

Real Time and Embedded Systems

by

Dr. Lesley Shannon

Email: Ishannon@ensc.sfu.ca

Course Website: <http://www.ensc.sfu.ca/~Ishannon/courses/ensc351>



Simon Fraser University

Slide Set: 9

Date: November 22, 2011

Slide Set Overview

- Priority
 - Inversion and Inheritance

Priority Inversion and Inheritance

Priority

- When threads have different priority levels, the higher priority threads will run first
 - This is good
- Unfortunately, sometimes a low-priority thread consumes all the available CPU time, blocking the higher-priority thread from running
 - This is known as priority inversion
 - This is bad
- How does it happen?

Priority Inversion

- Let's assume we have C1 with priority 5, C2 with priority 15, and a Semaphore with priority 20

Priority Inversion

- Let's assume we have C1 with priority 5, C2 with priority 15, and a Semaphore with priority 20

- We can fix this with priority inheritance

Priority Inversion

- We can fix this with priority inheritance
- Basically, the semaphore will inherit the priority of the low priority client thread C1 and drop to priority 5

- Now we have a new problem

Priority Inversion

- Now we have a new problem, even with priority inheritance
- Let's assume the same threads using, priority inheritance

Priority Inversion

- This is another form of priority inversion- blocking the higher priority thread without consuming any CPU time

Priority Inversion

- Now we have a new problem, even with priority inheritance
- Let's assume the same threads using, priority inheritance

Priority Inversion

- This is another form of priority inversion- blocking the higher priority thread without consuming any CPU time

Priority Inversion

- Solution: boost the server's priority to the highest priority of all blocked clients
- Good news:
 - In some systems, you don't have to do this (e.g. QNX does it automatically for you when you use message passing – but only one client- to – server deep)
 - The second server inherits the first server's real priority
 - We then have a priority inversion problem again (Why?)
 - In Linux, there is also a PI-futex patch that provides "rt_mutex"
 - The futex call is "normal"
 - The rt_mutex call supports priority inheritance

Questions?

- What is priority inversion
- How do we solve priority inversion?