Multiprocessors in One Box

4-28-2003

Opening Discussion

- What did we talk about last class?
- Have you seen anything interesting in the news?

Humorous Pitfall

- "Pitfall: Moving functions from the CPU to the I/O processor, expecting to improve performance without a careful analysis."
  - This is a bit funny to see as the current trend is actually to go in the other direction. With all the transistors we have on a chip we are moving more things closer into the processor for better performance and to reduce costs.
  - This is a nice example of how things change with time in our field.

Major Questions of Parallel

- Do they share memory?
  - This is the biggest architectural distinction. Shared memory machines have multiple processors in one box. Programmers can use a single address space.
- How do the different processors coordinate?
  - Do they used shared variables or networked messages?
- How many processors are there?

Memory Model

- Programming wise, shared memory means we have to synchronize data access, typically with locks.
- If all processors can get to any memory in the same amount of time it is a uniform memory access (UMA) multiprocessor. If a processor can get to some memory faster than others and if differs by processor it uses nonuniform memory access (NUMA).

Shared Memory Schematic

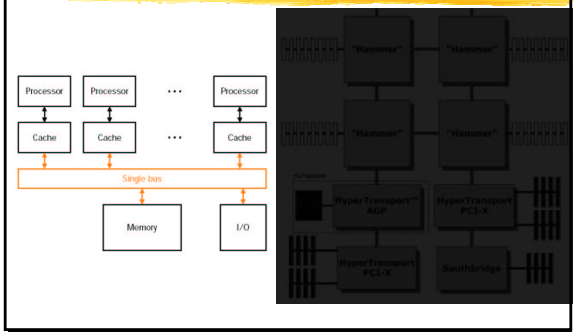

Cache Coherency

- Snooping is one way of dealing with this where every cache is listening to the bus to see if data that it holds is being written.
  - Write-invalidate: when a word is written in one cache, any other cache holding that word invalidates their version. It monitors for others getting read misses on it.
  - Write-update: new data is broadcast to all processors which update their versions if they have it.

FSM for Write-Invalidate

- Read Only state means it hasn't been written.
- Read/Write state means it has been written.
- Invalid means it must be read before use.

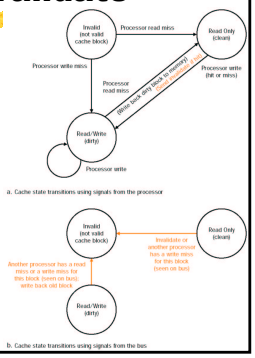

Minute Essay

- Have you ever done any shared memory programming? This type of task is likely to become much more common in the future. What do you think that we should do as computer scientists to help deal with this?