

# CS452/652 Real-Time Programming Course Notes

Daniel M. Berry, Cheriton School of Computer Science  
University of Waterloo

# Three Levels of Task Switching

1. High: Unbounded number of processors
2. Middle: Unbounded number of processes, one CPU, and kernel routines
3. Low: Unbounded number of processes, one CPU, and kernel task


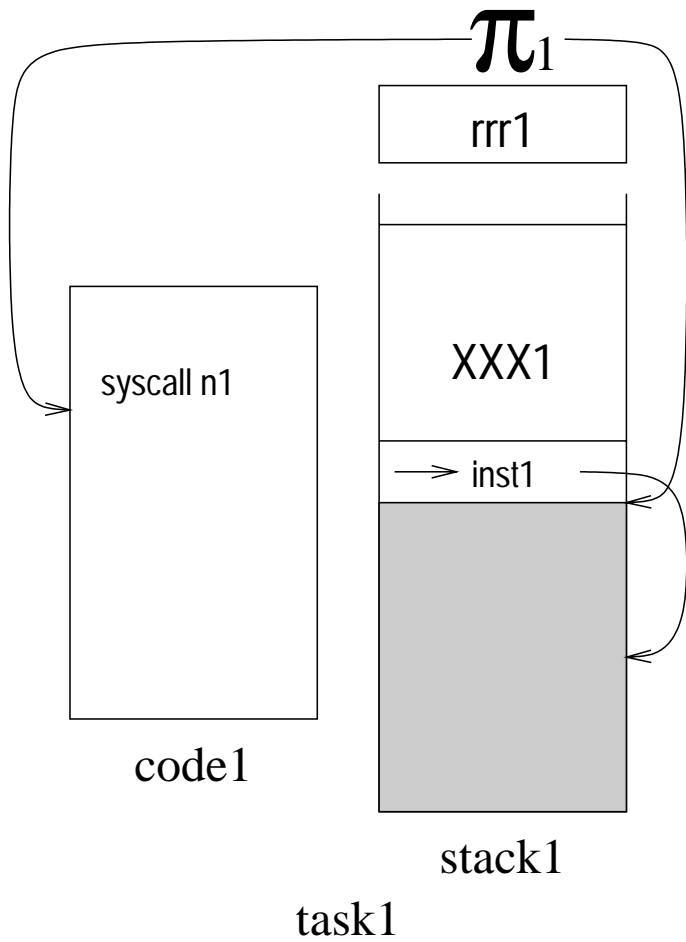
# High Level View

Unbounded number of processors:

A system service is either

- instantaneous, like an ordinary instruction, or
- blocking, causing the executing processor to be blocked.

