

## Course Web Page New Students?

1. From Department HomePage:  
<http://www.sce.carleton.ca>
2. Pick Course Materials (on left)
3. Pick SYSC 5701  
<http://www.sce.carleton.ca/dept/sce.php/courses/sysc-5701>  
protected content: user: **sysc-5701**  
password: **rtos**

Jan 7, 2014

1

## SYSC 5701 Operating System Methods for Real-Time Applications

### Motivation

*Winter 2014*

## Broad Background

- systems concepts, computer systems
- time
- software engineering: development, design
- concurrency
- interrupts

Jan 7, 2014

3


## System

- a set of components that interact to accomplish an objective
- can be applied to just about anything! ☺

Jan 7, 2014

4

## Uniprocessor Computer System



- **Objective:** involves maintaining input/output relationships at the I/O / External interface

Jan 7, 2014

5

## Variations: Multiprocessor


- more than one processor → shared bus
- processors share global resources
- a processor may also have private local resources connected via a secondary (private) bus structure  
(not shown below)

**multicore !!**

Jan 7, 2014

6

## Multiprocessor (con't)



Jan 7, 2014

7

## Variations: Network

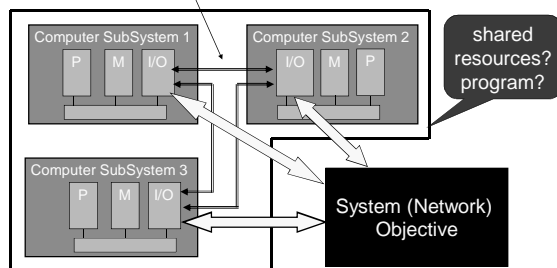
- computer subsystems interconnected via I/O components
- subsystems do not share resources via shared bus
- sharing a resource is more complicated!  
→ requires co-operation of subsystems
- subsystems co-operate to accomplish network-wide objective

Jan 7, 2014

8

## Network (con't)


subsystems must communicate to co-operate



Jan 7, 2014

9

## Real-Time Systems



**Objective:**

- maintain **time-constrained** input/output relationships between computing system and external devices/systems
- How should these be **described**?

Jan 7, 2014

10

## (Typical) Hard vs. Soft Real-Time

- **Hard Real-Time**
  - failure to meet time constraints is catastrophic
  - recovery may be difficult, or futile
  - e.g. reactor melt-down, plane crash, loss of life
- **Soft Real-Time**
  - occasional failure to meet time constraints is inconvenient but not catastrophic
  - try again, or be patient
  - e.g. no dial tone, lost voice packet




Safety Critical

Jan 7, 2014

11

## Describing Systems

- **Requirements:** specify the objectives in terms of behaviour at the interface to the external devices/systems




- **Implementation:** describes how the computing system is utilized to meet the requirements
- Why is it useful to describe both? What is a system "design"?

Jan 7, 2014

12

## Requirements vs. Implementation

- “**Ideally**”: the requirements are independent of the implementation  
→ **Abstraction** ← engineering



Jan 7, 2014

13

## Concurrent Activities


- are in progress at the same time
- **dependent activities**: interact to complete a higher objective
- **independent activities**: do not interact

May have concurrency in the requirements behaviour and in the implementation

Jan 7, 2014

14

## Stream-2-Pipe Example




concurrent activities at interface:

- input (slow) data streams: A, B, C, D
- output (fast) data pipes: Pipe1, Pipe2

Jan 7, 2014

15

## Example (con't)



- streams A and B are compressed/multiplexed into stream Pipe1
- streams C and D are compressed/multiplexed into stream Pipe2

Jan 7, 2014

16

## Example (con't)

Concurrency at **requirements** level:

- A, B & Pipe1 are **dependent** activities
- C, D & Pipe2 are **dependent** activities
- { A, B, Pipe1 } activities are **independent** of { C, D, Pipe2 } activities

Concurrency in **implementation**?

How might the system be implemented?

Jan 7, 2014

17

## Concurrency in Physical Implementations

- **real** concurrency: active h/w components that operate in parallel to support concurrent activities  
– e.g. processors, active I/O components
- **apparent** concurrency: active devices are shared to give the impression (over time) that external activities are being carried out concurrently

Jan 7, 2014

18

## Important Distinction!

- concurrency in requirements is part of the objective
  - **cannot** be altered by design decisions
- concurrency in implementation is a design decision
  - **not imposed** by requirements

As a result: Concurrent activities in requirements are often at a different granularity than concurrent activities in implementation.

