
Concurrency


Notes for CSC 100 - The Beauty and Joy of Computing
The University of North Carolina at Greensboro

What is concurrency?

From Wikipedia page Concurrency (computer science):

*In computer science, **concurrency** is a property of systems in which several computations are executing simultaneously, and potentially interacting with each other. The computations may be executing on multiple cores in the same chip, preemptively time-shared threads on the same processor, or executed on physically separated processors.*

Traditional computer model



One Central Processing Unit (CPU)


Basic Property: Can do *one task (or thread) at a time*

What if want to do multiple tasks?

- Web browser running
- Movie player running
- MS Word running to take notes
- Checking for instant messages
-

Question: How do we do all these things simultaneously?

Traditional computer model



One Central Processing Unit (CPU)

Basic Property: Can do one task at a time

Answer: Time-Division Multiplexing

Switch back and forth so quickly that it looks like things are happening simultaneously.



If that's "traditional", what's current?

What are some common CPUs on the market?

From newegg.com, Oct 2, 2012




If that's "traditional", what's current?

What are some common CPUs on the market?


Question: What do the circled parts actually mean?

From newegg.com, Oct 2, 2012



Multicore model

Just like having multiple independent CPUs sharing a common memory



Now multiple tasks can run at the same time!
 Can still do time-division multiplexing - might run dozens of threads concurrently

Clearly multicore is the "new normal" - why?

Moore's Law


Prediction: Transistor count per chip doubles every two years

- Due to Gordon Moore (Intel co-founder)
- Prediction from 1965
 - Has stayed remarkably accurate




Many versions of Moore's Law, dealing with

- Transistor count (this is the real Moore's Law)
- Processing speed
- Storage capacity



Different CPU characteristics


Clock speed and single-thread performance have flattened out!



Graph from Dan Garcia's Lecture Notes (UC Berkeley)

Different CPU characteristics

More transistors - even if can't make faster, we can duplicate cores with all these extra transistors!



Other kinds of concurrency

Multiple processors (even multi-core) in a single computer

Multiple computers in a rack with high speed communication (cluster computing)

Multiple computers or clusters geographically separated (grid computing)

And even more possible configurations...

Beyond this class - just be aware that there are more possibilities!

Parallel algorithms

Previous examples: Running several independent tasks (solving several independent problems) simultaneously.



Parallel algorithms make use of multiple processors working on a single problem.

Ideally, with P processors you'd like to solve problems P times faster, but:

- Some tasks are *inherently sequential*, and must be done one step at a time (e.g., maze searching - not everything that seems inherently sequential is though!)
- Some tasks are *embarrassingly parallel*: obviously independent sub-tasks with little inter-task communication (e.g., a lot of image processing, graphics rendering, brute force crypto attacks, ...)
- Most tasks are somewhere in between these two extremes (but might have some embarrassingly parallel components!)


An example - a big sum

Think "big" even if only 20 numbers here - imagine a million numbers



What to notice...

This algorithm has a "parallel part" and a "sequential part"
More processors makes parallel part faster, but not sequential part!



Note: Sometimes "Sequential Time" can increase as number of processors increases, since there's more to coordinate. In last example:

- With 4 processors: sequential time was 4 steps because 4 partial results
- 100 processors really speeds up parallel time, but gives 100 partial results.

Dangers - Race Conditions

Race conditions occur when multiple threads are working on the same data, and the interleaving of individual steps causes problems.


A (only slightly contrived) BYOB example:

What if BYOB didn't allow you to add a specific item to a list, but only allowed you to add space and then fill in that space using "replace" (this is actually what happens behind the scenes!). Then you might want to define something like:

```
add pItem to end of pList
script variables @NewIndex
set @NewIndex to length of pList + 0
append empty item to pList
replace item @NewIndex of pList with pItem
```

Dangers - Race Conditions - cont'd

Next: You use your block in a game with Alonzo and the dragon. Each character "registers" with a characters list, as follows:



After we click the green flag to start the game, both characters do the add, and this is the resulting list:



Question: Can you explain why isn't Alonzo in the list?

Dangers - Race Conditions - cont'd

Consider what Alonzo does, step by step...

Alonzo's sequence of steps




Dangers - Race Conditions - cont'd


Consider what Alonzo does, step by step...

Then the dragon starts a fraction of a second after Alonzo...

Alonzo's sequence of steps



Dragon's sequence of steps



Dangers - Race Conditions - Solutions?

Can we rearrange things to solve the problem?

```
add pItem to end of pList
script variables sNewIndex
append empty item to pList
set sNewIndex to length of pList
replace item sNewIndex of pList with pItem
```

What's good: setting sNewIndex is closer to its use in "replace item..."

But: **This does not fix the problem!** "Closer" makes it *harder* to have a race condition, but it's still possible - and if it's possible to fail, it will fail at the worst possible time (see Murphy's Law).

Real solution: Use *locks* to give a thread temporary exclusive access to data (like the character list) - all other threads must wait.

Dangers - Deadlock

Deadlock is when threads hold resources that other threads need, but don't have everything they need to make progress.




Image from CS Unplugged project

Example: A banking program where to transfer money from account A to account B, records for A and B are locked to avoid race conditions before updating balances.

