

Building Unreliable Systems out of Reliable Components: *The Real Time Story*



Edward A. Lee

Professor, Chair of EE, and Associate Chair of EECS

CHESS: Center for Hybrid and Embedded Software Systems

UC Berkeley

Monterey Workshop Series

2005 Theme: Workshop on Networked Systems: realization of
reliable systems on top of unreliable networked platforms

September 23-25, 2005

Laguna Beach, CA



Electronics Technology Delivers Timeliness

... and the overlaying abstractions discard it.

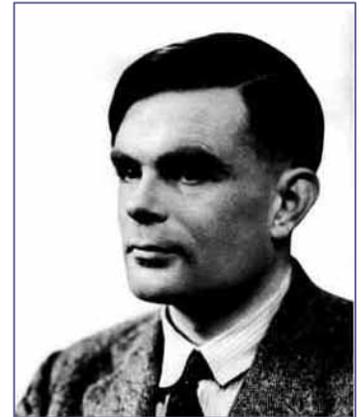
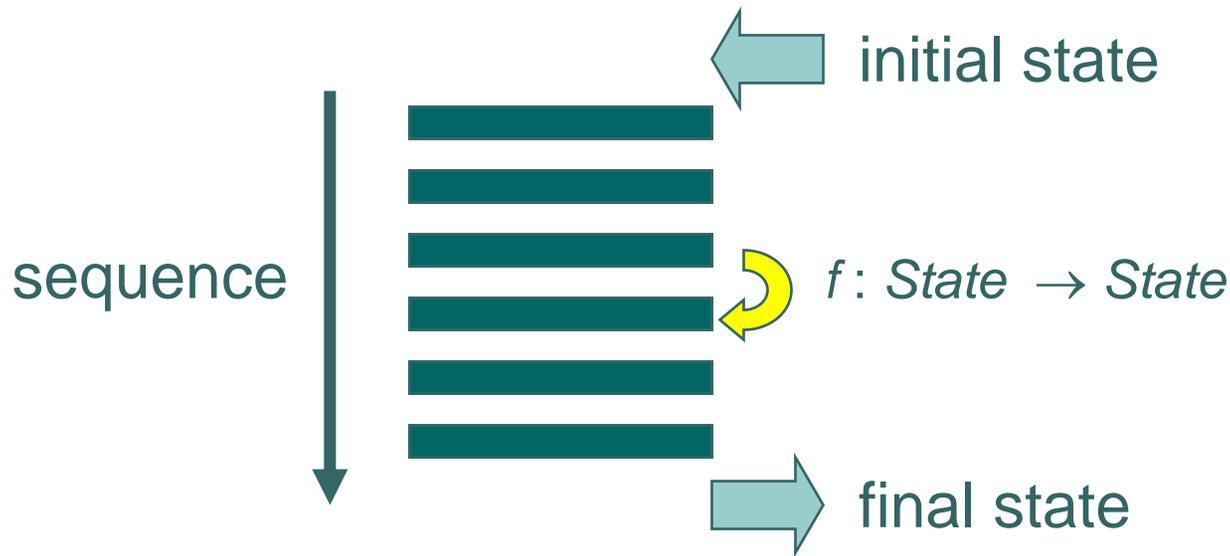


Computation in the 20th Century

$$f: \{0,1\}^* \rightarrow \{0,1\}^*$$



Computation in the 20th Century



Alan Turing

- Time is irrelevant
- All actions are ordered
- Nontermination is a defect
- Concurrency is an illusion



Exploiting the 20th Century Abstraction

- Programming languages
- Debuggers
- Virtual memory
- Caches
- Dynamic dispatch
- Speculative execution
- Power management (voltage scaling)
- Memory management (garbage collection)
- Just-in-time (JIT) compilation
- Multitasking (threads and processes)
- Networking (TCP)
- Theory (complexity)



What about timeliness?



Moore's
law has
saved us!



In Core Software Abstractions: Real-Time is Not

- Time is not in the semantics of programs.
 - Have to step outside the semantics to specify timing.
- Timing is a consequence of implementation not a property of design.
 - Measured on the bench
 - For a particular realization
- Resulting systems are brittle.
 - Small changes have big consequences
 - Ports to new platforms require redesign



The Myth of WCET

Worst-Case Execution Time

- True WCET can be thousands of times bigger than actual execution time.
- In many implementations, true WCET is not a useful number.
- Dubious WCET is what is actually used.
- Correctness of even safety-critical systems depends on WCET being correct.



What is Done in Practice

- Real-time systems are boxes, not software services.
- Critical real-time systems use idiosyncratic, non-mainstream processors (like DSPs).
- Designs are bench tested, then encased.



APOT

The question: What would have to change to achieve *absolutely, positively on time* (APOT)?

The answer: *nearly everything.*



What to do?

- Put time into programming languages
 - *Promising start:* Simulink, Giotto, Discrete-event models
- Rethink the OS/programming language split
 - *Promising start:* TinyOS/nesC
- Rethink the hardware/software split
 - *Promising start:* FPGAs with programmable cores
- Memory hierarchy with predictability
 - *Promising start:* Scratchpad memories vs. caches
- Memory management with predictability
 - *Promising start:* Bounded pause time garbage collection
- Predictable, controllable deep pipelines
 - *Promising start:* Pipeline interleaving + stream-oriented languages
- Predictable, controllable, understandable concurrency
 - *Promising start:* Synchronous languages, SCADE
- Networks with timing
 - *Promising start:* Time triggered architectures, time synchronization
- Computational dynamical systems theory
 - *Promising start:* Hybrid systems



Recall: Computation in the 20th Century

$$f: \{0,1\}^* \rightarrow \{0,1\}^*$$



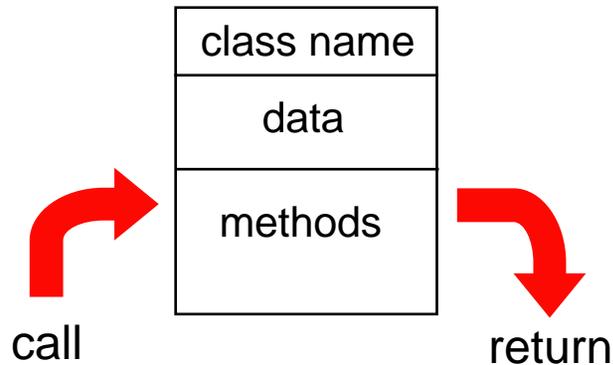
Computation in the 21st Century

$$f: [T \rightarrow \{0,1\}^*]^P \rightarrow [T \rightarrow \{0,1\}^*]^P$$



We Need Component and Composition Models with Time and Concurrency

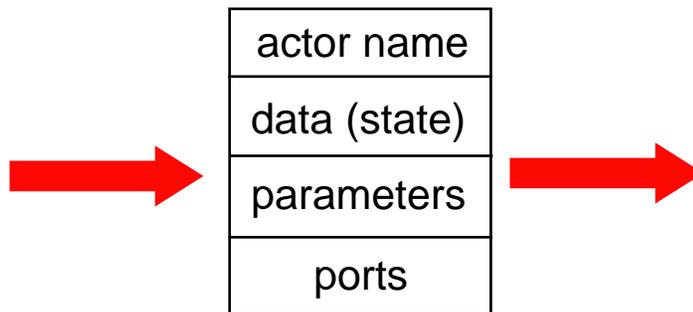
Object-oriented:



What flows through an object is sequential control

Stuff happens to objects

Actor oriented:



Actors make things happen

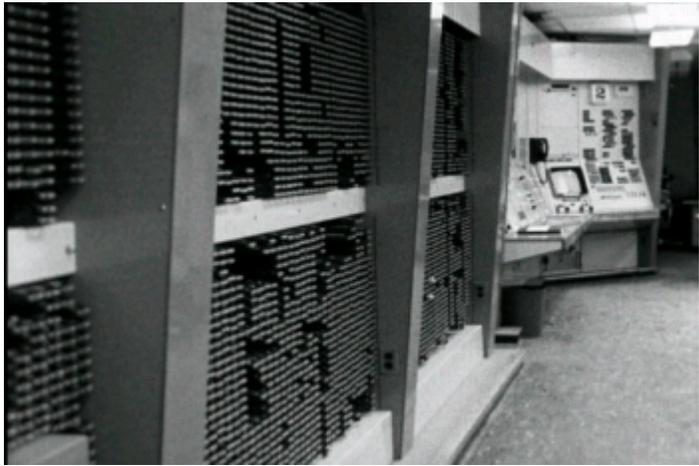
What flows through an object is streams of data



The First (?) Actor-Oriented Platform

The On-Line Graphical Specification of Computer Procedures

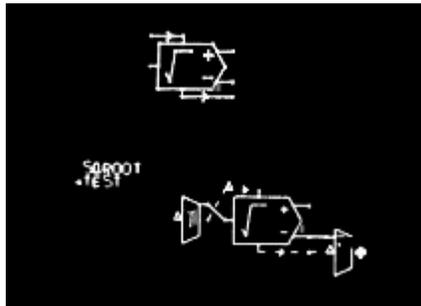
W. R. Sutherland, Ph.D. Thesis, MIT, 1966



MIT Lincoln Labs TX-2 Computer



Bert Sutherland with a light pen

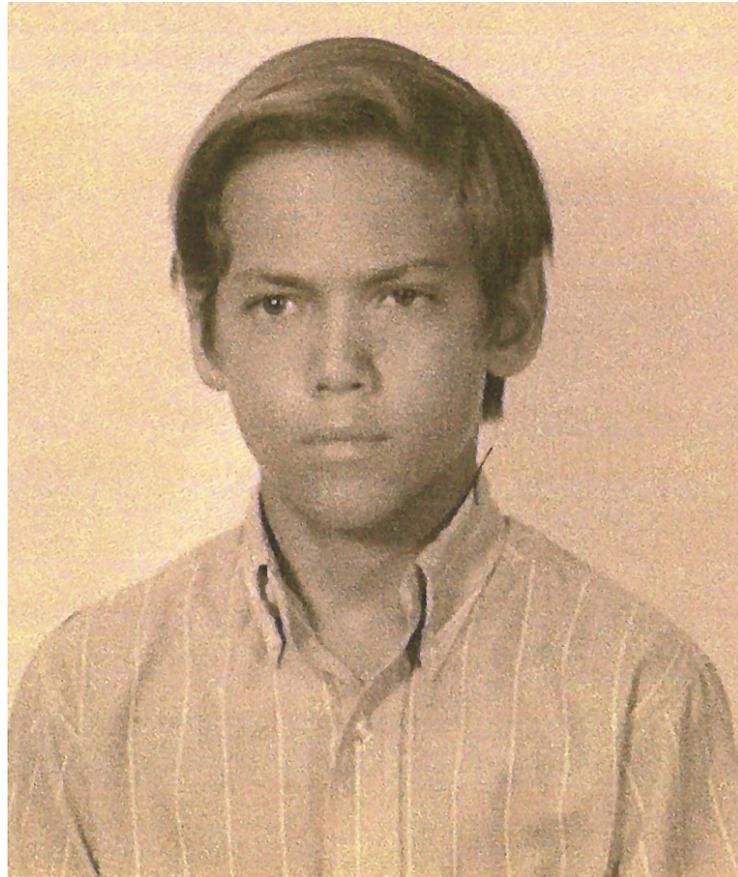


Bert Sutherland used the first acknowledged object-oriented framework (Sketchpad, created by his brother, Ivan Sutherland) to create the first actor-oriented programming framework.

Partially constructed actor-oriented model with a class definition (top) and instance (below).



