How ASCI changed the face of modern computing

Pat Miller – pmiller@quantlab.com

SUBHEAD

Abstract

The purpose of ASCI/ASC was to demonstrate that modeling incredibly complex systems was not just feasible, but practical and effective. The program provided a willing "first-mover" that let vendors build systems at (then) unprecedented scales for computation, storage, and networking. Those components were key to the program's success, but they had interesting side-effects.

The ground work done made it tractable for non-government users to pick up the tools and solve their own problems. Clusters became commonplace rather than something used only by fringe groups. More importantly, high performance computing and the associated techniques became accepted and expected practice for all forms of mathematical thinking. HPC ideas, techniques, software, and hardware started showing up everywhere! Resource management and storage formats/models/systems/hardware flowed from vendors to high end users outside the government labs. GPGPU became so commonplace that it dropped the GP. People and ideas streamed out from the Labs to ignite innovation up and down the tech stack.

I'll talk about how I've seen industry transition from a 1980's style throughput computing model to a real HPC model and discuss some of the consequences for the Labs as these new "commodity" HPC actors come of age.

SUBHEAD

So why did we launch ASCI (now ASC) in the first place?

- It's all about the stockpile…
- But we couldn't dig holes in the desert anymore
- So we had to prove our methodology worked

SUBHEAD

Scientific simulation

- The scientific method invites us to make and test hypotheses
- Many things are much easier in models than in “real” life
- But we want models with sufficient fidelity to answer real world questions
  - Very true in my industry!

SUBHEAD

First movers

- There was a real chicken and egg problem that ASCI solved
- There was an untapped market for high-end computing once the kinks were worked out
- But who pays the R&D?
  - For graphic cards in the 90's, it was the gaming market
  - For really large scale clusters/storage/interconnect, it was DOE
  - ASCI-White was really a video server as was the early GPFS!

SUBHEAD

So we built some big, 3-D codes… so what

- We had the lab codes
- We had the ASCI center codes (real rocket science!)

- It led to a change in mindset about how to approach big problems.

SUBHEAD

So why did it catch on?

- Well, the simultaneous rise of open-source was part of that
  - Richard Stallman of all people came as part of a LLNL Lab lecture series ☺
- The push for “dual use” and industry-lab partnerships was part of it
- The ASCI Centers created new pools of like-minded people
- Flow of people from the Labs to industry didn’t hurt

SUBHEAD

Everybody is modeling

- Big data is all about modeling
- A lot of “knowledge discovery” is really projecting models on to existing data

- Biology is about modeling
  - Less and less wet lab and more simulation
  - DE Shaw Research’s molecular modeling supercomputer

- Social science is about modeling

- Computer skills are required in grad school now

SUBHEAD

Computation and Wall St

- Money and trading go way back!
  - But, often slow to move forward
  - The last “pits” in Chicago Mercantile closed less than 3 years ago!
- Excel probably did more to change Wall St than any other tool
- Well, there is Black-Scholes ☺
- Computational work, though, was very dreary for a long time
  - Dominated by main frame thinking
  - Centralized throughput based tasks

SUBHEAD

Simple beginnings

- DE Shaw was founded on the idea that you could build computational models of the interactions in the market
- Bridgewater purportedly wanted to build models of the whole economy
- Early, simplistic models were fine because they were either run outside of trading or worked with small data
  - Mark-to-market
  - Risk computations
  - Suitable for batch-style offline throughput style

SUBHEAD

Then the problems started

- There's was an arms race in models
- Data sizes started growing (markets and consolidation)
- Trades became electronic
  - Which led to an arms race in models which led to…

- The problems were solved by aggressively growing computational, storage, and interconnect resources
  - Using open source tools
  - Using high commodity components developed for our kind and deep pocketed first movers
  - Staffed and developed by former National Lab scientists

SUBHEAD

So what did it get for us?

- We got a lot of bang for the ASCI bucks
  - Everyone uses the technology that ASCI accelerated to simulate everything
  - People think big rather than thinking around the edges
- The hardware and software to do this are pervasive and (relatively) cheap
- Linux commodity clusters are commonplace
  - You can rent them in the cloud!
  - Hey, linux isn't just a bunch of weird OS kernel hackers (well, maybe it still is)