \dots \\ \\ \end{array} \right.$$



## Boolean and Integer Dataflow (BDF, IDF) (Lee and Buck, 1993)

- Balance equations are solved symbolically in terms of unknowns that become known at run time.
- An annotated schedule is constructed with predicates guarding each action.
- Existence of such an annotated schedule is undecidable (as is deadlock & bounded memory)



• • •

#### Dynamic Dataflow (DDF)

- Actors have firing rules
  - Set of finite prefixes on input sequences
  - For determinism: No two such prefixes are joinable under a prefix order
  - Firing function applied to finite prefixes yield finite outputs
- Scheduling objectives:
  - Do not stop if there are executable actors
  - Execute in bounded memory if this is possible
  - Maintain determinacy if possible
- Policies that fail:
  - Data-driven execution

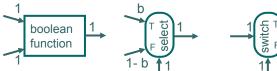
DDF, like BDF and IDF is undecidable (deadlock, bounded memory, schedule)

- Demand-driven executionFair execution
- Many balanced data/demand-driven strategies
- Policy that succeeds (Parks 1995):
  - Execute with bounded buffers
  - Increase bounds only when deadlock occurs

## Undecidability (Buck '93)

- Sufficient set of actors for undecidability:
  - boolean functions on boolean tokens
  - switch and select
  - initial tokens on arcs

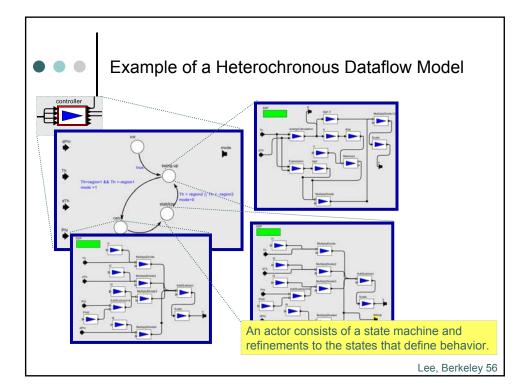
These four parts are sufficient to build any computable function.





- Undecidable:
  - deadlock
  - bounded buffer memory
  - existence of an annotated schedule

BDF, IDF, DDF, and PN are all undecidable in this sense. Fortunately, we can identify a large decidable subset, which we call heterochronous dataflow (HDF).

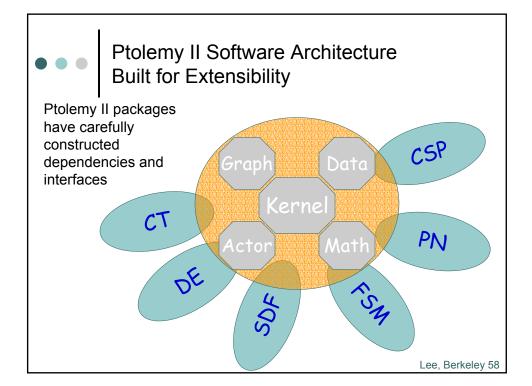




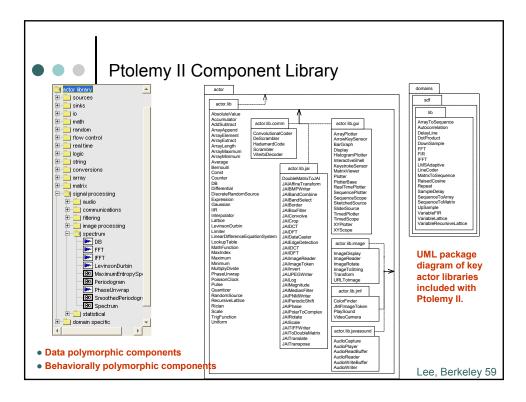
## Heterochronous Dataflow (HDF) (Girault, Lee, and Lee, 1997)

- o An interconnection of actors.
- An actor is either SDF or HDF.
- o If HDF, then the actor has:
  - a state machine
  - a refinement for each state
  - where the refinement is an SDF or HDF actor
- Operational semantics:
  - with the state of each state machine fixed, graph is SDF
  - in the initial state, execute one complete SDF iteration
  - evaluate guards and allow state transitions
  - in the new state, execute one complete SDF iteration
- HDF is decidable
  - but complexity can be high

Lee, Berkeley 57



Related to "parameterized dataflow" of Bhattachrya and Bhattacharyya (2001).