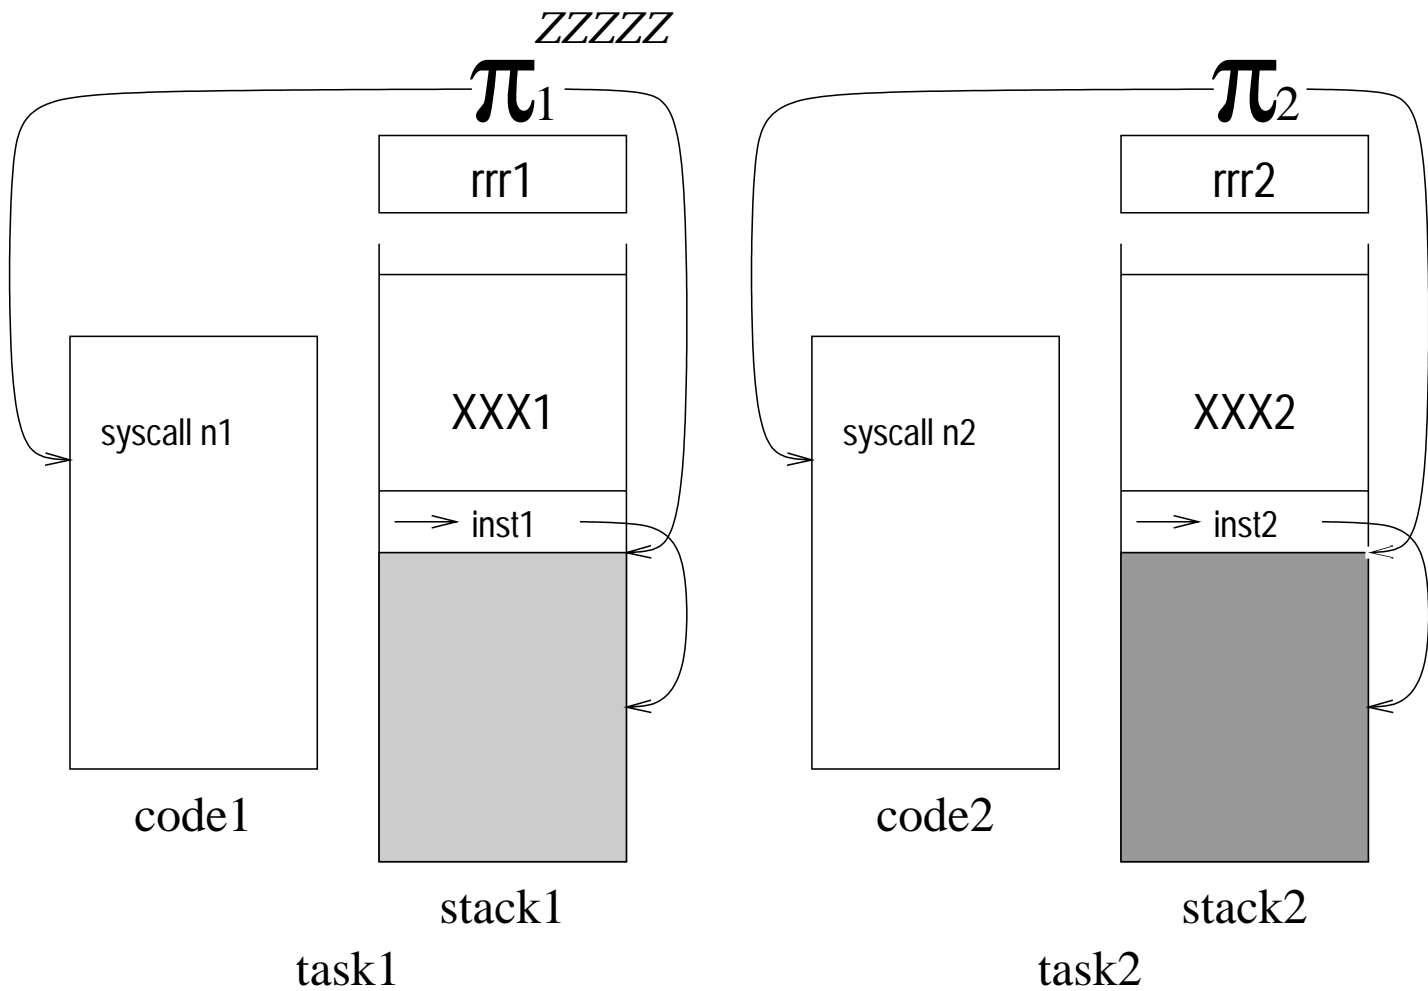
Distinction between running and ready does not exist.

