SC2000 Review

Stephen Wolbers, Jim Simone, Ruth Pordes

January 23, 2000

SC2000

- Held November 4-10, 2000
- Attendees number in the thousands, from industry, academia and national labs/research centers.
- Located in 2000 at the Dallas Convention Center, Dallas, Texas.
- Conference :
  - Tutorials
  - Exhibits
  - Technical Program
  - Vendor parties
- I'll talk about the technical program
- Web pages: www.sc2000.org, sc2000.fnal.gov


GUN SHOW &
2000. 11. 4

SC2000 Plenary and Parallel Talks

- Plenary talks
  - These were given by experts who have worked in the field for many years.
  - Talks were given at the beginning of each day.
  - I'll review the ones I was able to attend.
- Parallel talks
  - Primarily given by those who "did the work".
  - I can only discuss those that I attended and can only hit the highlights anyway.
  - A lot of interesting work was described in the parallel talks.

First Day

- Guest appearance, Jack Kilby, nobel prize winner in physics in 2000.
- Keynote address, Steven J. Wallach
  - "Petaflops in the Year 2009"
  - http://www.sc2000.org/webcasts/keynote.ram
  - No surprises, expect Moore's law through 2009.
  - However, this talk dealt with Supercomputing ($150 Million).
  - All UNIX (LINUX), no NT, still FORTRAN.
  - Will need parallel programming of some sort.
  - A large computer is composed of lots of big computers tied together with network/switches.
  - Emphasis on optical, switches, and networking (Terabit-scale).

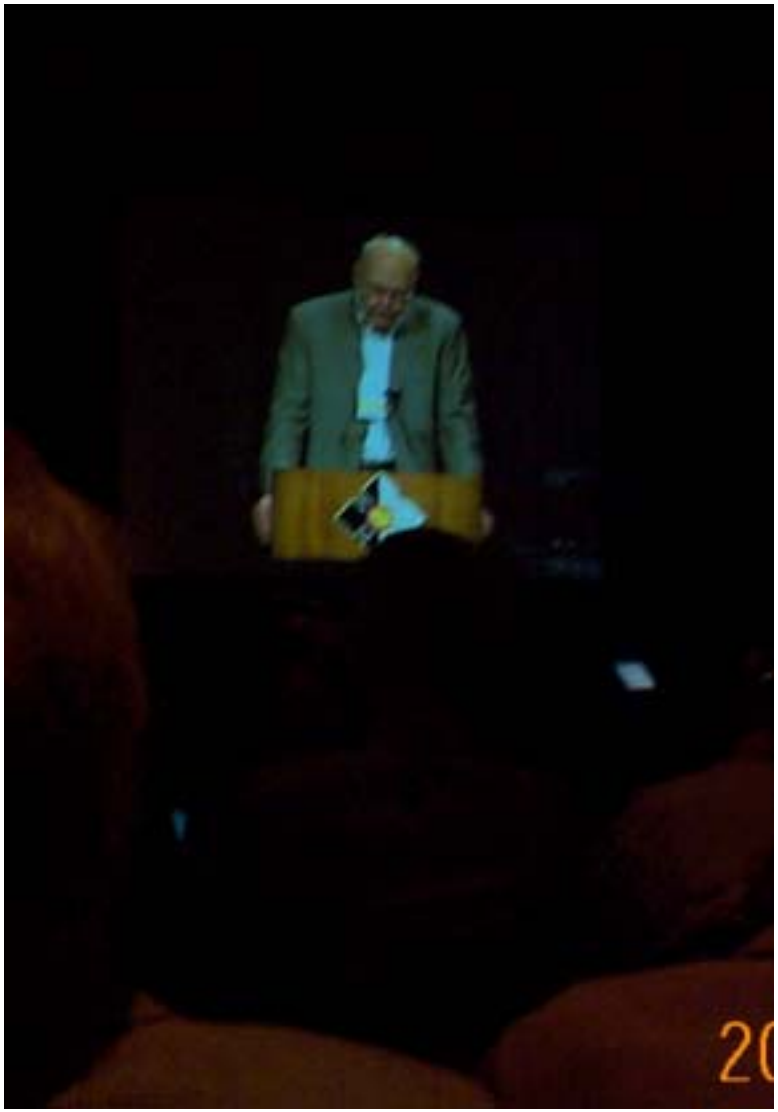
20


2000

Thomas Sterling – COTS Clusters

- http://www.sc2000.org/webcasts/sterling.ram
- General overview of “PC Clusters” (Beowulf)
- Emphasis on the impact on “Supercomputing”.
  - Tightly coupled applications were the focus.
  - Myrinet, SCI, mesh, hypercube, tree topologies.
  - Some discussion of PVM and CONDOR.
- Clusters now appearing on top 500 Supercomputer list.
  - 28 on the current list.
  - Claim is that this is a small but growing contribution.
- Price/performance is one of the main drivers for this trend.
  - Chip speed continues to rise.
- Open Source was emphasized.
  - A common theme?

