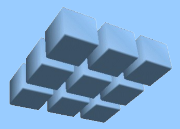

Building a Multicore Multiprocessor (MCMP) Supercomputer

Cmpware, Inc.

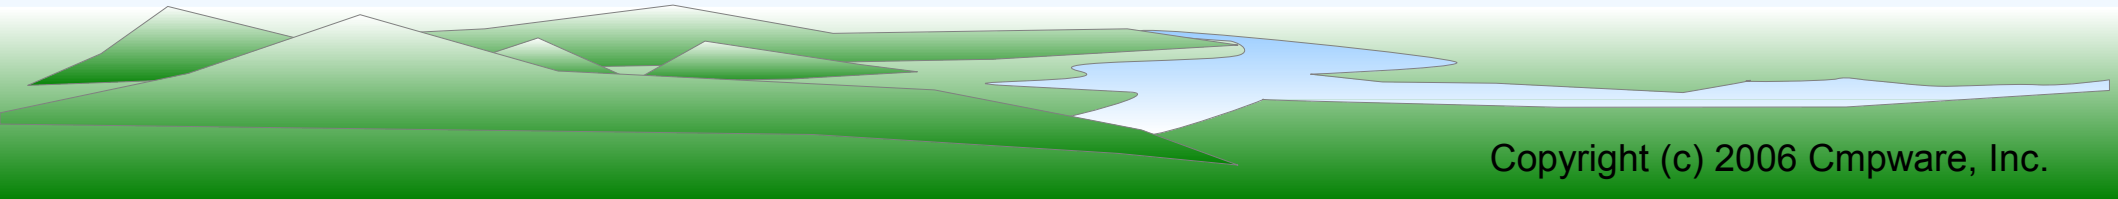

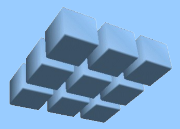

Supercomputers 2006

- Mostly multiprocessors
- Few Cray-style vector machines
- Approximately 10,000 processors
- Approximately 100 TFLOPS Linpack
- Standard desktop CPUs (80W)
- Power estimate: approx. 1 MegaWatt (!)
- See: http://www.top500.org/

Multicore Multiprocessing (MCMP)

- Hundreds of CPUs on a die possible
- High performance: Thousands of MFLOPs
- Simple HW design
- Simple HW verification
- Flexible
- Scalable
- Reprogrammable
- ***Reduces power***