Your Speaker in 1966





Modern Examples of Actor-Oriented Platforms

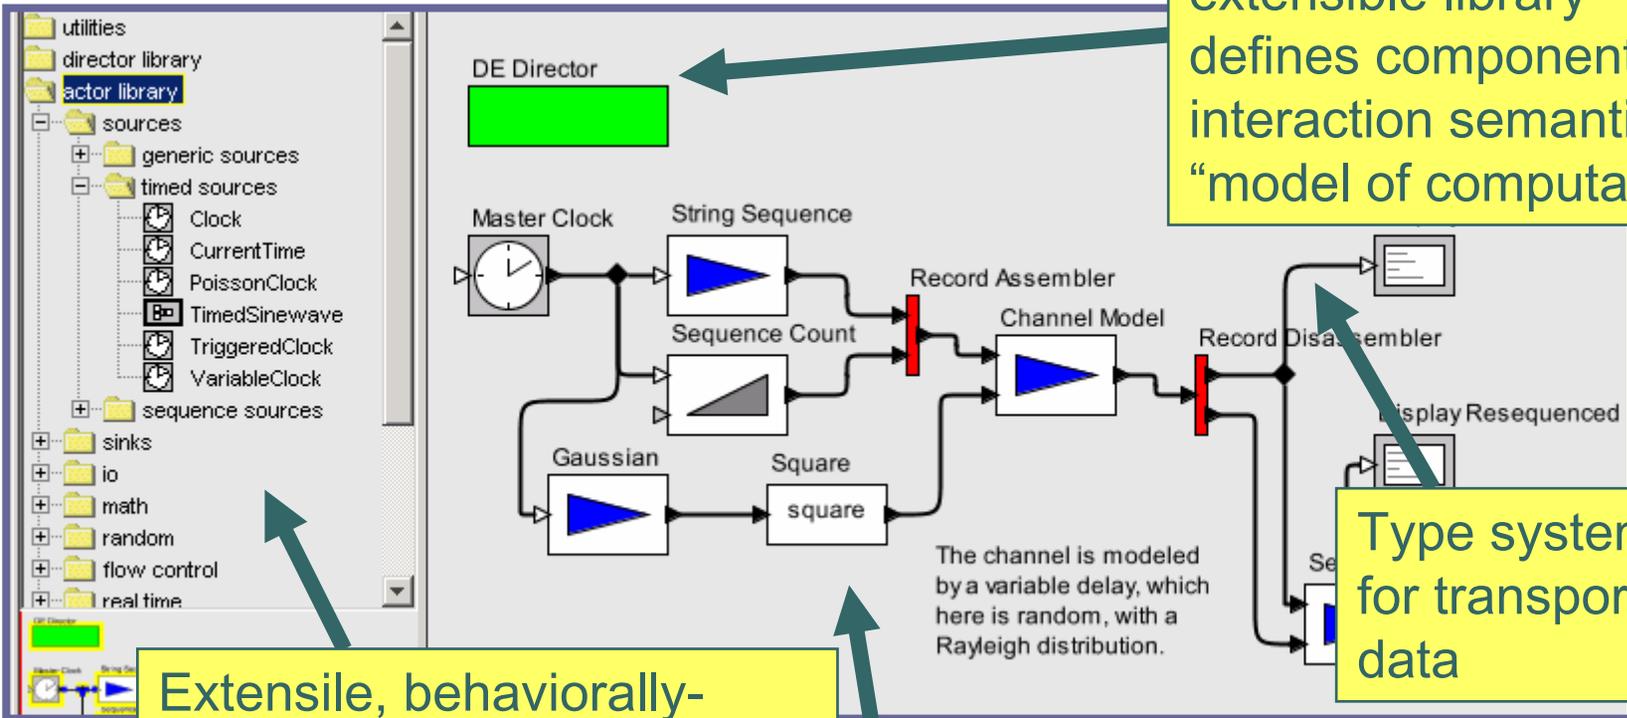
- Simulink (The MathWorks)
- LabVIEW (National Instruments)
- Modelica (Linkoping)
- OPNET (Opnet Technologies)
- Giotto and xGiotto (UC Berkeley)
- Polis & Metropolis (UC Berkeley)
- Gabriel, Ptolemy, and Ptolemy II (UC Berkeley)
- OCP, open control platform (Boeing)
- GME, actor-oriented meta-modeling (Vanderbilt)
- SPW, signal processing worksystem (Cadence)
- System studio (Synopsys)
- ROOM, real-time object-oriented modeling (Rational)
- Easy5 (Boeing)
- Port-based objects (U of Maryland)
- I/O automata (MIT)
- VHDL, Verilog, SystemC (Various)
- ...



Ptolemy II: Our Laboratory for Actor-Oriented Models of Computation

Concurrency management supporting dynamic model structure.

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



Extensible, behaviorally-polymorphic component library.

Type system for transported data

Visual editor supporting an abstract syntax



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
- DDE – distributed discrete events
- DDF – Dynamic dataflow
- DPN – distributed process networks
- DT – discrete time (cycle driven)
- FSM – finite state machines
- Giotto – synchronous periodic
- GR – 2-D and 3-D graphics
- PN – process networks
- SDF – synchronous dataflow
- SR – synchronous/reactive
- TM – timed multitasking

Most of
these are
actor
oriented.



A Start on a 21st Century Theory of Computation: The Tagged Signal Model

[Lee & Sangiovanni-Vincentelli, 1998]

- A set of *values* V and a set of *tags* T
- An *event* is $e \in T \times V$
- A *signal* s is a set of events. I.e. $s \subset T \times V$
- A *functional signal* is a (partial) function $s: T \rightarrow V$
- The set of all signals $S = 2^{T \times V}$

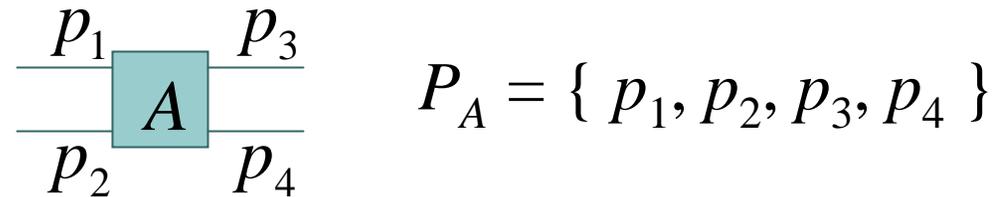
Related models:

- Interaction Categories [Abramsky, 1995]
- Interaction Semantics [Talcott, 1996]
- Abstract Behavioral Types [Arbab, 2005]



Actors, Ports, and Behaviors

An *actor* has a set of *ports* P



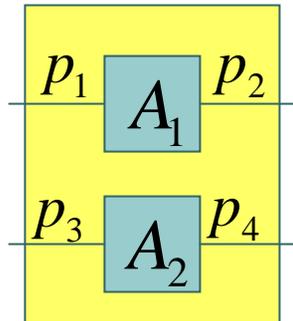
A behavior is a function $\sigma: P_A \rightarrow S$

An *actor* is a set of *behaviors* $A \subset [P_A \rightarrow S] = S^{P_A}$



Actor Composition

Composition is simple intersection
(of sets of functions)



$$P_1 = \{ p_1, p_2 \}$$

$$P_2 = \{ p_3, p_4 \}$$

$$A = A_1 \wedge A_2 \quad P = P_1 \cup P_2$$

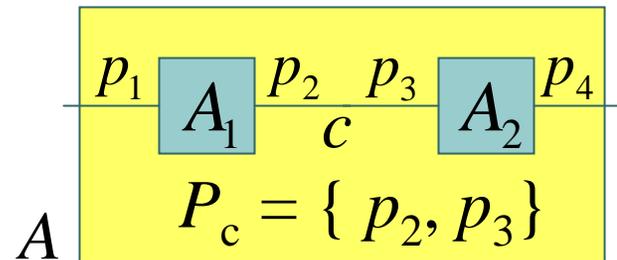
$$A = A_1 \wedge A_2 = \{ \sigma \mid \sigma \downarrow_{P_1} \in A_1 \text{ and } \sigma \downarrow_{P_2} \in A_2 \} \subset [P \rightarrow S]$$



Connectors

Connectors are trivial actors.