Eugene Spafford – Computer Security

- http://www.sc2000.org/webcasts/sterling.ram
- Main Message – Security is an after-thought, not an integral part of system and software design.
- There are many many vulnerabilities and the number increases all the time.
- The standard security approaches are not very good.
  - Firewalls
  - Virus checkers
  - Intrusion detection
- The problem will get worse as machines, disks, networks, etc. grow in capacity and speed. (Wireless doesn't help.)
- It will also get worse because high quality software is not a high priority.
- The solution, according to Eugene, is to improve software engineering.
  - Better quality and security should be built-in.

Margaret Wright – Insight and Numbers

- http://www.sc2000.org/webcasts/thursday.ram
- Interesting talk – fairly mathematical/theoretical.
- Most interesting part was the example of modeling antenna placement and orientation for cellular systems.
  - Margaret Wright works for Lucent so this is a real world problem.
  - Sophistication required to do this modeling properly is quite high.
    - Need building, topography, etc. in great detail
    - Lots of data is required and lots of computing to analyze the data.
  - Goal is to replace the “receiver in a truck” technique currently used to optimize antenna placement.

J.C. Browne – Parallel/Distributed Programming: Research Success – Application Failure?

- http://www.sc2000.org/webcasts/thursday.ram
- The name sort of tells the story.
- Many many parallel programming languages and strategies over the past years : MPI, OpenMP, HPF,…
- However, the successes have not been overwhelming.
- Proposes some new approaches.
  - These are meant, I think, to produce a system which will work for real-life problems and people.
  - The details are a bit fuzzy (to me), this case has been made by many people in the past but I suspect it is a hard problem.

Interesting parallel talks

- Keiji Tani, “Status of the Earth Simulator Project in Japan”
  - Interesting talk about a “custom” very high performance non-defense computer in Japan
  - >5 Tflops, 640 processors.
  - Targeted for earthquake, weather, climate, etc. research.
  - Custom machine built by NEC (definitely not commodity!)
  - Total cost approx. $600 Million, I believe.
    - (includes new building, power, etc.)
  - No paper in proceedings (this was part of the Masterworks series).

W. Feng – The Failure of TCP/IP in High-Performance Computational Grids

- TCP gives poor performance because of:
  - Flow Control and congestion control.
- 3 kinds of TCP/IP described:
  - RENO
  - Vegas
  - Tahoe
- Different behavior of the three because of the strategy employed.
- This needs to be understood and studied to understand behavior of grid computation.
  - GRID computing relies on the network and network performance, QOS, etc. are probably crucial to successful GRID computation.
- Some plots were shown of behavior given current and near-future networks.

David Anderson – SETI@home: Internet Distributed Computing for SETI

- Some numbers:
  - 50 GB/day on DLT from Aracebo (raw data)
  - Data split into 10 kHz x 107 s chunks
  - Average computing used is 20 Tflops (constantly increasing)
  - Total cost about $500K., 5 people.
- Not open source, for various reasons.
  - Has caused much discussion.
- Each data set is processed >1 time as a check.
  - People hack the code, do other funny things.

Seti
http://setiathome.ssl.berkeley.edu/

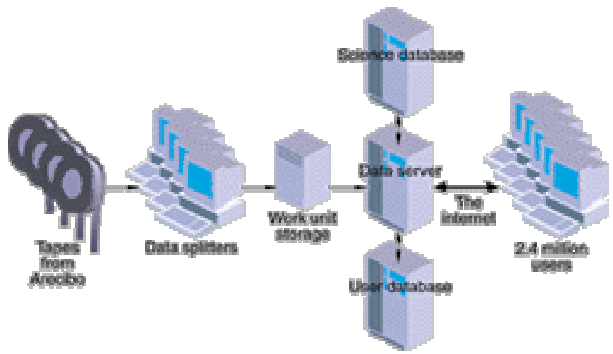

Conclusions

- SC2000 consists of both the display floor and technical talks, though not all of the technical talks are excellent.
- There were many topics covered, and it was not possible for me to attend all of them, especially the parallel talks.
- It would be good if Fermilab technical work was represented at future SC conference technical sessions and I encourage people to contribute.
- SC2001 is in Denver, Colorado.