Jan 7, 2014

19

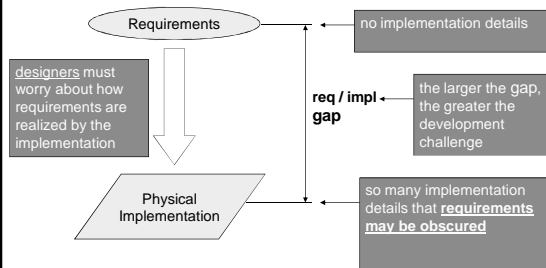
## Design for Concurrency

- mapping concurrency in requirements onto implementation resources is a design decision
  - **goal:** allocation of system (implementation) resources to achieve concurrency in requirements
- many tough design issues here!  
(more later!)

Jan 7, 2014

20

## Development Problem: requirements/implementation gap



Jan 7, 2014

21

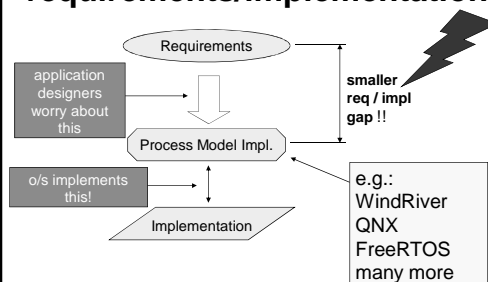
## To reduce/manage the requirements/implementation gap:

- introduce an intermediate level between requirements and implementation
  - resides “above” implementation
- virtual machine: deals with **concurrency** explicitly!
- introduce an abstract **process model**
- design implementation in terms of the process model
- **operating system** provides process model support

Jan 7, 2014

22


## Modified Development Problem: reduced requirements/implementation gap



Jan 7, 2014

23

## Before & After



Jan 7, 2014

24

## What SYSC 5701 Is ....

- concerned with using a **process model** to help reduce the development challenges for real-time applications
- primary concern: **designer's perspective!**
- Goals:
  - simplify the implementation of concurrency
  - hide some machine details
  - use “standard” process model
  - simplify the mapping of concurrency in requirements onto concurrency in implementation

Jan 7, 2014

25

## What SYSC 5701 Is Not ....

- **NOT** concerned with particular real-time applications
- **NOT** about Linux or Windows

Jan 7, 2014

26

## So ... what's so hard about concurrency? ☺

- event-driven vs. sequential mindset
- interference – shared resources
- synchronization – mutual exclusion, coordinate progress
- communication among concurrent activities
  - for application purposes & synchronization

Will elaborate on these in the rest of these slides

Jan 7, 2014

27

## Sequential Mindset

- control is managed sequentially
  - **only one thread of control**
- hardware/state is polled to decide when to perform work
- response to events depends on when event sources are polled

Jan 7, 2014

28

## Sequential Mindset: Polling

General form of polling-only implementation:

```
loop (forever)
{
    poll for next event/work to do
    process events/work as needed
}
```

Jan 7, 2014

29

## Polling & Priority

- for polled events, can often give work relative priorities
- e.g. poll all devices and decide on processing order
- higher-priority work: performed a.s.a.p.
  - e.g. service I/O hardware
- lower-priority work: after higher-priority work

Jan 7, 2014

30

## Timing Example:

- suppose a **h/w timer** is being used to implement a displayed clock
- h/w timer “tick” every millisecond
  - can poll for tick
- update **display clock** every second

Jan 7, 2014

31

## Polling Approach

```
poll h/w timer
if ( tick )
{ count++;
  if ( count == 1000 )
  { count = 0; }
  update display;
}
```

Jan 7, 2014

32

## Priority in Timer Example

- manipulating count is **higher-priority** processing
- failure to sense every tick = lost time !
- must poll "often enough" to sense all ticks
- update clock display is **lower-priority** processing
- could be delayed “a bit” in favour of higher-priority processing

Jan 7, 2014

33

## Event-Driven Mindset: H/W Interrupts

- high-priority processing performed by h/w Interrupt **Service Routines** ( **ISRs** )
- h/w generates interrupt (signal) when event occurs
  - e.g. h/w timer tick
- **signal** causes processor to execute ISR
  - no s/w involved in invocation of ISR!