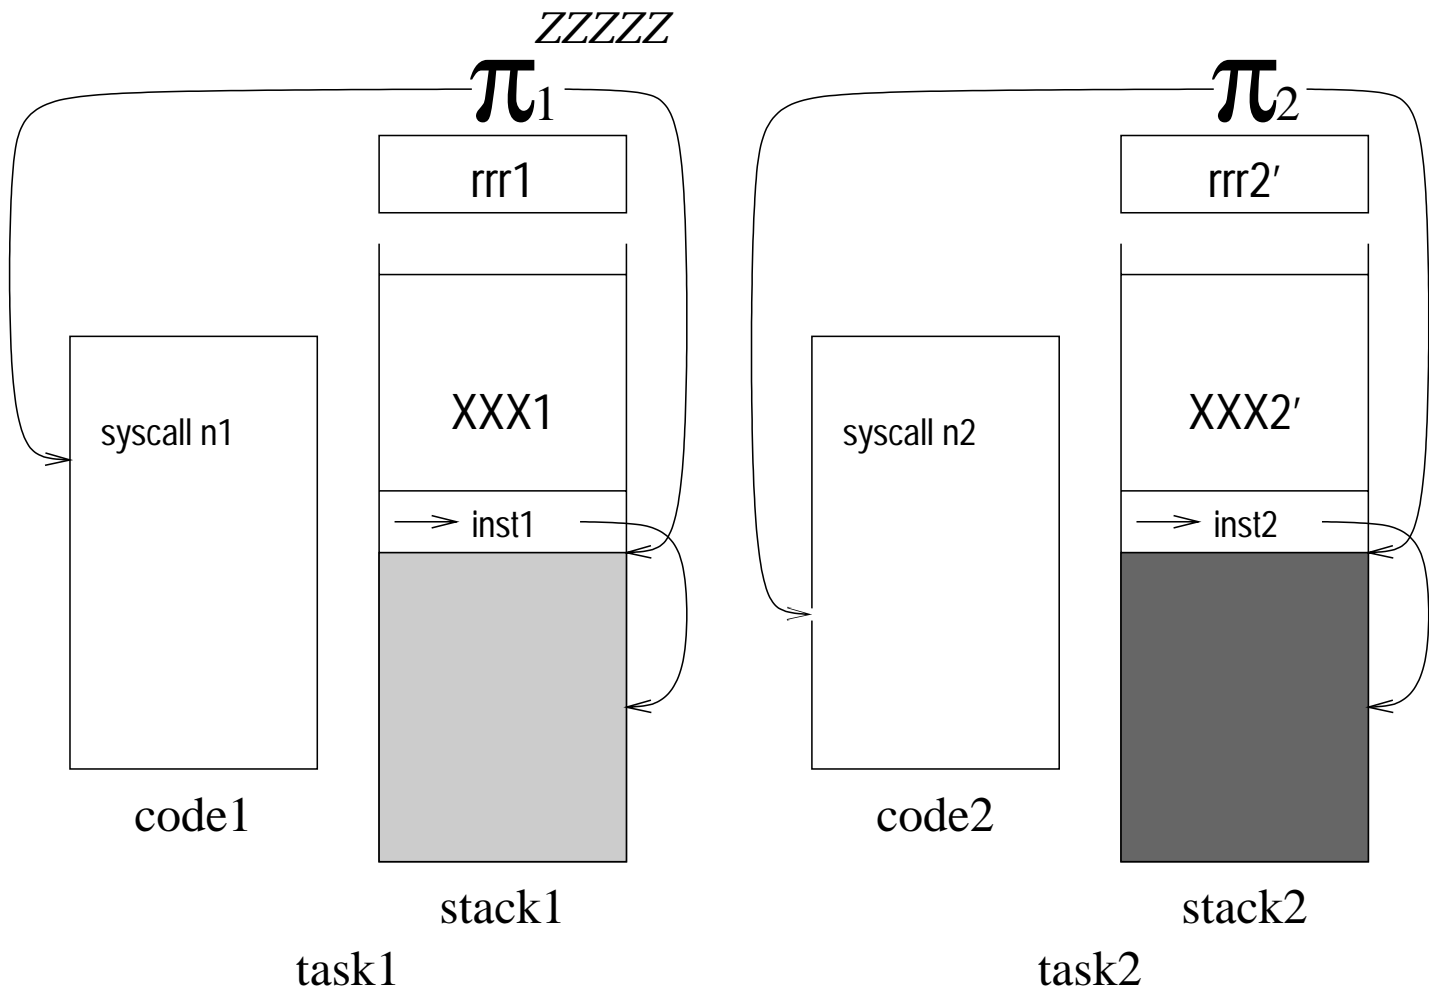


# High Level View, Cont'd

In the previous snapshot, both processors are running.

Next two snapshots shows one processor **blocked** (asleep) and the other running.





# Middle Level View

Unbounded number of processes, one CPU, and kernel routines

That is, a system service is done in a non-task kernel that is made up of procedures that can be called by the CPU.



