

Programming Models Breakout

With Architecture Comments!!!

What does the general MC community need? A better MPI? Better threads? Streaming programs?

Can't just look at HPC. But we can draw on our expertise. Lots of cross over if we look for it. What support can we ask for.

We can help industry figure out 8, 16 cores. And tell them about the pitfalls at 8000, 16000. Demonstrate the growth path from 8 to 10000.

Stir...Stir...Stir...

Transport & memory power
& area dwarf FP

- Si on a CPU is mainly memory
- DDR Pads are about the size of a core & not scaling
- Chip BW not scaling, number of contacts flatlining

perspectives

Control" (CPU/Pull)

Memory/Push)

Only because of Intel/
AMD's scale!
Man-Millennia

- *Get rid of all things called "CPUs" because they are "free"*

Should processors be the "name" for thread placement? or Memory?

ly need both perspectives

"MIPSWars" and "FLOPSWars" are not a good metric. Drive CPUs in odd directions. Need better metrics!

Stuff

- *Diagnostics*
- *Memory Synchronization*
- *Programmer access to enhancements*
- *Smart"er" memory controller*
- *Memory Hierarchy*

Diagnostics

- *Instrumentation*

Putting in a counter is easy... if you know what you are counting. Also, there is impact for process context switching (much less per thread!).

- *user, runtime, and tool accessible and useable*

- *programmer hints*

- *Hotspots, races, cross thread data corruption, memory mismanagement*

Some are hard for HW to detect. (Can't measure program intention). HW would need more data (tagged memory?)

Memory Synchronization

Need Weak ordering as well to scale.

- *AMOs* Good.
- *F/E bits and friends like transactions*
- *Order-less Synchronization*

Programmer access

- *Memory hints (cache and like)*

Smart'er' memory controller

- *Scatter/Gather*
- *AMOs*
- *User programmable*
- *Virtualization*



We Like these.

Memory Hierarchy

- *Abstraction of hierarchy*
- *Cache aware/oblivious*
- *Use local memory in a more transient way*
- *Handle memory hierarchy expansion*

We like these too.

- How do YOU want to handle FLASH or other NV/Slow memory? (block addressing, streaming, Load/Store)? They can look more like disks.
- Access characteristics require different program model.
- How do we make it “user” visible? (not just OS)