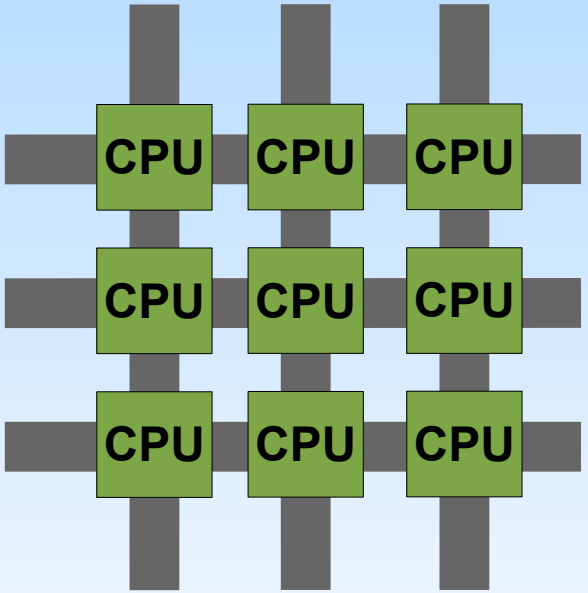

Power and MCMP

Desktop CPU

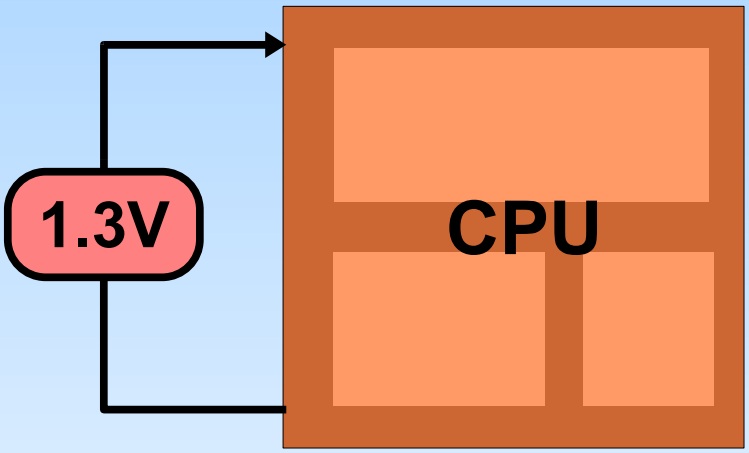

2 GHz

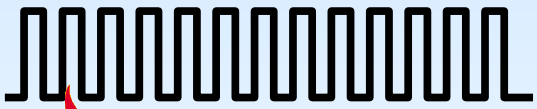

75 Watts

CMP Architecture

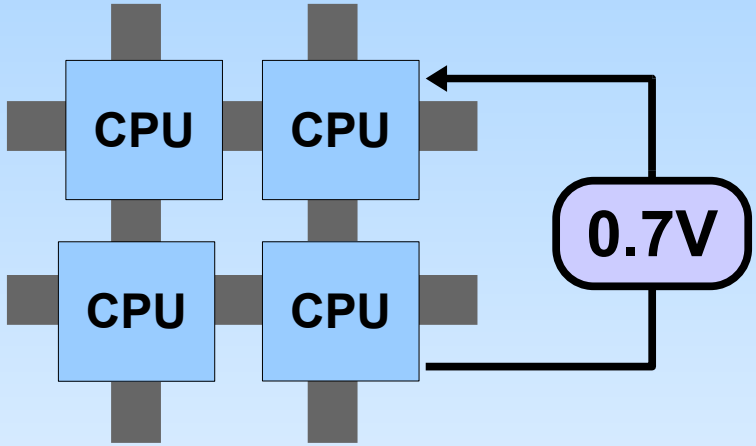

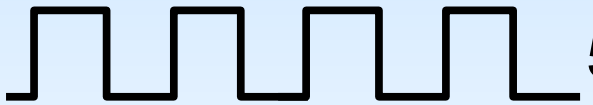

500 MHz

0.3 Watts

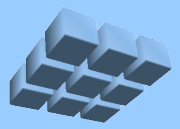

MCMP Devices (2006)

- Many CPU cores available (Arc, Arm, Tensilica, MIPS, PowerPC, others)
- Typical example: Arc 700 family
  - 30K gates CPU + 30k gates FPU = 60k gates
  - 0.15 mW/MHz @ 500 Mhz = 75 mW
  - *1000x* less power than desktop CPUs
  - *Estimate*: Assuming 60K gates @ 5 transistors per gate and 500 M transistors per die, 2x for on-chip memory ==> **833 CPUs**
  - 833 x 75 mW = **62W** per device

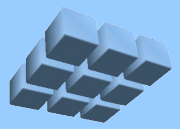

Desktop CPUs

- Dedicated to the serial, uniprocessor model
- Millions of transistors used for small, incremental performance gains over decades (*superscalar*, *reordering buffers*, *out of order execution*, *various caches*, etc ...)
- Getting data on and off chip can take 100s of cycles. Context switches can take 1000s.

==> *Once uniprocessor model broken, these millions of transistors may be better deployed as extra processors for calculations*

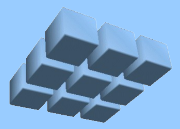

MCMP System Level Parameters

100 TFLOPS

15.5 KiloWatts

$25k

- 833 x 500 Mhz = 400+ GFLOPS per device
- 100 TFLOPS = 250 devices (peak)
- 250 device x 62W = 15.5 KWatts
- MCMP cost: $100 x 250 ==> $25k total

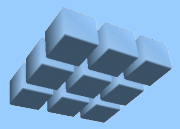

Open Issues

- Memory bandwidth
  - Can memory keep processors busy?
  - Do CMP devices permit new techniques?
  - Is this better than existing desktop CPUs?
- Software
  - How easily will existing applications run / port?
  - Software development tools
- Device and System Architecture
  - How to size and connect CMP processors?

