Final

CSE 30321 Computer Architecture I							
Niemier, Michael Thaddeus	01	52	12/13/2010	10:30 AM	12:30 PM	127	Hayes-Healy Center

If you haven't done so already, do your CIF

Final is open book, open notes

What to take away from last 4-5 lectures

Historical sources of performance gain

- Pipelining
- Technology Scaling
- Memory Hierarchies
- Hardware support for threads
- …
- All the easy stuff done
  - (And to some extent being taken away!)

Technology drive to multi-core computing

Why the clock flattening? POWER!!!!

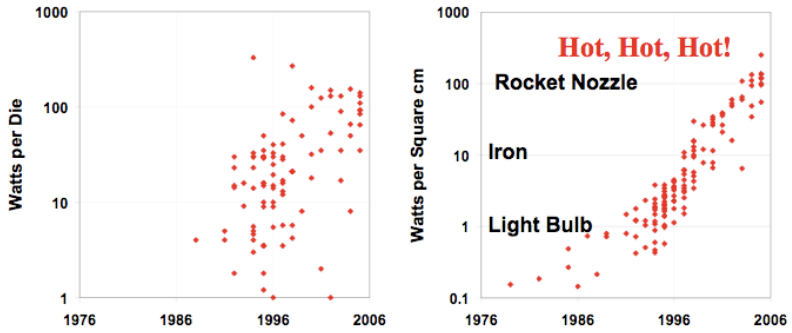

The multi-core solution:

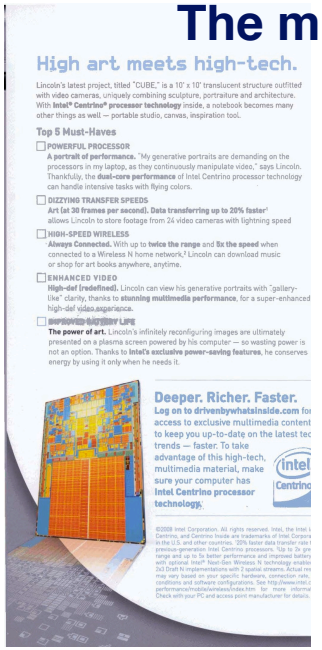

Fundamentally, built on datapath, memory hierarchy discussed in class.

- Processor complexity is good enough
- Transistor sizes can still scale
- Slow processors down to manage power
- Get performance from...

Parallelism

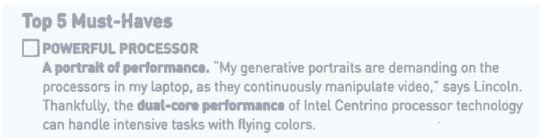

(i.e. 1 processor, 1 ns clock cycle
vs.
2 processors, 2 ns clock cycle)

Would like to do deliver speedups equal to the number of cores

- Time(1) / Time(P) for P processors
  - note: must use the best ***sequential*** algorithm for Time(1); the parallel algorithm may be different.

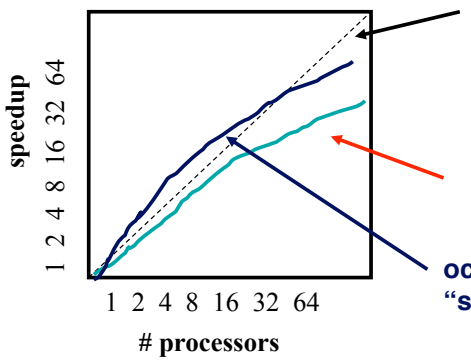

"linear" speedup (ideal)

typical: rolls off w/some # of processors

occasionally see "superlinear"... why?

Could program P processors as if your off for a nice pleasant hike…

But you should always be "bear aware"…

What are "the bears" ???