$$P_1 = \{p_1, p_2\} \quad P_2 = \{p_3, p_4\}$$

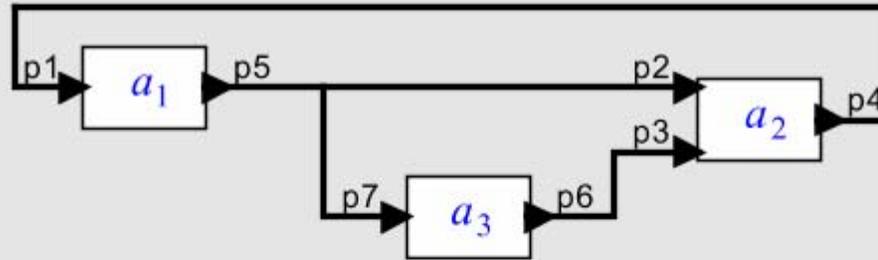


$$c \subset [P_c \rightarrow S], \quad \forall \sigma \in c, \quad \forall p_1, p_2 \in P_c, \quad \sigma(p_1) = \sigma(p_2)$$

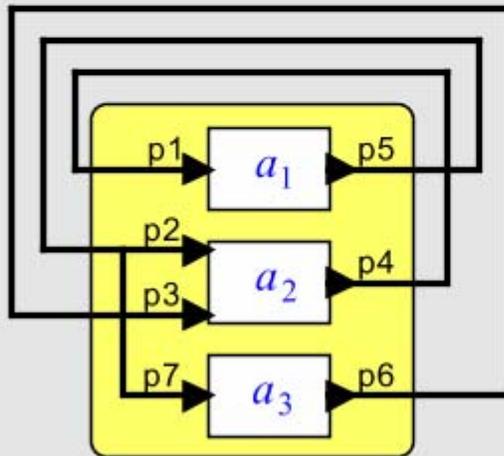
$$A = A_1 \wedge A_2 \wedge c$$



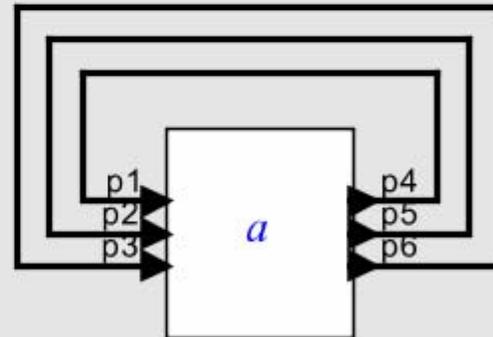
Tagged Signal Model Gives a Fixed-Point Semantics to Arbitrary Composition



(a)



(b)



(c)



Tagged Signal Model can be used on a Wide Variety of Concurrent and Timed Models of Computation

- CSP – concurrent threads with rendezvous
- CT – continuous-time modeling
- DE – discrete-event systems
- DDF – Dynamic dataflow
- DT – discrete time
- Giotto – synchronous periodic
- PN – process networks
- SDF – synchronous dataflow
- SR – synchronous/reactive



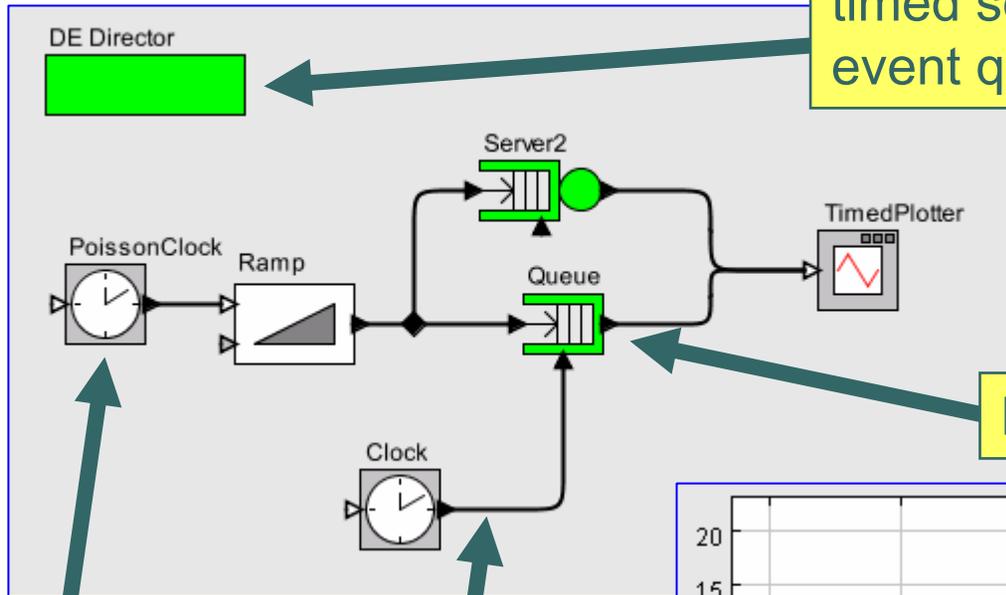
Application of this Theory of Computation: Discrete-Event Systems

