

## Topic Notes: Introduction and Overview

Welcome to CSIS 400!

---

### Why Parallel Computing?

Parallel computing sounds simple enough – when one computer isn't powerful enough to solve your problem, use more than one.

Before we start to think about how to use parallelism on a computer, let's think about a parallel approach to solving a “real-world” problem.

- Taking a census of Loudonville.

One person doing this would visit each house, count the people, and ask whatever questions are supposed to be asked. This person would keep running counts. At the end, this person has gathered everything.

If there are two people, they can work concurrently. Each visits some houses, and they need to “report in” along the way or at the end to combine their information. But how to split up the work?

- Each person could do what the individual was originally doing, but would check to make sure each house along the way had not yet been counted.
  - Each person could start at the town hall, get an address that has not yet been visited, go visit it, then go back to the town hall to report the result and get another address to visit. Someone at town hall keeps track of the cumulative totals. This is nice because neither person will be left without work to do until the whole thing is done. This is the *master-slave* method of breaking up the work.
  - The town could be split up beforehand. Each could get a randomly selected collection of addresses to visit. Maybe one person takes all houses with even street numbers and the other all houses with odd street numbers. Or perhaps one person would take everything west of Route 9 and the other everything east of Route 9. The choice of how to divide up the town may have a big effect on the total cost. There could be excessive travel if one person walks right past a house that has not yet been visited. Also, one person could finish completely while the other still has a lot of work to do. This is a *domain decomposition* approach.
- Grading a stack of exams. Suppose each has several questions. Again, assume two graders to start.

- Each person could take half of the stack. Simple enough. But we still have the potential of one person finishing before the other.
  - Each person could take a paper from the “ungraded” stack, grade it, then put it into the “graded” stack.
  - Perhaps it makes more sense to have each person grade half of the *questions* instead of half of the exams, maybe because it would be unfair to have the same question graded by different people. Here, we could use variations on the approaches above. Each takes half the stack, grades his own questions, then they swap stacks.
  - Or we form a *pipeline*, where each exam goes from one grader to the next to the finished pile. Some time is needed to start up the pipeline and drain it out, especially if we add more graders. These models could be applied to the census example, if different census takers each went to every house to ask different questions.
  - Suppose we also add in a “grade totaler and recorder” person. Does that make any of the approaches better or worse?
- Adding two  $1,000,000 \times 1,000,000$  matrices.
    - Each matrix entry in the sum can be computed independently, so we can break this up any way we like. Could use the master-slave approach, though a domain decomposition would probably make more sense. Depending on how many processes we have, we might break it down by individual entries, or maybe by rows or columns.

In each of these cases, we have taken what we might normally think of as a *sequential* process, and taken advantage of the availability of *concurrent processing* to make use of multiple workers (processing units).

Parallelism adds complexity (as we will see in great detail), so why bother?

- we want to solve the same problem but in a shorter time than possible on one processor – goal: speedup
- we want to solve larger problems than can currently be solved at all on a single processor – goal: scale-up
- some algorithms are more naturally expressed or organized concurrently
- and now: that’s where performance gains come from in modern processors!