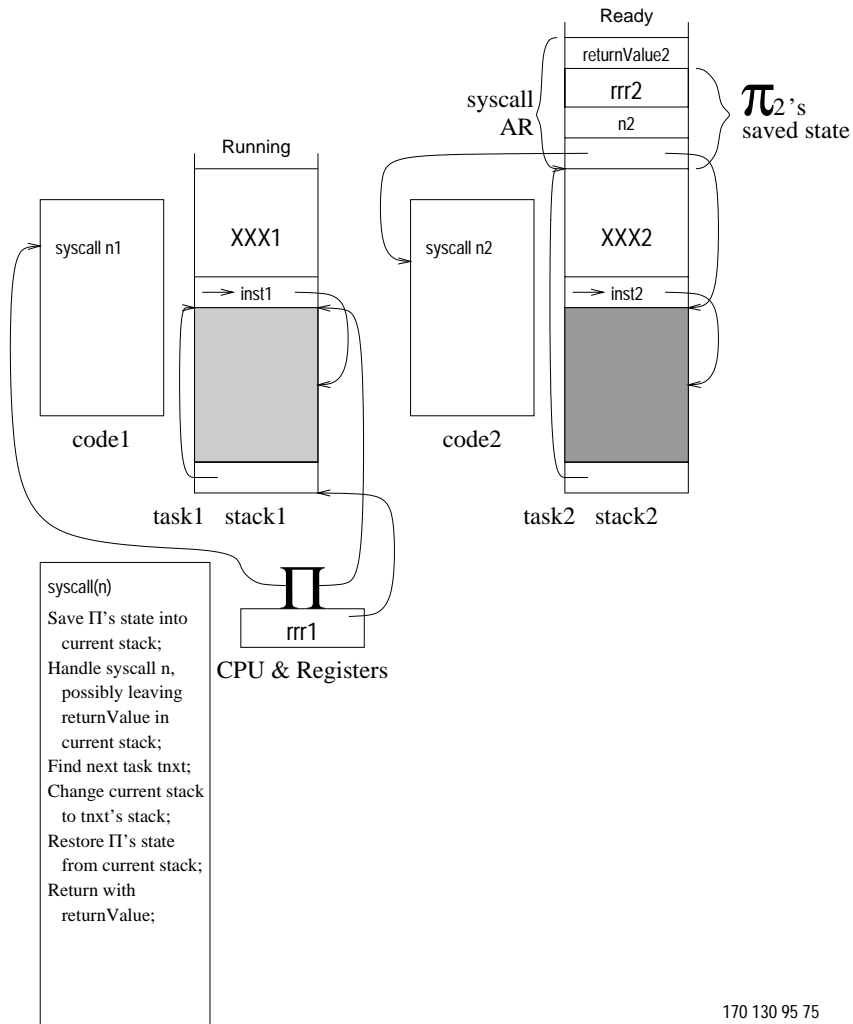
# Middle Level View, Cont'd

Task switching is done as a side effect of making a `syscall`, which does both the requested service and switches to the task with the highest priority, usually leaving the calling task blocked or ready.