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



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

DE Director implements timed semantics using an event queue

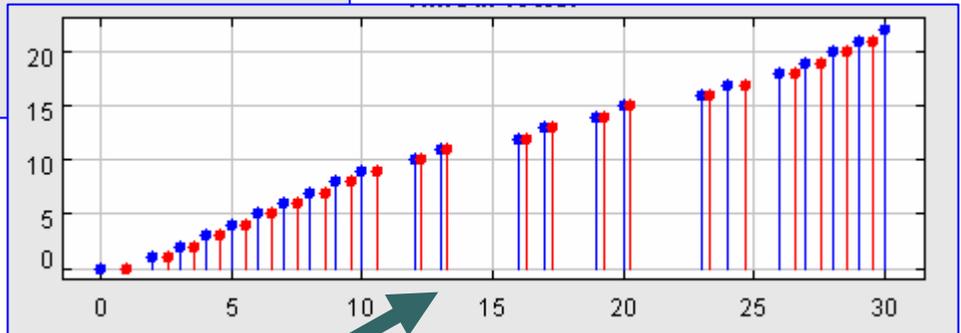


Reactive actors

Event source

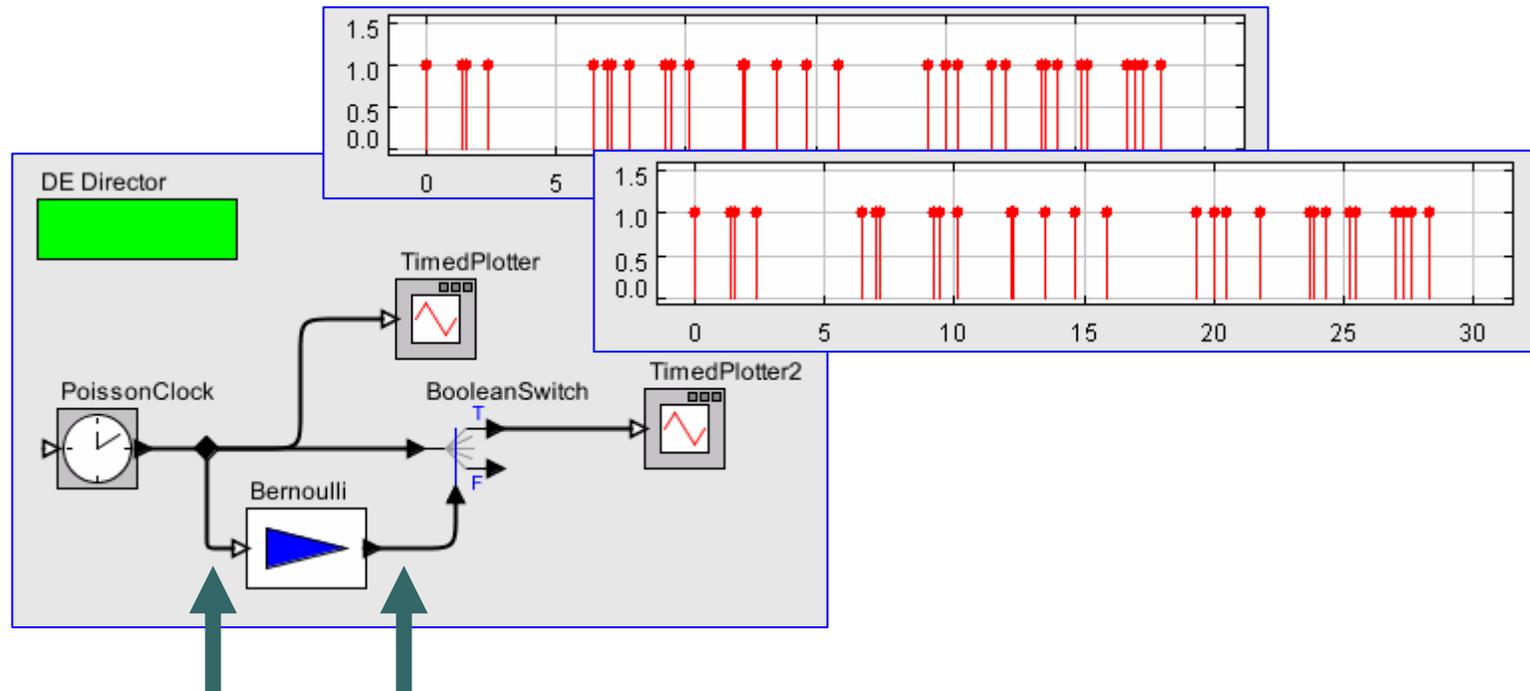
Signal

Time line





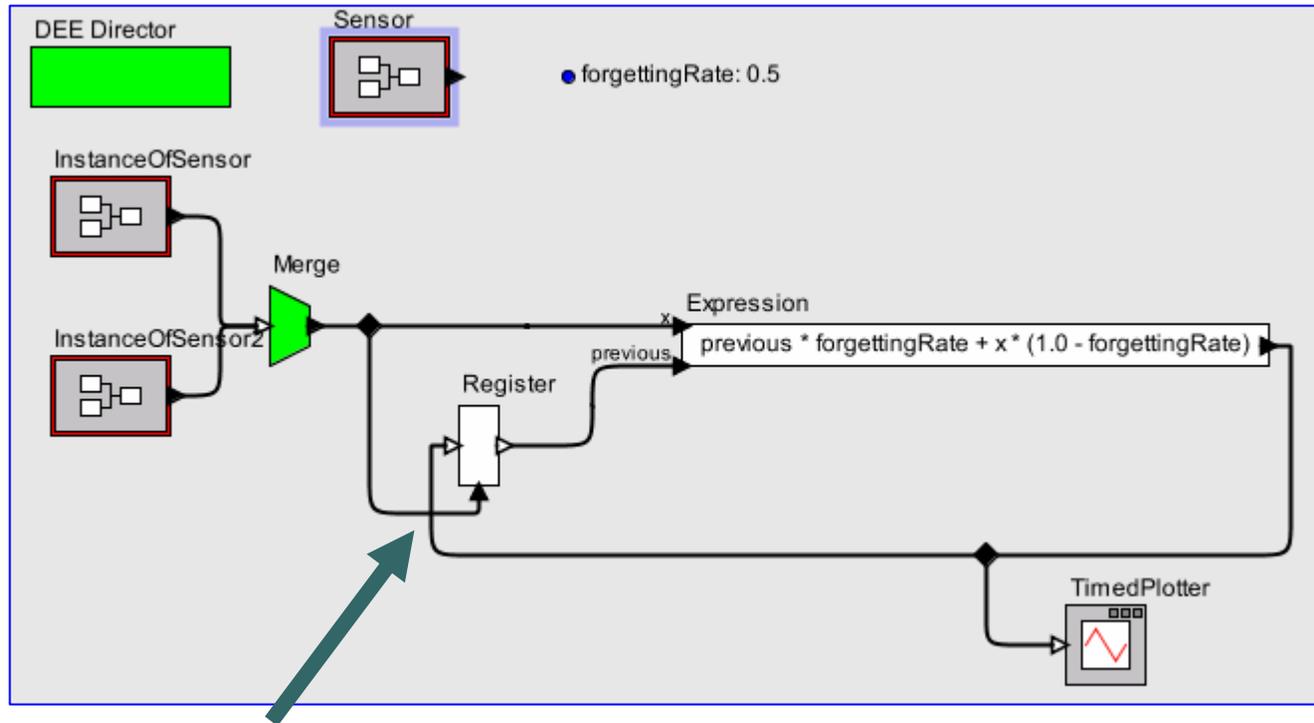
Semantics Clears Up Subtleties: Simultaneous Events



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.



