ECE 574 – Cluster Computing
Lecture 19

Vince Weaver

http://www.eece.maine.edu/~vweaver

vincent.weaver@maine.edu

10 November 2015

Announcements

- Projects
- HW extended

MPI Review

- MPI is *not* shared memory
- Picture having 4 nodes, each running a copy of your program *without* MPI.
  Things initialized the same in all will have same values, no need to initialize.
  Things initialized in only one node will need to be somehow broadcast for the values to be the same in all.
- Problems debugging memory issues.

Valgrind should work, but Debian compiles MPI with checkpoint support which breaks Valgrind :(
Mpirun supposed to have -gdb option, doesn't seem to work.

- `MPI_Gather( sendarray, 100, MPI_INT, rbuf, 100, MPI_INT, root, comm);`
  rbuf ignored on all but root

- All collective ops are blocking by default, so you don't need an implicit barrier

- MPI_Gather(), same as if each process did an

MPI_Send() and the root note did in a loop MPI_Receive() incrementing the offset.

- MPI_Gather() aliasing

  cannot gather into same pointer, will get an aliasing error

  Can use `MPI_IN_PLACE` instead of the send buffer.

  Why is this an error? Partly because you cannot alias in Fortran. Just avoids potential memory copying errors.

SLURM update

Be careful with letting jobs run infinitely. I put a 5 minute timelimit on because some jobs were taking forever and locking people out of the queue. Will need to modify your slurm files.

Hadoop Update

Can set up Hadoop on single machine, even the name and data servers. Just download big chunk of Java, have Java and ssh installed. Didn't get a chance to try to run it yet.

Reliability in HPC

Good reference is a class I took a long time ago, CS717 at Cornell:

`http://greg.bronevetsky.com/CS717FA2004/Lectures.html`

Sources of Failure

- Software Failure
  Buggy Code
  System misconfiguration

- Hardware Failure
  Loose wires
  Tin whiskers
  Lightning strike
  Radiation
  Moving parts wear out

- Malicious Failure
  Hacker attack

Types of fault

- Permanent Faults – same input will always result in same failure
- Transient Faults – go away, temporary, harder to figure out

What do we do on faults?

- Detect and recover?
- Just fail?
- Can we still get correct results?

Metrics

- MTBF – mean time before failure
- FIT (failure in Time)
  One failure in billion hours. 1000 years MTBF is 114FIT.
  Zero error rate is 0FIT but infinite MTBF Designers just
  FIT because additive.
- Nines. Five nines 99.999% uptime (5.25 minutes of downtime a year)
  Four nines, 52 minutes. Six nines 31 seconds.

- Bathtub curve

Things you can do Hardware

Hardware Replication

- Lock step – Have multiple machines / threads running same code in lock-step Check to see if results match. If not match, problem. If replicated a lot, vote, and say most correct is right result.
- RAID
- Memory checksums
- Power conditioning, surge protection, backup generators, UPS

- Hot-swappable redundant hardware

Lower Level

- Replicate units (ALU, etc)
- Replicate threads or important data wires
- CRCs and parity checks on all busses, caches, and memories

Lower-Level Problems

Soft errors/Radiation

- Chips so small, that radiation can flip bits. Thermal and Power supply noise too.
- Soft errors – excess charge from radiation. Usually not permanent.
- Sometime called SEU (single event upset)

Radiation

- Neutrons: from cosmic rays, can cause "silicon recoil" Can cause Boron (doped silicon) to fission into Li and alpha.

- Alpha particles: from radioactive decay

- Cosmic rays – higher up you are, more faults Denver 3-5x neutron flux than sea level. Denver more than here. Airplanes. Satellites and space probes are radiation-hardened due to this.

- Smaller devices, more likely can flip bit.

Architectural Vulnerability factor

- Some bit flips matter less
- (branch predictor) others more (caches) some even more (PC)
- Parts of memory that have dead code, unused values

Shielding

- Neutrons: 3 feet concrete reduce flux by 50%
- alpha: sheet of paper can block, but problem comes from radioactivity in chips themselves

Case Studies

- "May and Woods Incident" first widely reported problem. Intel 2107 16k DRAM chips, problem traced to ceramics packaging downstream of Uranium mine.

- "Hera Problem" IBM having problem. ^{210}Po contamination from bottle cleaning equipment.

- "Sun e-cache" Ultra-SPARC-II did not have ECC on cache for performance reasons. High failure rate.

Hardware Fixes

- Using doping less susceptible to Boron fission
- Use low-radiation solder
- Silicon-on-Insulator
- Double-gate devices (two gates per transistor)
- Larger transistor sizes
- Circuits that handle glitches better.
- Memory fixes
  - ECC code
  - spread bits out. Right now can flip adjacent bits, flip

too many can't correct.

- Memory scrubbing: going through and periodically reading all mem to find bit flips.

Testing

- Single event upset characterization of the Pentium MMX and Pentium II microprocessors using proton irradiation, IEEE Transactions on Nuclear Science, 1999.

- Pentium II, took off-shelf chip and irradiated it with proton. Only CPU, rest shielded with lead. Irradiate from bottom to avoid heatsink

- Various errors, freeze to blue screen. no power glitches or "latchup 85% hangs, 14% cache errors no ALU or

FPU errors detected.

Things you can do Software

Algorithm Based

- Parity checks, CRC
- Spread out work so that if one gives wrong result it can be checked. Overlap work.
- Add some extra values to calculation that can be checked, can tell if something went wrong

Control Flow Checking

- Knows where code should be allowed to jump to
- If you jump somewhere impossible, checker stops things

Checking Data Structures

Extra state in data structure or checksum so can tell if it gets corrupted.