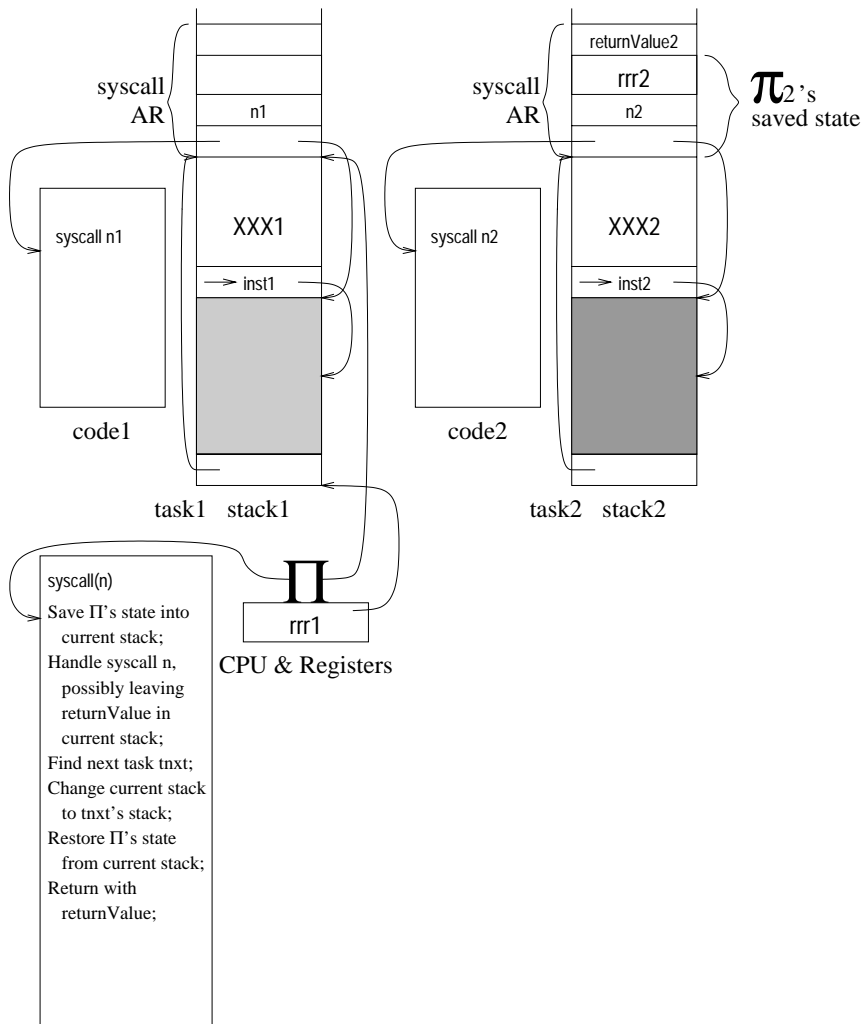
It is possible that if the calling task were to be left ready, it might be chosen as the task to switch to, but only if it is the task with the highest priority.

# Middle Level View, Cont'd

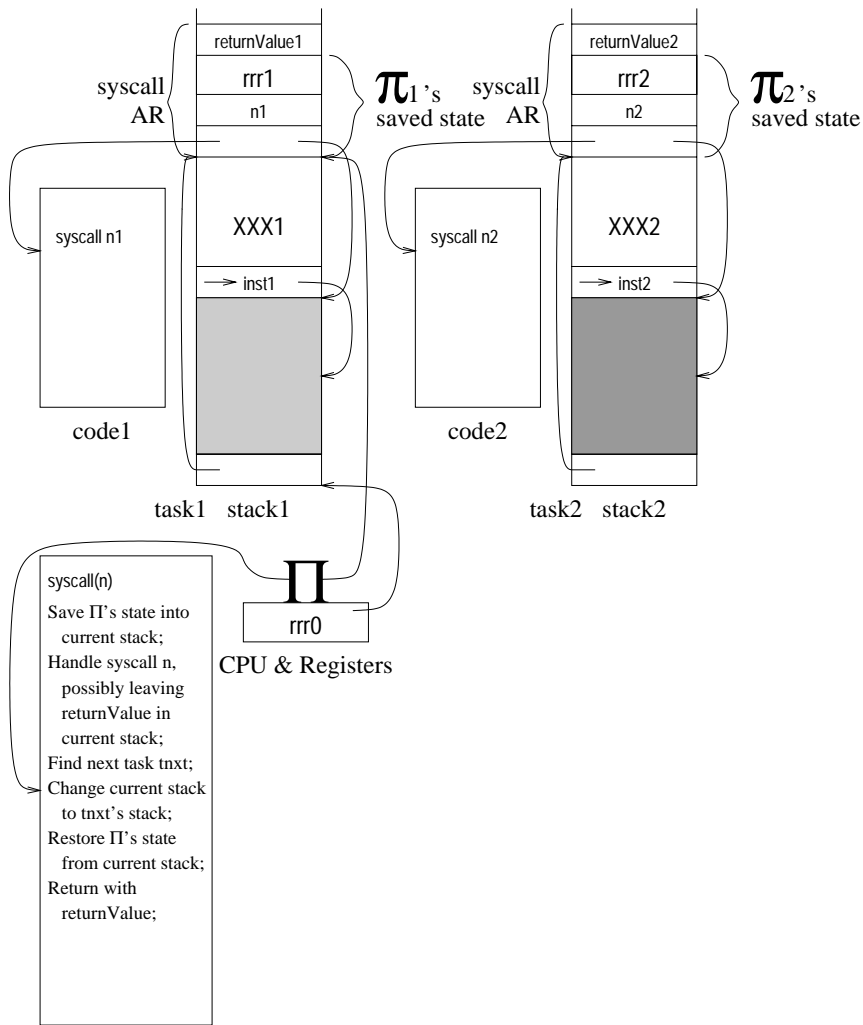
Note that any task that is not running is left **blocked** or **ready** in the middle of the **syscall** routine. So task switching is done by changing the stack and base pointers while keeping the instruction pointer pointing into the **syscall** routine.