If you don't recall about **interrupts** – be sure to read about them in any microprocessor system text! See doc link on webpage.

Jan 7, 2014

34

## ISR Related Control Flow

1. current s/w state is saved on stack (registers: including status (e.g. flags) and program counter)
  - the current software is suspended! (interrupted! pre-empted!)
2. ISR runs
3. prior state (1) is restored and s/w continues

If you don't recall about **interrupts** – be sure to read about them in any microprocessor system text! See doc link on webpage.

Jan 7, 2014

35

## Interrupt & ISR



- Similar to a h/w invoked function call
- **NO** s/w involved in invocation!!
- interrupted s/w (s/w<sub>x</sub>) does not “know” it was momentarily suspended or that the ISR executed! (i.e. that s/w<sub>x</sub> was pre-empted)

Jan 7, 2014

36

## Event-Driven Mindset: Interrupts & Concurrency

- **processor is shared** between the **threads** of control associated with ISRs and the sequential **thread** of the main program
  - shared processor = virtual concurrency
- h/w interrupts are **asynchronous**
  - the result of the actions of active hardware devices
- ISRs run due to h/w event handling, not due to sequential s/w sensing of events!

Jan 7, 2014

37

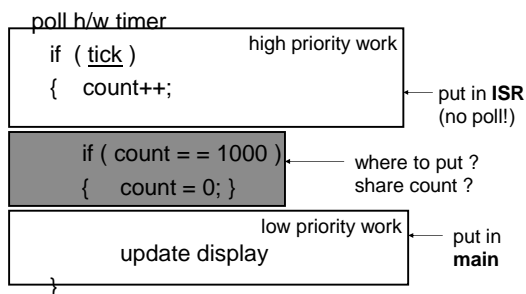
## To use Interrupt-Driven Approach:

- place high-priority processing in ISRs
- place low-priority processing in main (sequential) program
- ISRs and main must **communicate**
- main requests that high-priority work to be performed by ISRs
- ISRs inform main of completed work
- communicate using **shared variables**

Jan 7, 2014

38

## Recall Previous Timer Example



Jan 7, 2014

39

## Timer Example Revised

Suppose ISR and main share: boolean **SECOND**

- **in ISR:** `count ++;`  
`if ( count == 1000 )`  
`{ count = 0;`  
`SECOND = TRUE; }`
- **in main:** `poll: if ( SECOND )`  
`{ SECOND = FALSE;`  
`update display }`

count is not shared

shared variable  
initial value = FALSE

Jan 7, 2014

40

## Recall (Half of) Stream–2–Pipe Example:



- suppose streams and pipe are services by h/w ISRs:
  - **ISRA** – receives a Data packet of Stream A data
  - **ISRB** – receives a Data packet of Stream B data
  - **ISRP** – transmits Pipe packets

Jan 7, 2014

41

## Stream–2–Pipe Communication


- ISRs share a queue (**Packet\_Q**) to exchange packets
- **ISRA** and **ISRB** produce packets as they are received
- when packet of data is received it is put in **Packet\_Q**
- **ISRP** consumes packets by transmitting them
- when ISRP is idle, it gets a packet from **Packet\_Q**
- instance of classical **producer/consumer** problem

used widely to illustrate  
operating system issues

Jan 7, 2014

42

## Stream-2-Pipe: Pictorial Representation



Jan 7, 2014

43

## Issues to Expose: SYNCHRONIZATION

- among concurrent activities
- e.g. transmit on pipe cannot proceed without data from streams
  - pipe transmission must wait for work
- frequent requirement in concurrency!

Jan 7, 2014

44

## Issues to Expose: Buffer Management

- how do ISRA and ISRB obtain empty packet buffers for receiving packets?
- what does ISRP do with an empty packet buffer after transmitting a packet?
- static vs. dynamic schemes?
- what happens if no buffers/memory available?

Jan 7, 2014

45

## Issues to Expose: INTERFERENCE

### Potential for INTERFERENCE:

- concurrent activities share **Packet\_Q**

**INTERFERENCE** occurs when simultaneous concurrent activities corrupts a shared resource

- modification is concurrent with “other” access

Jan 7, 2014

46

## Critical Sections

- a region of code that has the potential to cause interference is called a **critical section**
- the existence of a critical section does not guarantee interference – often depends on specific access sequences and timing
- interference may not show up in testing !
  - hard to debug!