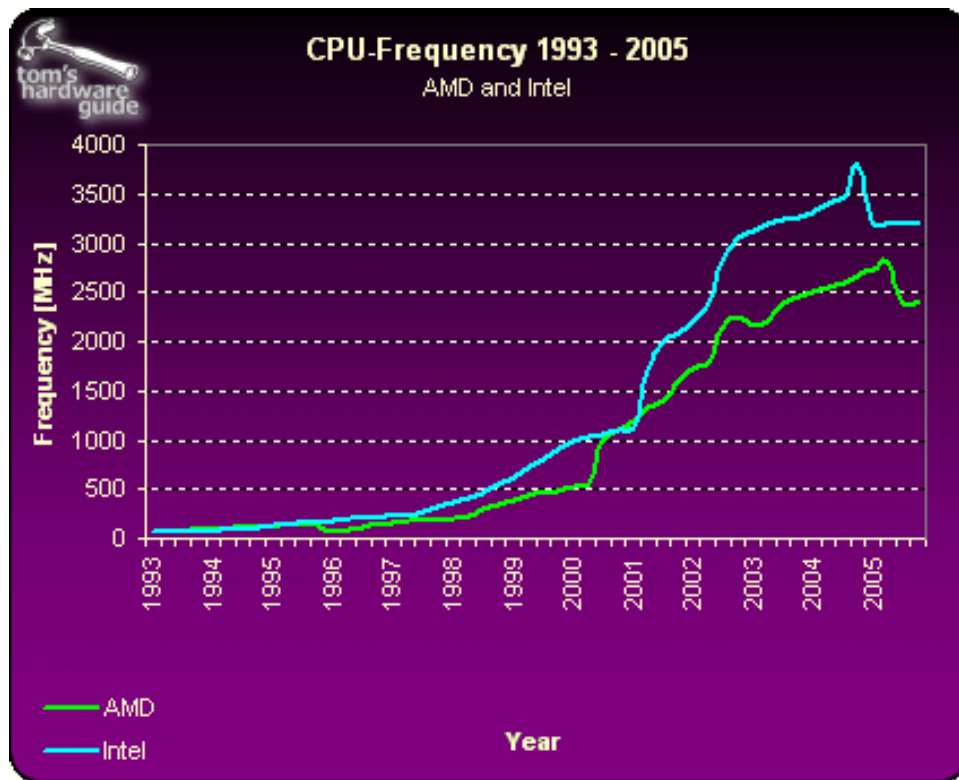


Figure used with permission from article *The Mother of All CPU Charts 2005/2006*, Bert Töpelt, Daniel Schuhmann, Frank Völkel, Tom's Hardware Guide, Nov. 2005, [http://www.tomshardware.com/2005/11/21/the\\_mother\\_of\\_all\\_cpu\\_charts\\_2005/](http://www.tomshardware.com/2005/11/21/the_mother_of_all_cpu_charts_2005/)

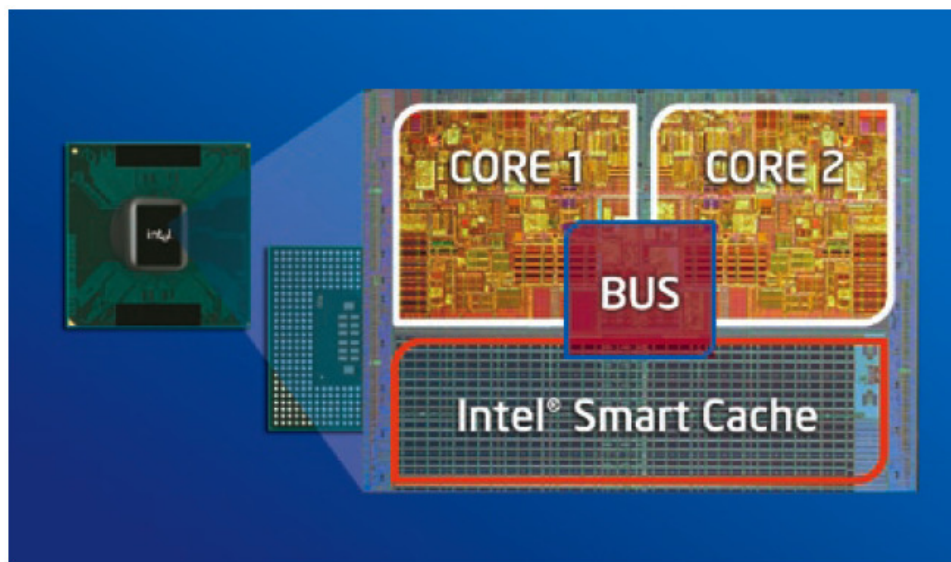


Image from Intel Core Duo Processor product brief.

## Some Basics

*Sequential Program*: sequence of actions that produce a result (statements + variables), called a process, task, or thread (of control). The state of the program is determined by the code, data, and a *single* program counter.

