

Multiple Processors in Many Boxes

4-30-2003

Opening Discussion

- What did we talk about last class?
- Have you seen anything interesting in the news?
- Have you ever heard of the site top500.org? They keep a list of the fastest computers in the world. The list is worth looking at, especially noting what the computers are like.

Minute Essay Comment


- Last time I asked you what you thought we needed to do as computer scientists about the fact that we will be seeing a lot more multiprocessors. One person commented that we should be making things easier not harder. That means we want tools (like compilers) that deal with the complexities for us. A lot of research has gone into parallel compilers, but progress hasn't been astounding.

Problems with a Single Bus

- Putting all the processors on a single bus and in one box has many advantages. In particular, it is fast. You can have true shared memory all with fast access and communication latency is really low.
- The biggest disadvantage of putting everything on one bus is expandability. The bus must be fairly short which means your maximum number of processors is limited. There are also bandwidth problems.

Distributed Memory

- Instead of using a bus to give all the processors access to the same memory, we can put them across a network with distributed memory.



Message Passing


- In a distributed memory environment we don't have shared variables to use for communication. Instead we have to send messages across the network when one processor needs to talk to another one.
- This is inherently slower, but it also serves as automatic synchronization so that we don't have to worry about locks or similar shared memory protections.

Shared, Distributed Memory

- Some distributed memory systems act to the programmer like they are shared memory systems. The programmer interacts with what seems to them to be a single large address space and the underlying system deals with getting the memory from other machines. This introduces a different set of coherency problems that HPC systems solve with directories.

Relative Performance

- Bandwidth limitations on buses cause the performance of multiple processors on a single bus to level off.
- As a result, there is a point of optimal price/performance.



Clusters and Grids

- What we have seen since just before the time your text was published is an incredible growth in the use of clusters of computers for parallel computation.
- Clusters have multiple full fledged computers connected by networks. All are "off the shelf" products which greatly reduces the cost.
- Grids are now becoming more popular and utilize even more distributed resources.

Minute Essay

- What did we talk about today? Would you be willing to learn a new language and paradigm if it gave better support to parallel processing?