Semantics Clears Up Subtleties: Feedback

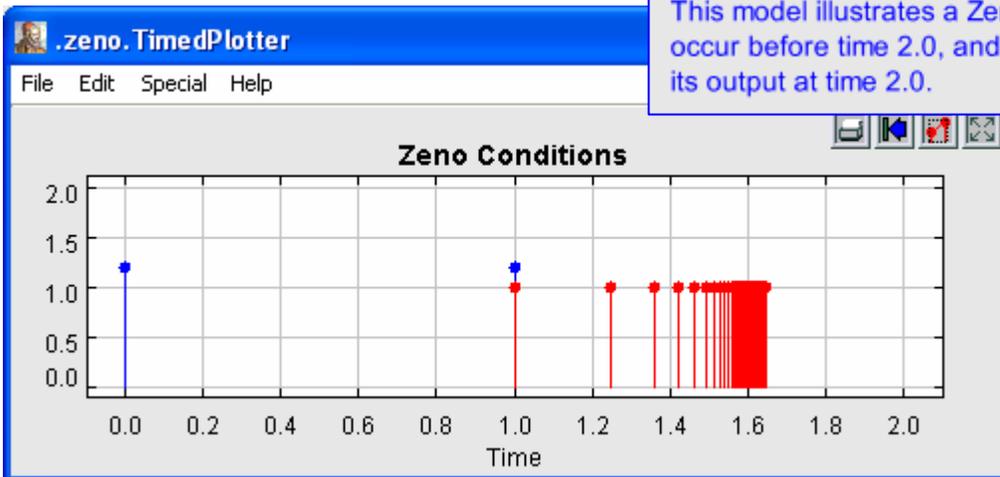
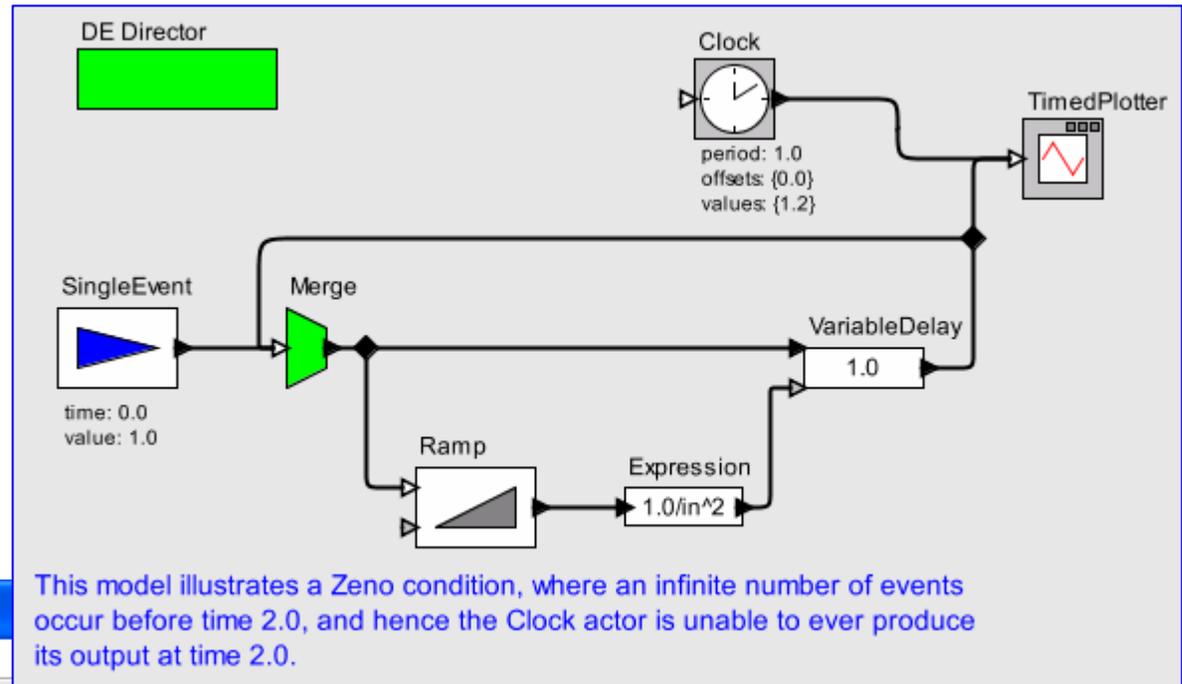


Data precedence analysis has to take into account the non-strictness of this actor (that an output can be produced despite the lack of an input).



Semantics Clears Up Subtleties: Zeno Systems

DE systems may have an infinite number of events in a finite amount of time. Carefully constructed semantics gives these systems meaning.





Example of Current Research Challenges

Use distributed discrete-event systems as a timed model of computation for embedded software in unreliable, sporadically connected networks, such as wireless sensor networks.

The most interesting possibilities are based on distributed consensus algorithms (as in Croquet, Reed, Lamport).

Research challenges include:

- Defining the semantics
- Combining the semantics heterogeneously with others. E.g.:
 - Signal processing for channel modeling
 - TinyOS for node functionality
- Creating efficient runtime environments
- Building the design environment



Application of this Theory of Computation: Hybrid Systems

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



Standard Model for Continuous-Time Signals

The usual formulation of the signals of interest is a function from the time line T (a connected subset of the reals) to the reals:

$$p: T \rightarrow \mathbb{R}$$

$$\dot{p}: T \rightarrow \mathbb{R}$$

$$\ddot{p}: T \rightarrow \mathbb{R}$$

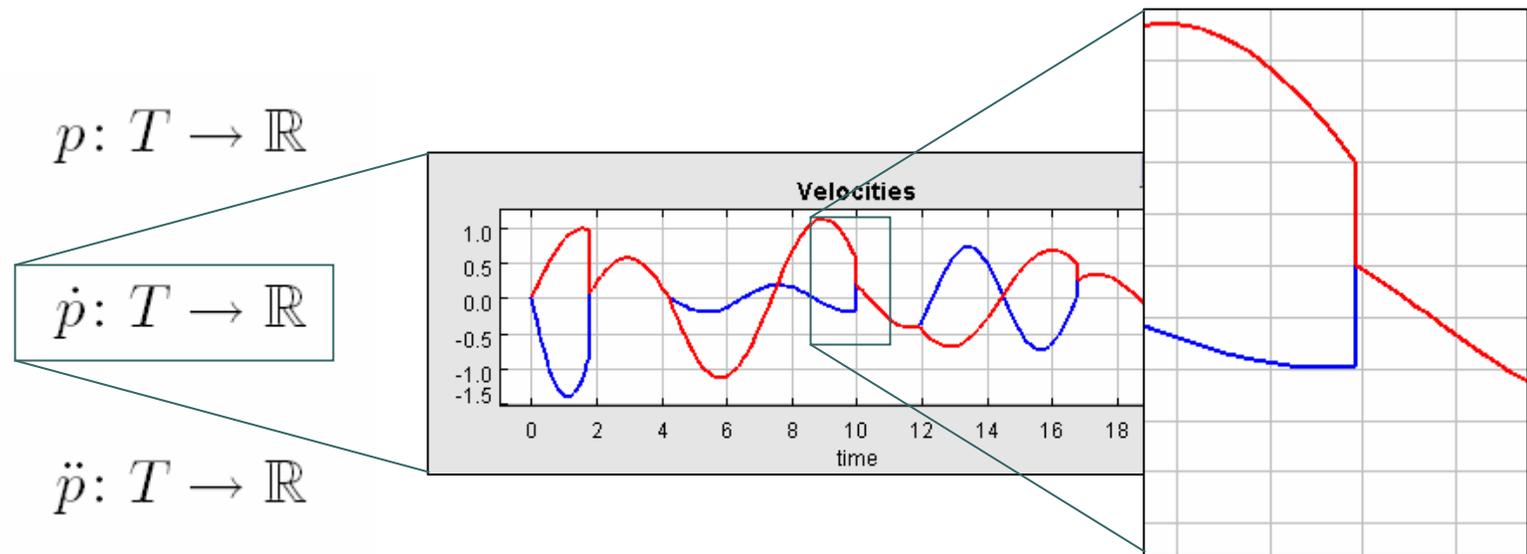
Such signals are continuous at $t \in T$ if (e.g.):

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



Piecewise Continuous Signals

In hybrid systems of interest, signals have discontinuities.