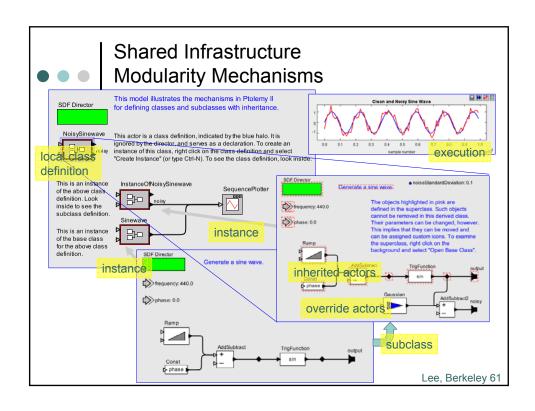


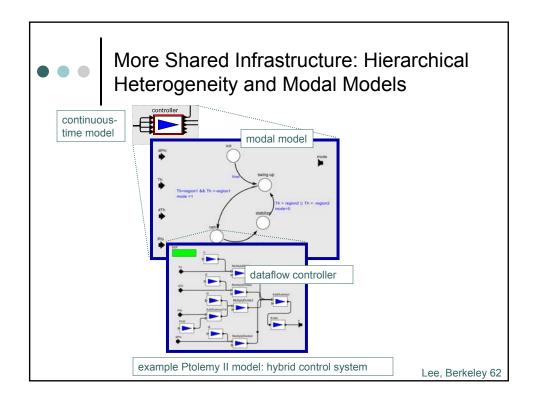
## Polymorphic Components - Component Library Works Across Data Types and Domains

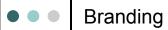
- Data polymorphism:
  - Add numbers (int, float, double, Complex)
  - Add strings (concatenation)
  - Add composite types (arrays, records, matrices)
  - Add user-defined types
- Behavioral polymorphism:
  - In dataflow, add when all connected inputs have data
  - In a time-triggered model, add when the clock ticks
  - In discrete-event, add when any connected input has data, and add in zero time
  - In process networks, execute an infinite loop in a thread that blocks when reading empty inputs
  - In CSP, execute an infinite loop that performs rendezvous on input or output
  - In push/pull, ports are push or pull (declared or inferred) and behave accordingly
  - In real-time CORBA, priorities are associated with ports and a dispatcher determines when to add



By not choosing among these when defining the component, we get a huge increment in component re-usability. But how do we ensure that the component will work in all these circumstances?





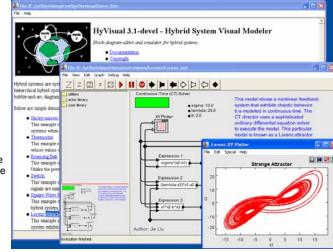


Ptolemy II configurations are Ptolemy II models that specify

- o welcome window
- o help menu contents
- library contents
- o File->New menu contents
- o default model structure
- o etc.

A configuration can identify its own "brand" independent of the "Ptolemy II" name and can have more targeted objectives.

An example is HyVisual, a tool for hybrid system modeling. VisualSense is another tool for wireless sensor network modeling.



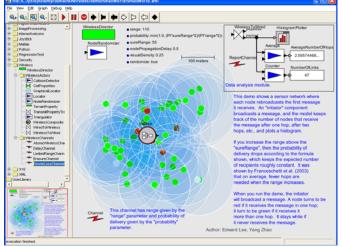
Lee, Berkeley 63

#### Ptolemy II Extension Points

- Define actors
- Interface to foreign tools (e.g. Python, MATLAB)
- Interface to verification tools (e.g. Chic)
- Define actor definition languages
- Define directors (and models of computation)
- Define visual editors
- Define textual syntaxes and editors
- Packaged, branded configurations

All of our "domains" are extensions built on a core infrastructure.

#### Example Extension: VisualSense



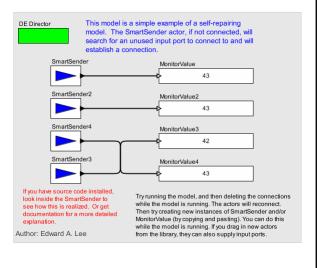
- Branded
- Customized visualization
- Customized model of computation (an extension of DE)
- Customized actor library
- Motivated some extensions to the core (e.g. classes, icon editor).