Impediments to Parallel Performance

(*independent of problem/algorithm itself*)

- Contention for access to shared resources
  - i.e. multiple accesses to limited # of memory banks or shared cache may dominate system scalability
  - The interconnection network itself

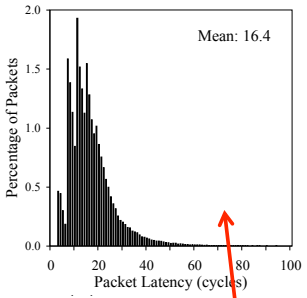

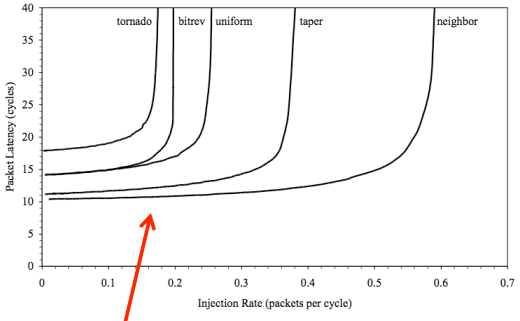

Cannot necessarily presume a uniform communication time form core-to-core

Impediments to Parallel Performance

(*independent of problem/algorithm itself*)

- **Latency**
  - **Multiple clock cycles to cross chip or to communicate from core-to-core even in absence of contention**

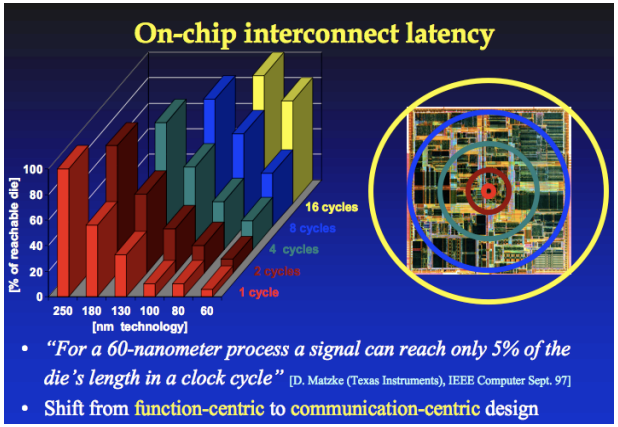

Impediments to Parallel Performance

(*independent of problem/algorithm itself*)

- **Cache coherency**
  - **P1 writes, P2 can read → need consistent data**
    - **Protocols can enable $ coherency but not free**

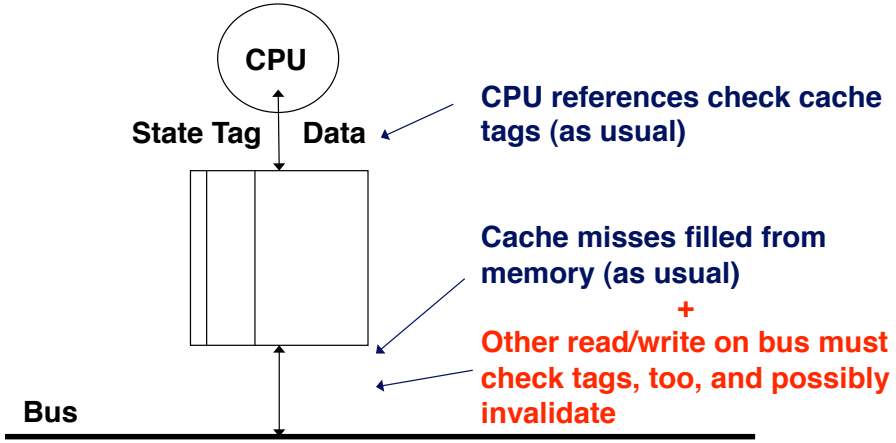

Impediments to Parallel Performance

(*independent of problem/algorithm itself*)

- **Programming languages, environments, & methods:**
  - **Need simple semantics that can expose computational properties to be exploited by large-scale architectures**

What if you write good code for a 4-core chip and then get an 8-core chip?

- i.e. tune code to deal with communication latencies, etc.

Impediments to Parallel Performance

(*independent of problem/algorithm itself*)

$$\text{Speedup} = \frac{1}{\left[1 - \text{Fraction}_{\text{parallelizable}}\right] + \frac{\text{Fraction}_{\text{parallelizable}}}{N}}$$

- All "bears" adversely affect $\text{Fraction}_{\text{parallelizable}}$ and hence speedup
- Should have quantitative sense of overhead when writing software
  - Tricky because it can change from chip-to-chip!

Computer Architecture…

Parallel software…

Can keep bears at bay, but need to think hard about how

We ran into a bear along the way…