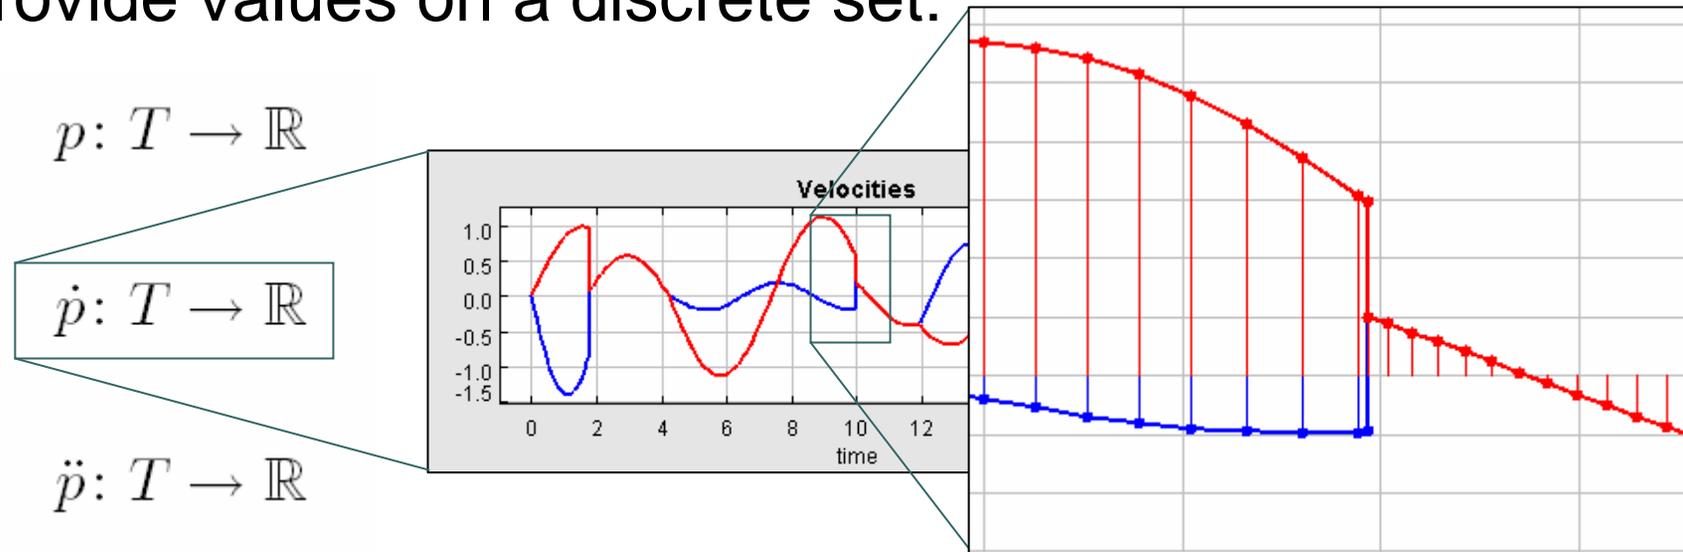
Piecewise continuous signals are continuous at all $t \in T \setminus D$ where $D \subset T$ is a *discrete set*.¹

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



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.



Definition: *Continuously Evolving Signal*

Change the domain of the function:

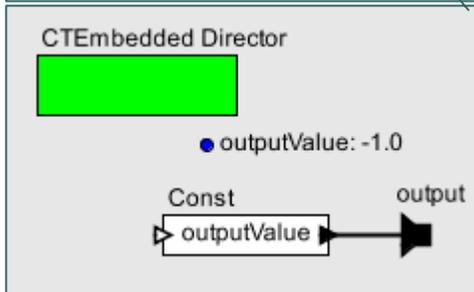
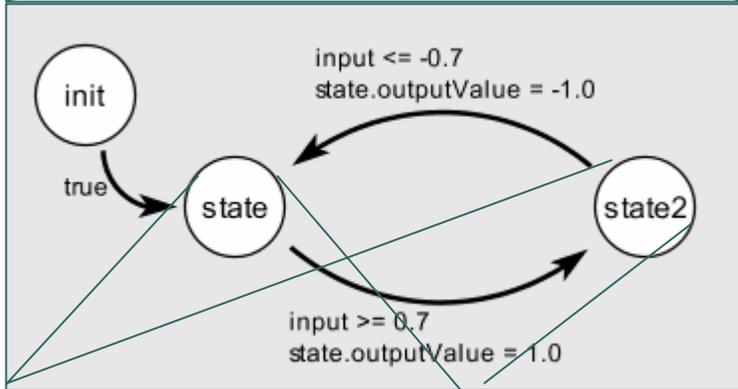
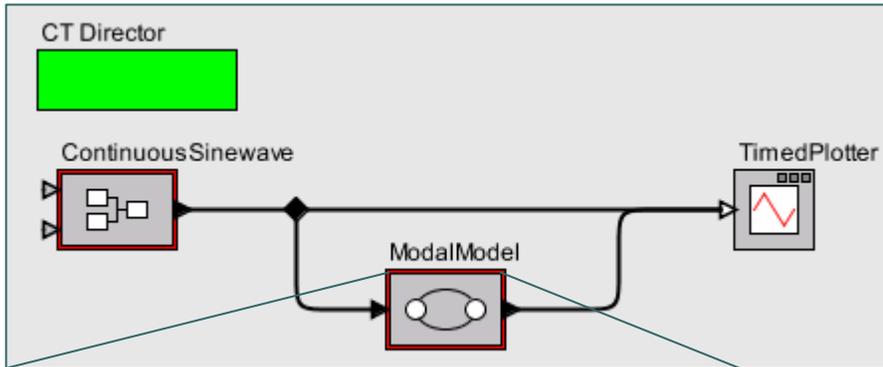
$$x : T \times \mathbb{N} \rightarrow V$$

Where T is a connected subset of the reals and \mathbb{N} is the set of natural numbers.

At each time $t \in T$, the signal x has a sequence of values. Where the signal is continuous, all the values are the same. Where is discontinuous, it has multiple values.



