



Introduction to Embedded Systems


Edward A. Lee & Sanjit Seshia

UC Berkeley
EECS 124
Spring 2008


Copyright © 2008, Edward A. Lee & Sanjit Seshia, All rights reserved

Lecture 21: Concurrency Models 1

Simple Example: Spectrum Analysis




How do we keep the non-time critical path from interfering with the time-critical path?



EECS 124, UC Berkeley: 2

Dataflow Models



Buffered communication between concurrent components (*actors*).


Static scheduling: Assign to each thread a sequence of actor invocations (*firings*) and repeat forever.

Dynamic scheduling: Each time `dispatch()` is called, determine which actor can fire (or is firing) and choose one.

May need to implement interlocks in the buffers.

EECS 124, UC Berkeley: 3


Buffers for Dataflow



- Unbounded buffers require memory allocation and deallocation schemes.
- Bounded size buffers can be realized as *circular buffers* or *ring buffers*, in a statically allocated array.
 - A *read pointer* r is an index into the array referring to the first empty location. Increment this after each read.
 - A *fill count* n is unsigned number telling us how many data items are in the buffer.
 - The next location to write to is $(r + n)$ modulo buffer length.
 - The buffer is empty if $n == 0$
 - The buffer is full if $n == \text{buffer length}$
 - Can implement n as a semaphore, providing mutual exclusion for code that changes n or r .

EECS 124, UC Berkeley: 4

Abstracted Version of the Spectrum Example: Non-preemptive scheduling




Assume infinitely repeated invocations, triggered by availability of data at A.

Suppose that C requires 8 data values from A to execute. Suppose further that C takes much longer to execute than A or B. Then a schedule might look like this:




EECS 124, UC Berkeley: 5

Uniformly Timed Schedule




A preferable schedule would space invocations of A and B uniformly in time, as in:




minimum latency

EECS 124, UC Berkeley: 6

Non-Concurrent Uniformly Timed Schedule




Notice that in this schedule, the rate at which A and B can be invoked is limited by the execution time of C.




EECS 124, UC Berkeley: 7

Concurrent Uniformly Timed Schedule: Preemptive schedule




With preemption, the rate at which A and B can be invoked is limited only by total computation:




EECS 124, UC Berkeley: 8

Ignoring Initial Transients, Abstract to Periodic Tasks

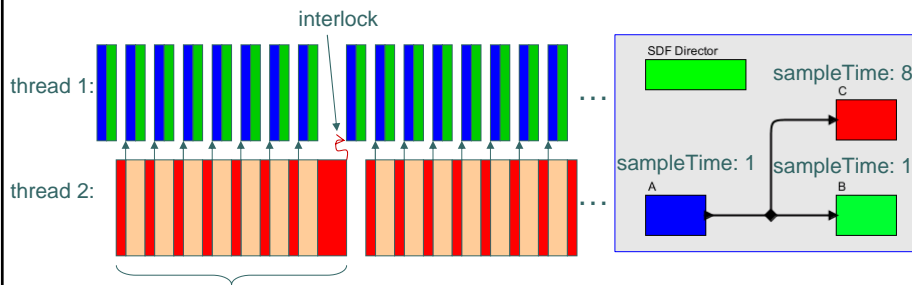


In steady-state, the execution follows a simple periodic pattern:



EECS 124, UC Berkeley: 9

Requirement 1 for Determinacy: Periodicity




If the execution of C runs longer than expected, data determinacy requires that thread 1 be delayed accordingly. This can be accomplished with semaphore synchronization. But there are alternatives:

- Throw an exception to indicate timing failure.
- "Anytime" computation: use incomplete results of C

EECS 124, UC Berkeley: 10


Requirement 1 for Determinacy: Periodicity



If the execution of C runs shorter than expected, data determinacy requires that thread 2 be delayed accordingly. That is, it must not start the next execution of C before the data is available.