Context Switching Code



Context Switching Code



Context Switching Code

# “possibly leaving returnValue”?

Why does the instruction say “Handle syscall n, *possibly leaving* returnValue in current stack;” instead of just “leaving”?

# “possibly ...”, Cont’d

If the `syscall` is an instantaneous kind, then handling `syscall n` will definitely leave `returnValue` in current stack.

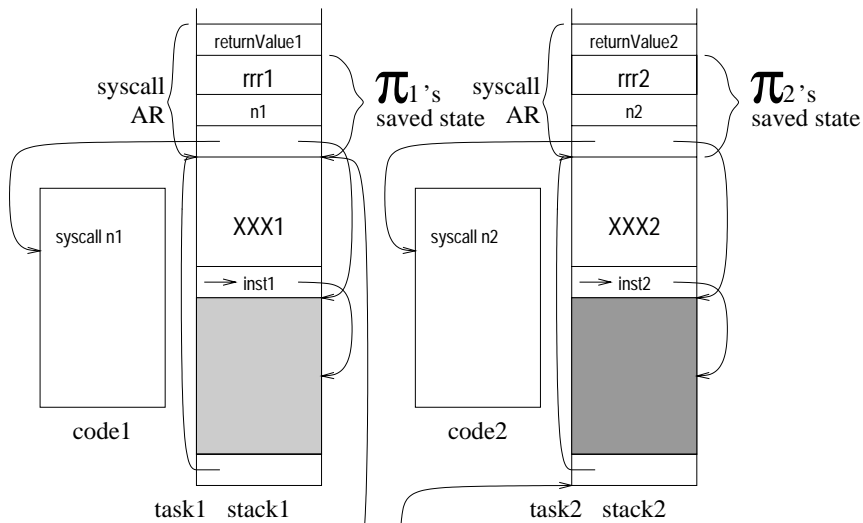
## “possibly ...”, Cont’d

If, on the other hand, the `syscall` is a blocking kind, then handling `syscall n` results in arranging that some later execution of `syscall` will discover that the data requested in this `syscall` are ready, that its value should be deposited at the top of *this stack*, via a pointer associated with the expected value, and that *this stack* should be changed from `blocked` to `ready`.



# “possibly ...”, Cont’d

The difficulty of keeping track of these pending completions of `syscalls` through multiple invocations of `syscalls` by one CPU, each with its own activation record, is the reason that it is nice to make the kernel a task that keeps its own data.

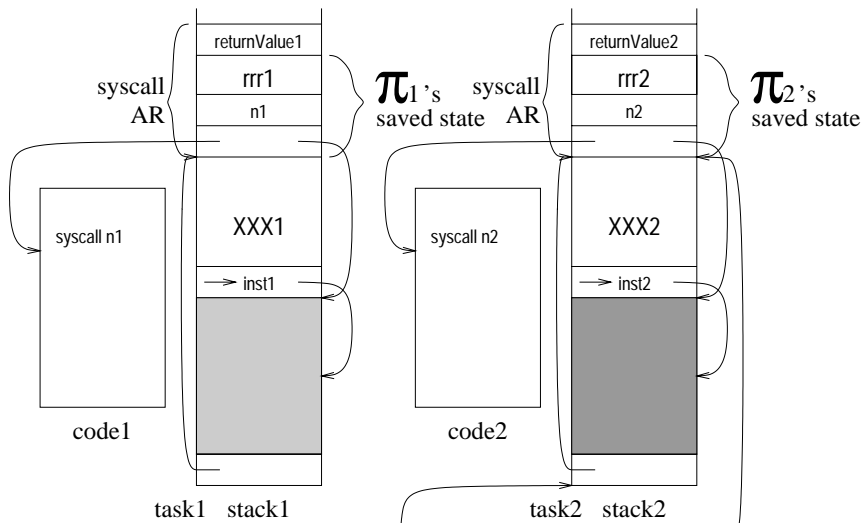


```

syscall(n)
Save  $\Pi$ 's state into
current stack;
Handle syscall n,
possibly leaving
returnValue in
current stack;
Find next task txnt;
Change current stack
to txnt's stack;
Restore  $\Pi$ 's state
from current stack;
Return with
returnValue;