Simple Example: Hysteresis

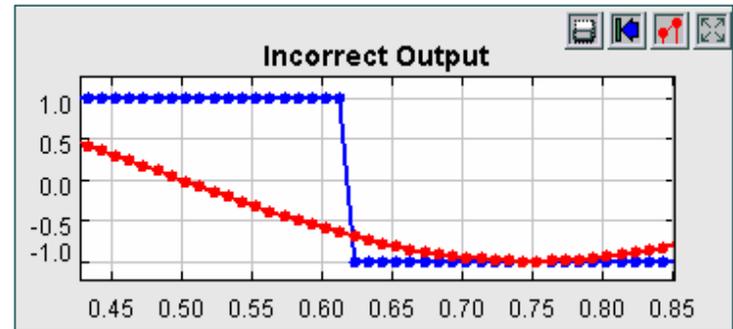
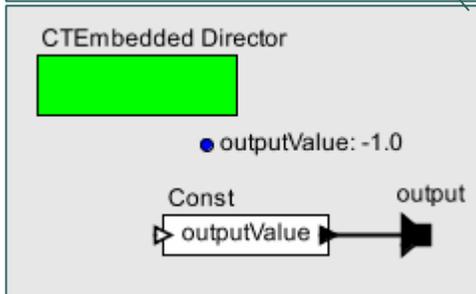
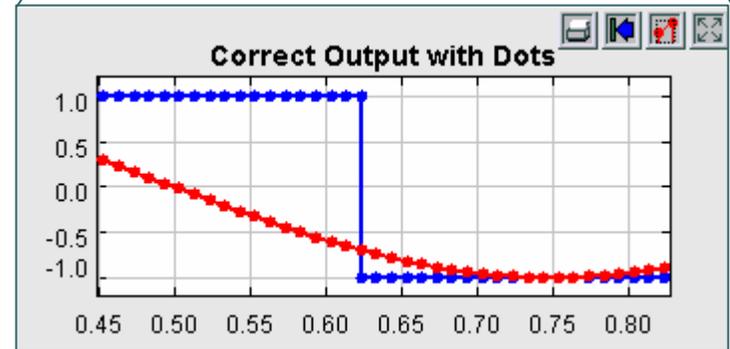
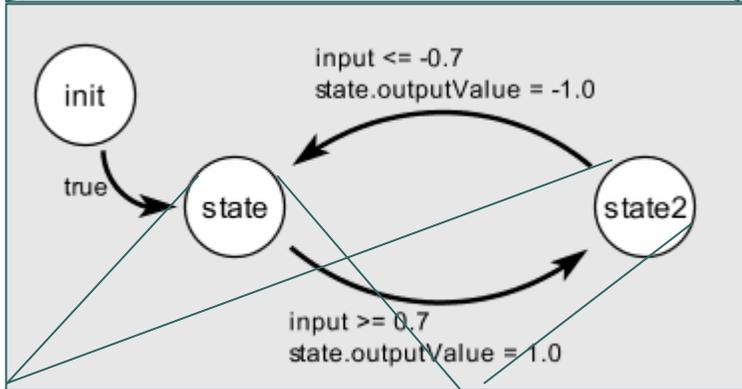
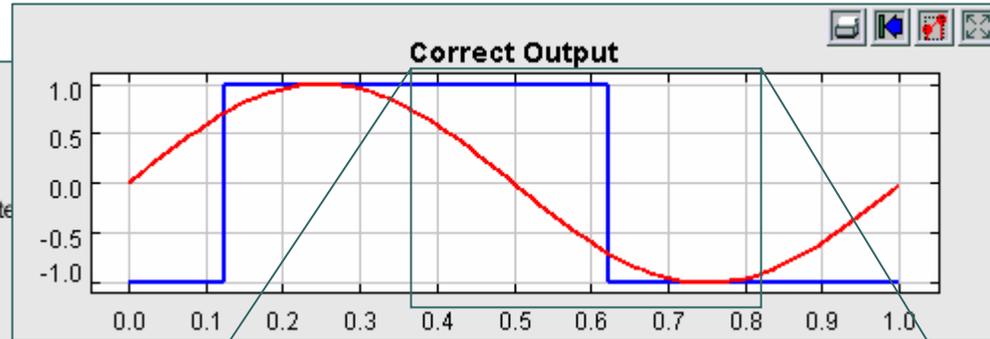
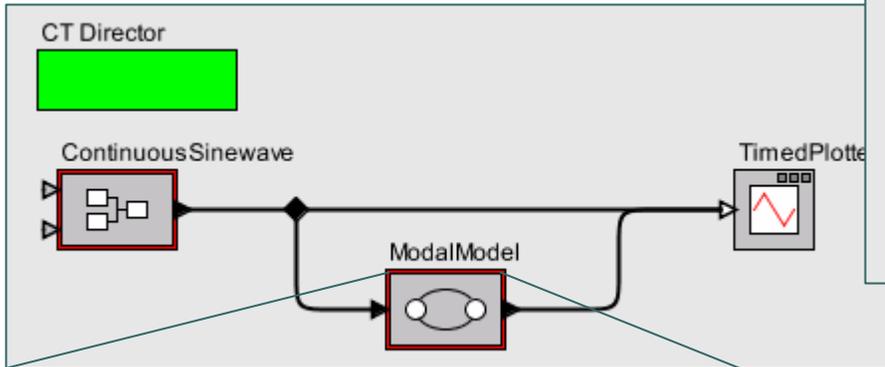


This model shows the use of a two-state FSM to model hysteresis.

Semantically, the output of the ModalModel block is discontinuous. If transitions take zero time, this is modeled as a signal that has two values *at the same time*, and *in a particular order*.



Signals Must Have Multiple Values at the Time of a Discontinuity



Discontinuities need to be semantically distinguishable from rapid continuous changes.



Initial and Final Value Signals

A signal $x: T \times \mathbb{N} \rightarrow V$ has no *chattering* *Zeno* condition if there is an integer $m > 0$ such that

$$\forall n > m, \quad x(t, n) = x(t, m)$$

A non-chattering signal has a corresponding *final value signal*, $x_f: T \rightarrow V$ where

$$\forall t \in T, \quad x_f(t) = x(t, m)$$

It also has an *initial value signal* $x_i: T \rightarrow V$ where

$$\forall t \in T, \quad x_i(t) = x(t, 0)$$



Piecewise Continuous Signals

A piecewise continuous signal is a non-chattering signal

$$x : T \times \mathbb{N} \rightarrow V$$

where

- The initial signal x_i is continuous on the left,
- The final signal x_f is continuous on the right, and
- The signal x has only one value at all $t \in T \setminus D$ where $D \subset T$ is a discrete set.



Our Current Projects

- Abstract semantics (Cataldo, Liu, Matsikoudis, Zheng)
 - Behavioral polymorphism
 - Actor semantics (prefire, fire, postfire)
 - Compositional directors
 - Time semantics
 - Causality interfaces
- Distributed computing (Feng, Zhao)
 - Robust distributed consensus
 - Data coherence (distributed caches)
 - Time synchronization
- Real-time software (Bandyopadhyay, Cheong, Zhou)
 - Time-based models vs. dataflow models
 - Deterministic, understandable multitasking
 - Memory hierarchy with scratchpad memory
 - Code generation
- Hybrid systems (Cataldo, Zheng)
 - Operational semantics
 - Stochastic hybrid systems
 - Aspect-oriented multi-view modeling
 - Code generation



Conclusion

The time is right to create the 21-st century theory of (embedded) computing.