Jan 7, 2014

47

## Example: consider a static array implementation of **Packet\_Q**

circular Q: (data structure)

- Head and Tail pointers (indices)
- remove @ Head
- Tail points to next available array element
- when reach end of array, wrap to start:
 
$$\text{index} = (\text{index} + 1) \bmod Q\_size$$

Jan 7, 2014

48



## Data Declarations

```
Q_Size = ***** ;           // some constant
Packet_Q :
  array [ 0 .. Q_Size - 1 ] of packet_buffer ;
Head : integer ; // index of packet to remove
Tail : integer ; // index of next free array element
Count : integer ; // # of packets in Packet_Q
```

SHARED data!

Jan 7, 2014

49

## Initial Values & Empty() Method

Initially:

```
Head = 0;
Tail = 0;
Count = 0;
```

```
boolean Empty() { return ( Count == 0 ); }
```

Jan 7, 2014

50

## Add Method

```
Add ( P : packet_buffer )
{ if Count >= Q_Size
  { /*exception! Q full!*/ exit ; }
  Packet_Q [ Tail ] = P ;
  Tail = ( Tail + 1 ) mod Q_Size ;
  Count = Count + 1 ;
}
```

NOTE: puts P in Q before  
adjusting Tail or Count !

Jan 7, 2014

51

## Remove Method

```
Remove ( var P : packet_buffer )
{ // assume Count > 0
  P = Packet_Q [ Head ];
  Head = ( Head + 1 ) mod Q_Size;
  Count = Count - 1;
}
```

NOTE: removes P from Q before  
adjusting Head or Count

Jan 7, 2014

52

## Scenario

in a uniprocessor implementation, suppose:

- ISRA and ISRB finish receiving packets at approx. the same time
- independent reception – no interference
- both may attempt to access **Packet\_Q.Add** concurrently
- accessing shared resource!

Jan 7, 2014

53

## Add Method Details

- suppose ISRA calls **Add** first and is executing:  
**Packet\_Q** [ Tail ] = P<sub>A</sub> ;  
Tail = ( Tail + 1 ) mod Q\_Size ;
- suppose the compiled implementation of the 2<sup>nd</sup> line is:  
temp = Tail<sub>old</sub> ; // temp might be a register  
temp = temp + 1 ;  
temp = temp mod Q\_Size ;  
Tail<sub>new</sub> = temp ;

Jan 7, 2014

54

## ISRB Interrupts ISRA!

- suppose ISRA has executed:  
    **Packet\_Q** [ Tail<sub>old</sub> ] = P<sub>A</sub> ;  
    temp<sub>A</sub> = Tail<sub>old</sub> ;  
    and is about to execute:  
    temp<sub>A</sub> = temp<sub>A</sub> + 1 ;  
    when an interrupt occurs and ISRB  
    begins to run

Jan 7, 2014

55

## Data Corruption!

- when ISRB runs, ISRA has placed a packet in **Packet\_Q**, but has not yet modified Tail and Count
- ISRB will overwrite the packet just added by ISRA, then adjust Tail, and then increment Count
- when ISRA resumes it will finish adjusting Tail<sub>old</sub>, and then increment Count

Jan 7, 2014

56

## Interference!

- net result: (after both ISRs complete)
- **lost packet** P<sub>A</sub> originally added by ISRA  
    – overwritten by P<sub>B</sub> added by ISRB
  - Tail is still correct (for the packets in Q)  
    but Count is **corrupted** (too large by one)
  - Are there other interference problems?

Jan 7, 2014

57

## Other Potential Interference

- **Add / Remove** concurrently  
    – potential interference with Count
- concurrent **Add** when only one space left in **Packet\_Q**  
    – both calls could pass the “full” test before incrementing Count  
    – overwrite a valid packet & increment Count beyond Q\_Size

Jan 7, 2014

58

## Race vs. Interference

- **race:** two concurrent activities have begun the process of accessing a shared resource
- one activity will get there first!
- a race is due to sharing resources, but a race (by itself) does not corrupt the resource
- race conditions are a **common occurrence** in event-driven systems

Jan 7, 2014

59

## Critical Section Protection

- ensure mutually exclusive access to relevant shared resource(s)

### Uniprocessor Solution:

- disable interrupts while processing critical sections
- keep critical sections short!
- which interrupts should be disabled?
  - all?
  - only those with potential to interfere?