```

Context Switching Code

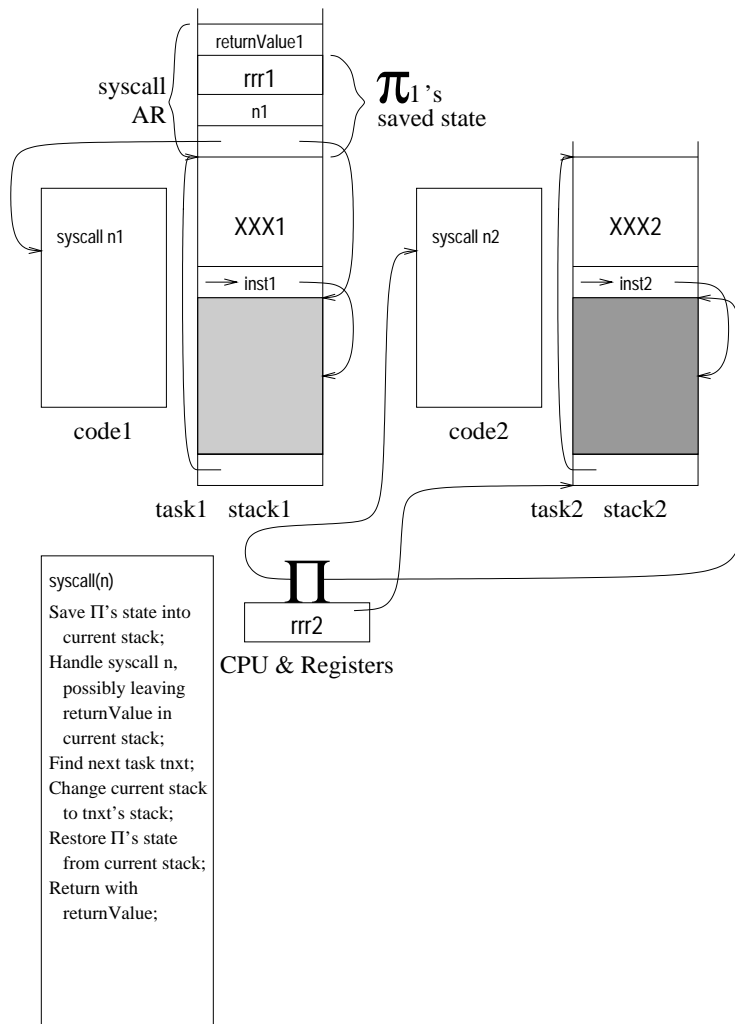


```

syscall(n)
Save  $\Pi$ 's state into
current stack;
Handle syscall n,
possibly leaving
returnValue in
current stack;
Find next task txnt;
Change current stack
to txnt's stack;
Restore  $\Pi$ 's state
from current stack;
Return with
returnValue;

```

Context Switching Code



Context Switching Code


# Low Level View

Unbounded number of processes, one CPU, and kernel task

That is, a system service is done in a kernel task.

# Low Level View, Cont'd

I tried to decompose the context switching code into the three pieces suggested by the explanation given of the diagram:



for the transition from task1 to the kernel and from the kernel to task2, but it did not quite work.

# Low Level View, Cont'd

Note that what I show is only *one* possible decomposition of the behavior.

Others will work, in particular e.g., making only one procedure instead of three.

# Low Level View, Cont'd

First, a view of the context switching and the kernel code.