EECS 124, UC Berkeley: 11

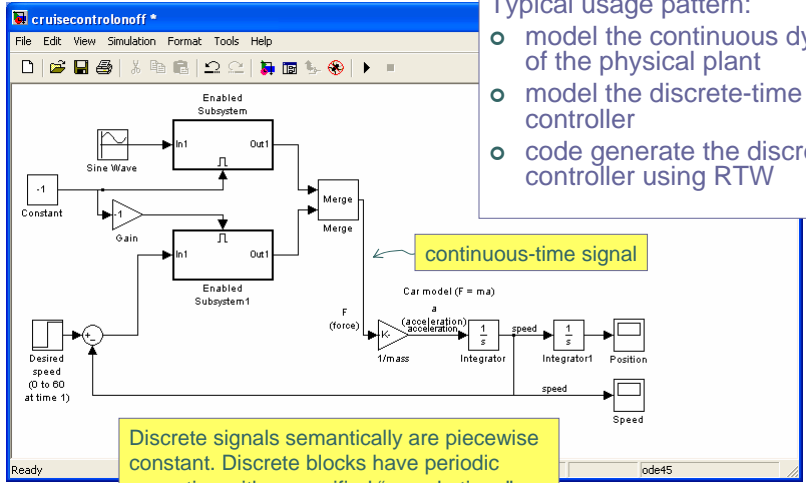
Semaphore Synchronization Required Exactly Twice Per Major Period



Note that semaphore synchronization is *not* required if actor B runs long because its thread has higher priority. Everything else is automatically delayed.

EECS 124, UC Berkeley: 12

Simulink and Real-Time Workshop (The MathWorks)

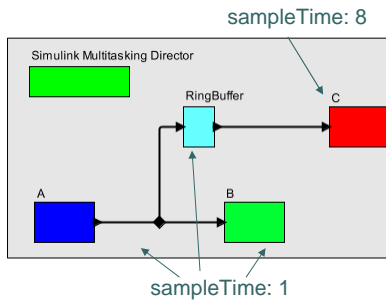


- Typical usage pattern:
- model the continuous dynamics of the physical plant
 - model the discrete-time controller
 - code generate the discrete-time controller using RTW

continuous-time signal

EECS 124, UC Berkeley: 13


Explicit Buffering is required in Simulink



In Simulink, unlike dataflow, there is no buffering of data. To get the effect of presenting to C 8 successive samples at once, we have to explicitly include a buffering actor that outputs an array.

EECS 124, UC Berkeley: 14


Requirement 2 for Determinacy: Data Integrity During Execution



It is essential that input data remains stable during one complete execution of C, something achieved in Simulink with a zero-order hold (ZOH) block.

EECS 124, UC Berkeley: 15


Simulink Strategy for Preserving Determinacy



In “Multitasking Mode,” Simulink requires a Zero-Order Hold (ZOH) block at any downsampling point. The ZOH runs at the slow rate, but at the priority of the fast rate. The ZOH holds the input to C constant for an entire execution.


EECS 124, UC Berkeley: 16

In Dataflow, Interlocks and Built-in Buffering take care of these dependencies





No ZOH block is required!

For dataflow, a one-time interlock ensures sufficient data at the input of C:



EECS 124, UC Berkeley: 17


Consider a Low-Rate Actor Sending Data to a High-Rate Actor



Note that data precedences make it impossible to achieve uniform timing for A and C with the periodic non-concurrent schedule indicated above.

EECS 124, UC Berkeley: 18

Overlapped Iterations Can Solve This Problem




This solution takes advantage of the intrinsic buffering provided by dataflow models.

For dataflow, this requires the initial interlock as before, and the same periodic interlocks.

EECS 124, UC Berkeley: 19

Simulink Strategy



Without buffering, the Delay provides just one initial sample to C (there is no buffering in Simulink). The Delay and ZOH run at the rates of the slow actor, but at the priority of the fast ones.

Part of the objective seems to be to have no initial transient. Why?

EECS 124, UC Berkeley: 20

Discussion Questions

- What about more complicated rate conversions (e.g. a task with `sampleTime 2` feeding one with `sampleTime 3`)?
- How can these ideas be extended to non-periodic execution?