Jan 7, 2014

60

## Uniprocessor Solution

Common solution:

```
disable;
critical section    // protected!
enable;
```

e.g.


```
disable;
Packet_Q.Add( myP )
enable;
```

Jan 7, 2014

61

## What about a Multiprocessor Solution ?

- recall stream-2-pipe example:
  - suppose the ISRs are implemented on independent processors & share memory
  - disabling ints on one processor won't stop interrupts on other processors!



Jan 7, 2014

62

## Multiprocessor Solution

- use busy waiting and shared variables to ensure mutual exclusion
  - busy waiting ☹
  - wastes CPU time!
- keep critical sections short ☺
  - minimize wasted time

Jan 7, 2014

63

## Busy Waiting (Version 1)

share a boolean variable Busy

TRUE == resource is busy

FALSE == resource is available

**Lock** ( var Busy : boolean )

```
{ while ( Busy ) { } // wait until available
  Busy = TRUE;      // indicate resource busy
}
```

Jan 7, 2014

64

## Busy Wait (version 1)

- PROBLEM!** non-atomic Lock!
- more than one processor could pass busy wait loop before setting Busy = TRUE
- each would proceed assuming mutually exclusive access to resource

```
while ( Busy ) { }
Busy = TRUE;
```

both processors  
could reach here  
before either sets  
Busy = TRUE

Jan 7, 2014

65

## Busy Wait (Version 2)

- use h/w enforced atomic operation to read and modify Busy
- Test-And-Set **TAS**
- functional syntax:
 

```
old_value TAS ( variable, new_value )
```
- returns original value of variable (old\_value), and sets variable to new\_value
- typically locks system bus for duration of instruction

Jan 7, 2014

66

## No Problem!

(as long as hardware supports TAS ☺)

```
myLock ( var Busy : boolean )
{
  while ( TAS ( Busy, TRUE ) ) { }
}
```

↑  
atomic operation

- Software-only solutions (no TAS) also exist for multiprocessor systems  
e.g. Lamport's bakery algorithm

Jan 7, 2014

67

## Summary of Motivation (1)

- **concurrency** has inherent difficulties:
  - potential for **interference**
  - need for **synchronization** of activities
  - need for **communication** among activities
  - race conditions (event-driven reality!)

Jan 7, 2014

68

## Summary of Motivation (2)

- concurrent activities can arise in the **requirements** of an application
  - i.e. the system must support more than one input/output relationship concurrently
- concurrency in an **implementation** is the result of design decisions

Jan 7, 2014

69

## Concurrency-Related Issues (1)

- **mindset:**  
~~sequential (polling)~~ vs.  
event-driven (interrupts, multiprocessor)
- **priority:** some activities are high-priority, while others have lower-priority
- **h/w:** determines extent of concurrent capabilities of components

Jan 7, 2014

70

## Concurrency-Related Issues (2)

- **culture:** "we do it this way here"
  - legacy
  - tools at hand
- designer's **artistic creation**
  - experience, problem solving
  - "on a previous project, a similar problem was solved by . . . "

Jan 7, 2014

71

## What SYSC 5701 Is ....

- concerned with using a process model to help reduce the development challenges for real-time applications
- primary concern: designer's perspective!
- simplifying the implementation of concurrency
- hide some machine details
- use "standard" process model
- simplifying the mapping of concurrency in requirements onto concurrency in implementation

Jan 7, 2014

72

## **Lamport on Concurrency (2009)**

“Education is not the accumulation of facts. It matters little what a student knows after taking a course. What matters is what the student is able to do after taking the course. I've seldom met engineers who were hampered by not knowing facts about concurrency. I've met quite a few who lacked the basic skills they needed to think clearly about what they were doing.”

Jan 7, 2014

73

## **So ... Why are you Here?**

- **IF**

Education is not the accumulation of facts. It matters little what a student knows after taking a course. What matters is what the student is able to do after taking the course.

- **THEN:**

What will you be able to do after completing a graduate degree?

What do you think a professor would answer?

Jan 7, 2014

74