Slide 1

Administrivia

- One purpose of the syllabus is to spell out policies (next slides).
- Most other information will be on the Web, either on my home page (here, office hours) or the course Web page (next slide). (here).

  A request: If you spot something wrong with course material on the Web, please let me know!

Slide 2

Course Web Page

- "Course Web page" (site?) is meant to point you to pretty much all information for the course — readings, assignments, etc.
- You can find it via TLearn, or via the link from my home page (findable from the department's Web site, `http://www.cs.trinity.edu` or — probably? — by doing a Web search on my name).

Slide 3

Course FAQ

- "What will my grade be based on?" (See syllabus.)
- "What happens if I can't turn in work on time, or I miss a class?" (See syllabus.)
- "What's your policy on collaboration?" (See syllabus.)

Slide 4

Course FAQ, Continued

- "When is the next homework due?" (See "Lecture topics and assignments" page.)
- "When are your office hours?" (See my home page.)

  Note that part of my job is to answer your questions outside class, so if you need help, please ask! in person or by e-mail.

Slide 5

Course FAQ, Continued

- Why are we using "my" book when there are books that are more textbook-like, and also more recent? because I think it comes closer than any other book I know to covering the material I think is perhaps best learned from a book.
- (And no, I'm not getting rich off the royalties! though extra income is not unwelcome.)

Slide 6

Course FAQ, Continued

- "What computer(s) can I use to do homework?"

  Easiest option may be department's Linux machines. You should have physical access via your TigerCard to all the classrooms and labs (probably not today but soon) any time the building is open. You should also be able to log in remotely to any that are booted into Linux, or to a cluster of Linux-only machines in ITS's server room (names `diasnn`, where `nn` ranges from 01 to 05). To log in from off-campus, we are currently recommending that you use ITS's VDI.

Slide 7

A Little About Me

- Up until last spring I hadn't thought about telling students a little about myself on the first day. A colleague mentioned that she does, so — okay!
- Short version of biography: Undergrad degrees from UT Austin, math and Plan II. More than ten years in what we now call IT. Back to school for master's and PhD in computer science. Two years as a postdoc, then at Trinity since Fall 1999.
- I teach a variety of courses, but currently focusing more on courses "close to the machine". My research area (sadly neglected for some years) is parallel computing.
- (What do I do for fun? well . . .)

Slide 8

What is Parallel/Distributed Computing?

- Some computational jobs are just too much for one processor — no way to get them done in reasonable time.
- For jobs done by people, what do you do when the job is too much for one person?

Slide 9

What is Parallel/Distributed Computing?

- For jobs done by people, if too much for one person you assign a team — but you have to figure out
  - How to divide up work among team members.
  - How to coordinate activities of team members.
- Same idea applies to computing — if too much for one processor, use multiple processors. Issues are similar — how to divide up work, how to coordinate.

Slide 10

Simple Examples

- "People job" examples:
  - Digging a hole.
  - Building a house.
  - Baby.

  What do you notice about the last one in particular?
- Computer examples:
  - Adding up a lot of numbers.
  - Computing Fibonacci numbers.

  But these don't seem too "big" . . .

Slide 11

How Much Calculating is "A Lot"?

- Dated examples from computational biology — how many operations per second are needed to get things done fast enough to be useful?:
  - Sequence the genome — 10^{12} ops/second (500 2-Gigahertz processors).
  - Protein/protein interactions — 10^{14} ops/second (25,000 4-Gigahertz processors).
  - Simulating whole-body response to a drug — 10^{16} operations/second (1,250,000 8-Gigahertz processors).
- (Source — Intel's former life sciences industry manager.)

Slide 12

What Are Some Other Hard Problems? Traditional HPC

- Crash simulation / structural analysis.
- Oil exploration.
- Explosion simulations (why Los Alamos is interested).
- Astrophysics simulations (e.g,, Dr. Lewis's work on Saturn's rings).
- Fluid dynamics.
- "Rendering" for computer-generated animation.
- And many others . . .

Slide 13

What Are Some Other Hard Problems? Recent

- "Big data"!
- AI, machine learning, what Dr. Hibbs and Dr. Lewis do . . .

Slide 14

The Need for Speed

- Solving the same problems faster — reducing wall-clock time.
- Solving bigger problems.
- Solving problems more exactly — to get better answers, need more detail, hence more processing.

Slide 15

Can't You Just Get a Faster Computer?

- Up to a point — yes. Moore's law predicts that number of transistors on a die roughly doubles every 1.5 years. Until recently, that meant doubling processor speed and memory. (Over 30 years, that's a factor of about a million!)
- But . . .

Slide 16

Can't You Just Get a Faster Computer? continued

- As you know — however fast processors are, it's never fast enough, and faster is more expensive.
- Eventually we'll run into physical limitations on hardware — speed of light limits how fast we can move data along wires (in copper, light moves 9 cm in a nanosecond — one "cycle" for a 1GHz processor), other factors limit how small/fast we can make chips.
- Maybe we can switch to biological computers or quantum computers, but those are pretty big paradigm shifts . . .
- In the past few years, chip makers are still able to put more transistors on a chip, but they seem to have run out of ways to exploit that to get more speed, and are instead making chips with multiple processing elements ("cores" for computational chips, other elements in GPUs).

Slide 17

"The Answer" — Parallel Computing

- Analogous to "team of people" idea — if one processor isn't fast enough, use more than one.
- Also useful when there's something "inherently parallel" about the problem — e.g., operating systems, GUI-based applications, etc.
- `http://www.top500.org` tracks fastest computers; for many years now all have been "massively parallel".

Slide 18

But I Don't Want To Solve Problems Like Those!

- What if you aren't interested in solving problems like these "grand challenge" problems, Is there still a reason to be interested in parallel computing?
- The hardware is there, and it's becoming mainstream — multicore chips, general-purpose computing GPUs, etc. (The Intel person says "the chip makers can put more and more transistors on a chip, and this is the best way to use that.")

  To get best use of it for single applications, will probably need parallelism.
- Also, for some applications, thinking of them as parallel/multithreaded can lead to a solution that lets you do something useful while waiting for I/O, etc.

Slide 19

About the Course

- Can think of this course as the equivalent of CS1 for parallel (and to some extent concurrent and distributed) programming. As with CS1, many things to learn all at once:
  - A new "box of tools" — or several boxes of tools (different languages/libraries/paradigms). Must learn syntax/functions, plus tools such as compilers and runtime systems.
  - How to use the stuff in the box of tools to solve interesting problems — from low-level "what is this syntax good for?" to algorithm design.
  - How to think about "does it work?"
  - How to think about "how fast is it?"
- Also as with CS1, the idea will be to teach a mix of technical skills and basic concepts, with emphasis on learning by doing.

Slide 20

Minute Essay

- (Most lectures will end with a "minute essay" — as a quick check on your understanding, a way for me to get some information, etc., and also to track attendance. Just put your answer in the body of the message; no Word documents please, and put "minute essay" and the course in the Subject line.)
- What are your goals for this course?
- Are you reasonably comfortable with C? How about C++? (There will be assignments using one or the other.) How about Java? In the past I've had students do an assignment in Java, but this year maybe it should be Scala?
- I'm told that our CS2 course now includes some exposure to multithreaded programming? What do you remember from that, and do you have other experience with any kind of parallel or distributed programming?
- Anything else you want to tell me? about the course, about what you did over the summer . . .