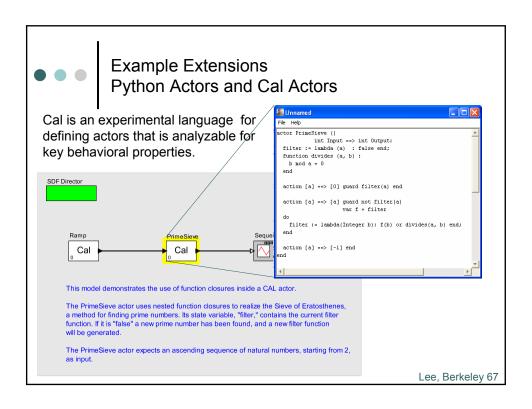
Lee, Berkeley 65

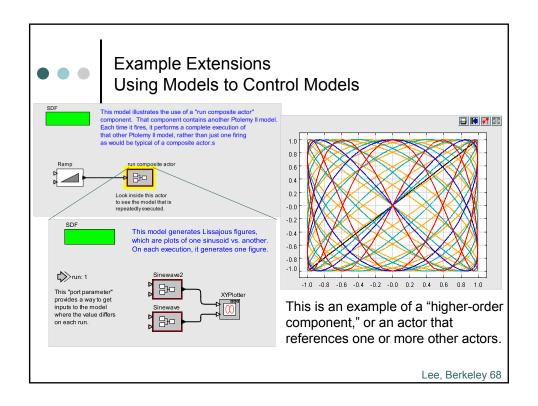
## • • •

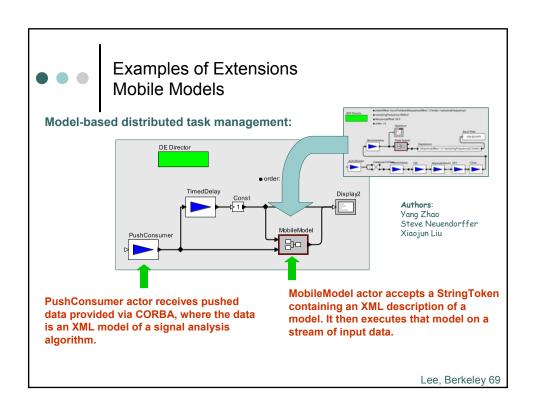
#### Example Extensions: Self-Repairing Models

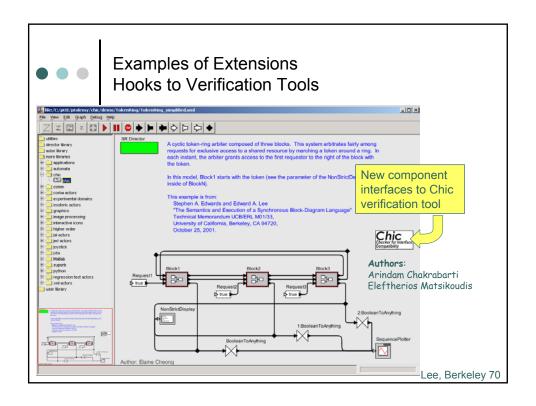
Concept demonstration built together with Boeing to show how to write actors that adaptively reconstruct connections when the model structure changes.

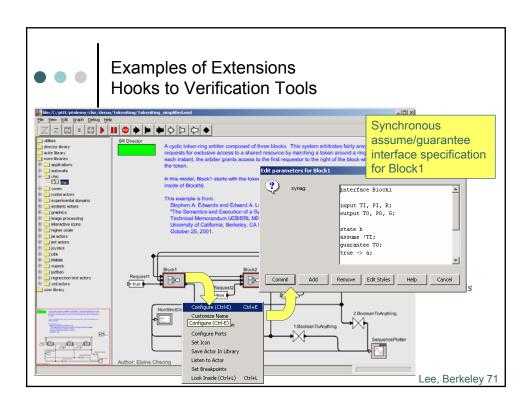


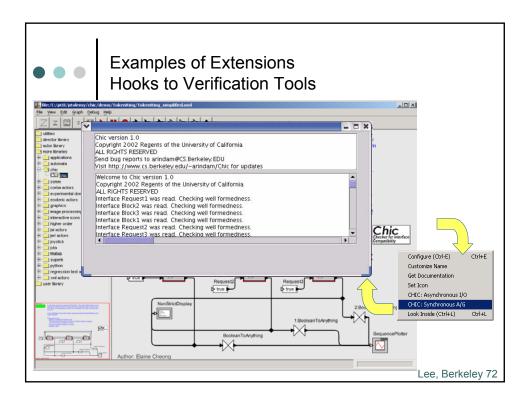












## Conclusion

- Threads suck
- There are many alternative concurrent MoCs
- o The ones you know are the tip of the iceberg
- Ptolemy II is a lab for experimenting with them
- Specializing MoCs can be useful
- o Mixing specialized MoCs can be useful.