*Concurrent Program*: two or more processes that work together. Big difference: *multiple* program counters.

To cooperate, the processes need *communication* and *synchronization*, which can be achieved through *shared variables*, or *message passing*

We will consider all of these

---

## Hardware to run concurrent processes

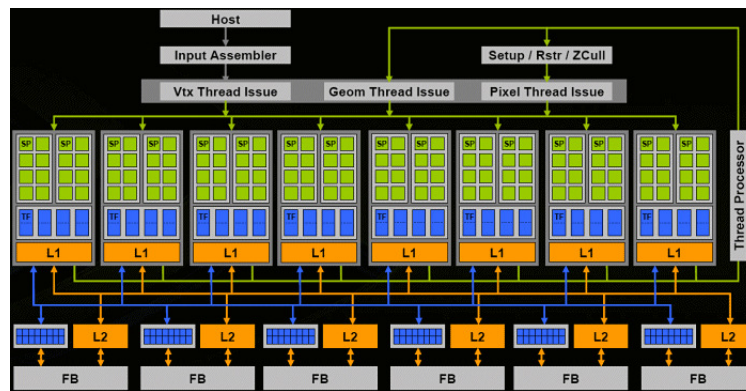
- single processor – logical concurrency (see Operating System course)
- multiprocessor – shared memory
- multicomputer – separate memories
- network – slower communication

Computers may be classified as:

- SISD: single instruction, single data – one processor doing one thing at a time to one piece of data at a time.
- SIMD: single instruction, multiple data – multiple processors all doing the same thing at the same time, but operating on different data. Also known as: vector computers. Program operates in “lock step” on each processor.
- MIMD: multiple instruction, multiple data – multiple processors each doing their own thing.
- SPMD: single program, multiple data – not really a classification of the computer, but of a model used to program a MIMD computer. Multiple processors run the same program, but do not operate in lock step. Also known as the “interacting peers” model. This is the model we will use most in this class.

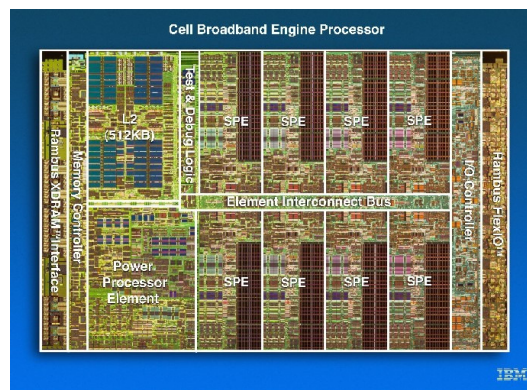
Some examples:

- SISD: Pre-”multi-core” desktops and laptops.
- SIMD: graphics cards that apply a single operation to an array of data points at the same time.



NVIDIA image

- MIMD: desktops and laptops with multiprocessors or multi-core chips – each processor can be executing any instruction and operating on any data.
- MIMD: Cell architecture (Sony PS3) – one general purpose processor and several special-purpose cores.



[http://www.research.ibm.com/cell/cell\\_chip.html](http://www.research.ibm.com/cell/cell_chip.html)

- MIMD: ASCI Red, Sandia National Labs: 4600+ nodes, each with 2 Intel Pentium II Xeon processors, first TeraOp machine in 1997.
- Hybrid: IBM Blue Gene systems – dense clusters of Cell processors.

See <http://www.top500.org/>.

Moral: from the desktop to the world’s largest supercomputers, it’s a world of parallel processing out there!

## How to Achieve Parallelism

- We need to determine where concurrency is possible, then break up the work accordingly

- This is easiest if a compiler can do this for you – take your sequential program and extract the concurrency automatically. This is sometimes possible, especially with fixed-size array computations.
- If the compiler can't do it, it is possible to give “hints” to the compiler to tell it what is safe to parallelize.
- But often, the parallelization must be done explicitly: the programmer has to create the threads or processes, assign work to them, and manage necessary communication.

We will consider all of these types of parallelism here.