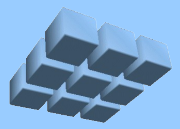

The Cmpware CMP-DK

- Useful for experimenting with multiprocessor and communications architectures
- Useful for parallel algorithm development
- Quickly build and program multiprocessors
- Redefine multiprocessor in seconds
- Speeds simulation (2M+ instructions / sec)
- Complete Eclipse development environment
- Standard models (Sparc, MIPS32 and more)

The Cmpware CMP-DK

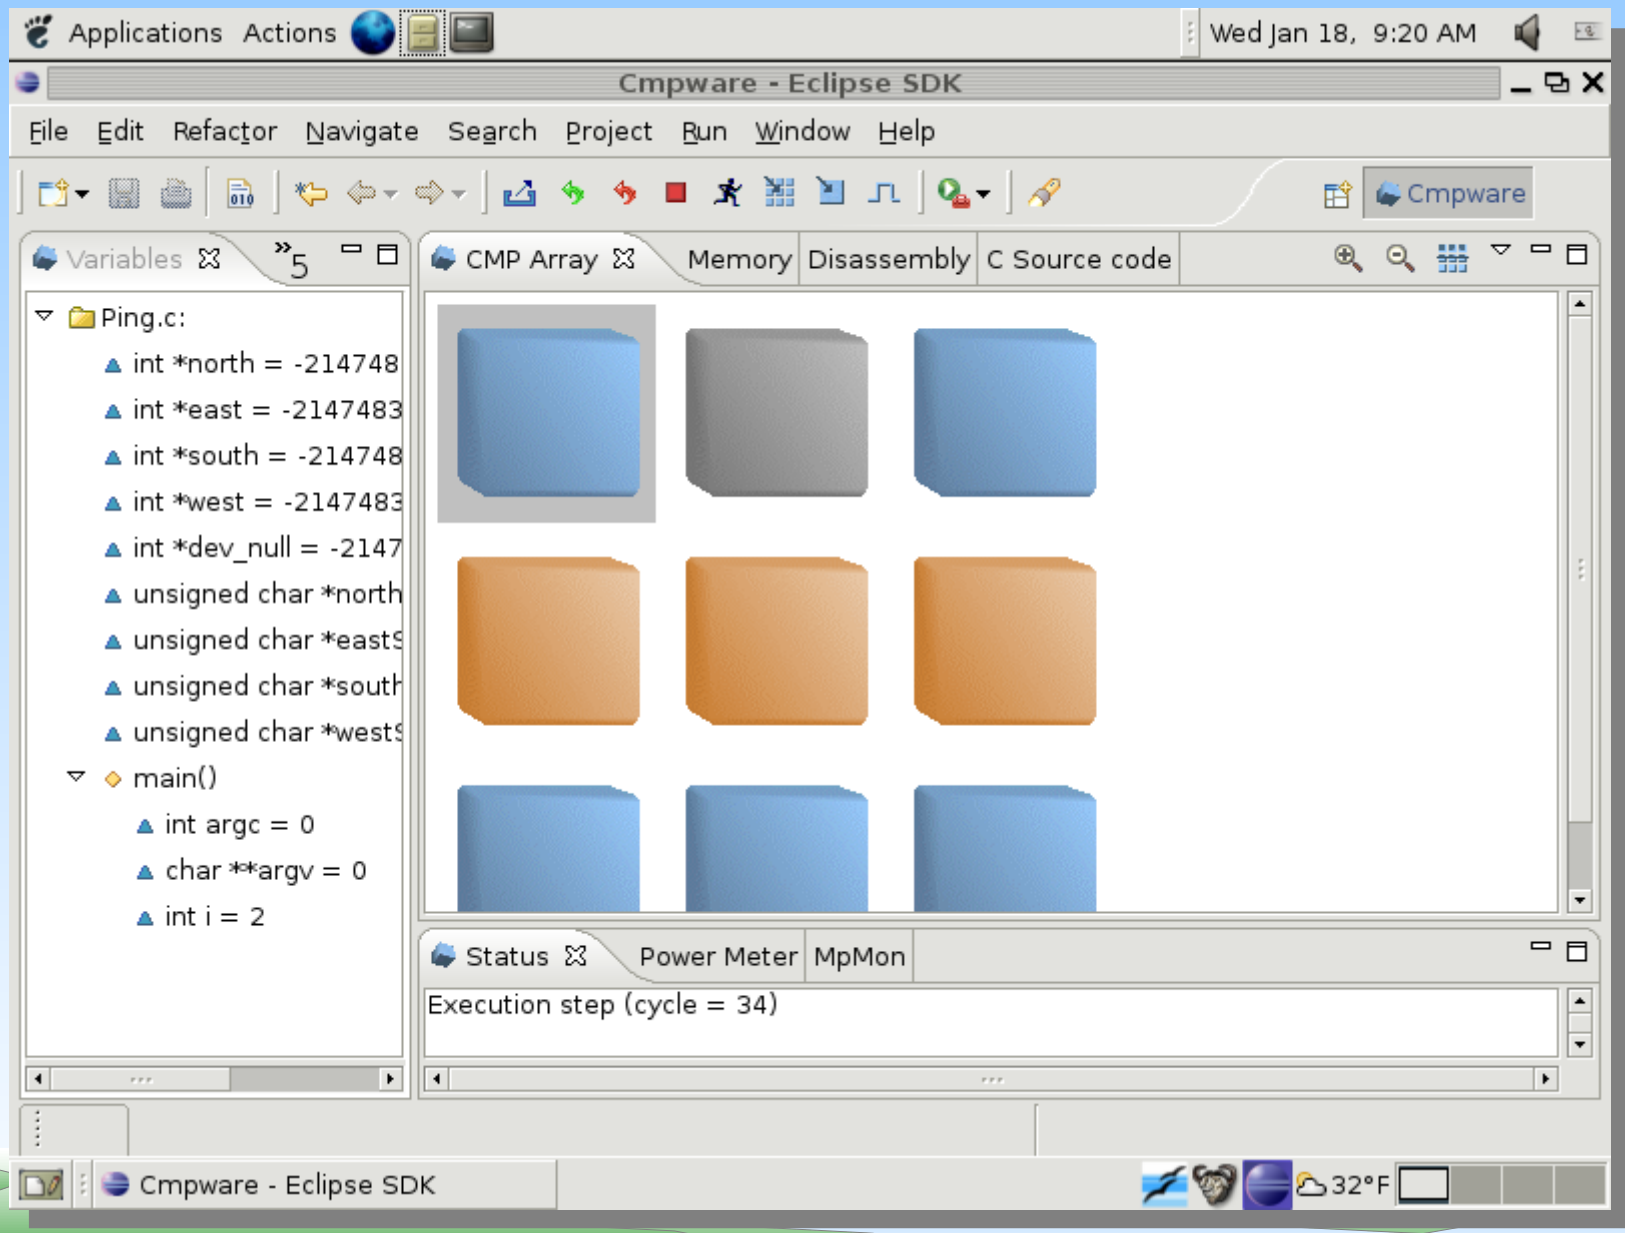

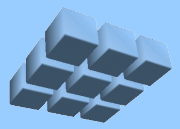

Conclusions

- Multicore Multiprocessor (MCMP):
  - Low power
  - High Performance
  - Inexpensive
- A new approach to large-scale computing
- Open questions on systems and software
- Very attractive potential Price / Power / Performance / Programmability (PPPP)