Lecture 1

January 8, 2008

Required background: computer architecture, programming, data structures, discrete math

Dictionary definitions:

Concurrent - 1. operating or occurring at the same time 2a. running parallel 3. acting in conjunction 4: exercised over the same matter or area by two different authorities

Concurrency - simultaneous occurence

Working definitions:

Sequential program - sequence of actions (statements) that produce a result (value). Variously called a *process* or *task* or *thread* (*of control*)

Concurrent program - two or more sequential programs, cooperating to accomplish some goal, and running (at least conceptually) at the same time.

Concurrent programming - the art and science of creating concurrent programs

Hardware terms: *uniprocessor, shared-memory multiprocessor, multicomputer* (separate memories)

Parallel program - a program intended to exploit the capabilities of a parallel computer to achieve increased performance

What makes concurrent programming worthwhile? What makes it challenging? Why do we study it as an independent topic?

- exponential state growth the blessing and the curse of concurrent programming. On the one hand CP provides compact and natural representation of real-world world problems with complicated state, yet on the other we have to reason about that complicated state to understand our concurrent programs
- interference
- common patterns and paradigms for dealing with the ensuing complexity
 - concurrency mechanisms

- programs
- reasoning
- true parallelism for performance

What situations give rise to concurrent programs?

- operating systems device drivers, time sharing
- embedded systems concurrency *plus* real-time
- distributed systems
- client-server systems
- user interface programs
- re-use existing applications (e.g. Unix pipes)
- true parallelism

We are going to ignore true parallelism as a motivation during much of this class. The techniques we study will be applicable to true parallelism but finding and exploiting parallelism in problems is different from what we will be doing. At the end we will revisit parallel programming.

State Explosion

Consider the program

x = 1; y = 2; z = 3

How many states does it have?

- discuss

Now consider (// means concurrently with, which implies arbitrary interleaving of steps of the components)

x = 1; y = 2; z = 3 // a = 1; b = 2; c = 3

How many states does it have?

- discuss

If there are n states in each of m concurrent processes, the program has n^m states.

Now consider

x = x+1

How many states?

- discuss

To answer this we have to know something about how such a statement would be compiled, e.g. in a simple accumulator machine it might be

LD x,A ADD 1 ST x,A

Now consider

x = x+1 / / y = x+1

How many states? If x is initially 0 what is the final value of y?

- discuss

As you can see, writing correct concurrent programs is going to be very hard if understanding what a tiny given program does is this difficult! So why isn't the whole notion of a concurrent programming a folly?

- discuss

The world is complicated – many states must be represented; if these states evolve along more-or-less independent trajectories, a representation as concurrent threads is more compact, and *easier* to understand and reason about than the alternative. (Which is not to say that the reasoning is *easy*!)

What we ask then of our systems for reasoning about concurrent programs, our programming notations and mechanisms for concurrent programming, and our patterns and paradigms for concurrent programs is that they help us manage the complexity that arises due to arbitrary interleaving and interference.

Plan for the semester: we will begin by working with *declarative concurrency* in which many of the thornier problems do not occur, but which has limited expressive power. We then move on to *message passing concurrency* where the our goal is to see how to ensure that the complications of concurrency don't impact parts of programs where they are really unnecessary, and finally we will look at *shared state concurrency* where the problems are in ones face much of the time.

Reading and problems for next time:

- Read Chapters 1 and 2. These provide necessary background for what follows. I will not talk about Chapter 1 explicitly. Start to read Chapter 3 which I will begin to lecture on next time.
- Retrieve the errata list for the textbook from http://www.info.ucl.ac.be/~pvr/bookerrata.html and mark all the corrections in your book.