UNIVERSITY OF SOUTH DAKOTA

Big Computing in High Energy Physics

Symposium on Data Harnessing

David Toback

Department of Physics and Astronomy
Mitchell Institute for Fundamental Physics and Astronomy

November 2019

Outline

- **Overview: High Energy Physics and Big Computing**
- **Big Collaborations and their needs:**
  - **What makes it hard? What makes it easy?**
  - **Individual Experiments**
    - **Collider Physics/LHC/CMS**
    - **Dark Matter/CDMS Experiment**
    - **Phenomenology**
- **How we get our computing needs met**
- **Lessons learned and Looking Forward**
- **Conclusions**

High Energy Physics and Big Computing
High Energy Physics = Particle Physics

- ***Big picture of our Science: Making discoveries at the interface of Astronomy, Cosmology and Particle Physics***
- **Some of the questions we work on:**
  - **What are the fundamental particles in our universe (building blocks of nature)? We've discovered 6 types of quarks and the Higgs Boson, are there more?**
  - **What is the Dark Matter that fills the Universe?**
  - **What is Dark Energy?**
  - **Why is there more matter in the universe than anti-matter?**
  - **What is the origin and evolution of our universe?**

Why bother with all this stuff?

- **Mostly because these are the most interesting questions that be (ok… I'm biased), and we ought to understand the nature of the universe in which we live (and how we came to be)**
  - **Nature is VERY different than we'd typically guess**
- **Rich history of fundamental understanding of nature impacting our everyday lives many years later**
  - **"One day, you will tax it" (Faraday)**

Some History of Particle Physics and Computing

Particle physics often leads the way

- **Punch cards and the like at Los Alamos**
  - **Nuclear physics and early particle physics**
- **Central computer systems for the big players, even through the 90's**
  - **Discovery of quarks and other particles**
- **Move to the Web (CERN is a particle physics place), and competitive desktop computing**
  - **Discovery of the 6th quark (top quark)**
- **Moving to clusters of computers, and Grid Computing**
  - **Discovery of the Higgs boson**
- **Next?**

The Large Hadron Collider (LHC)
(Discovery of the Higgs)

- **Huge particle accelerator colliding high energy particles trying to study interactions, and hopefully creating never before seen particles**
- **Collider Physics at CERN/Fermilab have often been the big computing drivers in the world (brought us the WWW and still drive Grid computing worldwide)**
- **5,000 scientists around the world per experiment, two experiments**
- **Experiments have a 3-tiered, distributed computing model on the Open Science Grid to handle the 10's of petabytes and hundreds of millions of CPU hours each month**

When I think of "Big Computing" this is the epitome of what I mean. Similarly, the epitome of "Big Science"

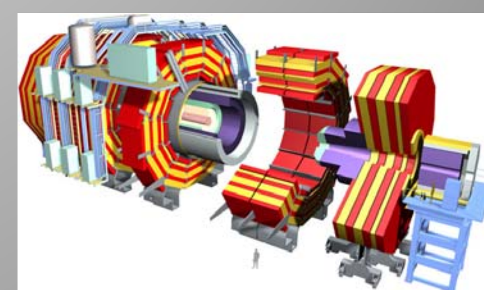

Dark Matter Searches with the CDMS Experiment

- Dark Matter fills our galaxy, so hopefully with a sensitive enough detector we can observe a dark matter particle interacting with a single nucleus
- Much smaller collaboration (~100 scientists), but similar computing issues
  - Only 100's of Tb, and just millions of CPU-hrs /Month
- Less data to analyze, but also fewer people to help do it
- Experiments are much more difficult to understand per unit volume
  - Lots of custom condensed matter physics simulations

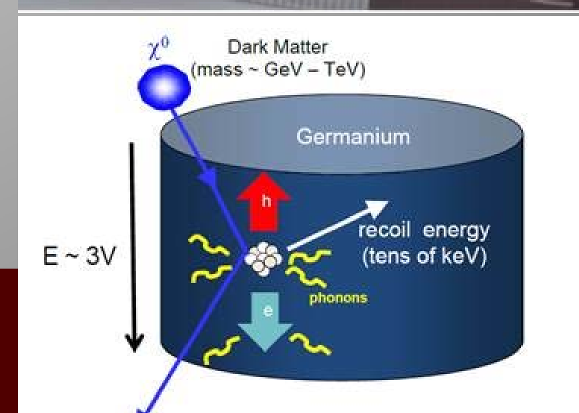

Particle Theory/ Phenomenology

- **In days of old a Theorist could just use a blackboard or pencil and paper. Those problems have mostly been solved**
- **These days they do extensive theoretical calculations and VERY large simulations of experiments:**
  - **Collider data to see what can be discovered at LHC**
  - **Dark Matter to help the best way to "discover" it, and if we did discover it, what could we learn about nature from those measurements**
  - **Interface between Astronomy, Cosmology and Particle Physics**
- **Just need high throughput, a few Gb of memory and a few 10's of Tb**

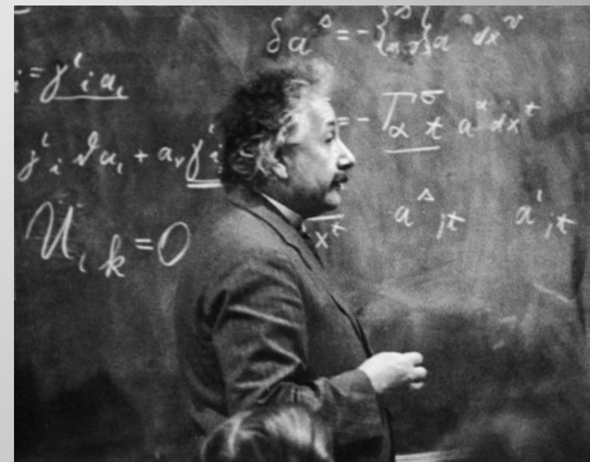

What makes it Easy? What makes it Hard?

Advantages:

- Physics goals well defined
- Algorithms well defined, much of the code is common
  - We already have the benefits of lots of high quality interactions with scientists around the world, so typically we don’t need more statisticians or computer scientists
- Lots of world-wide infrastructure and brain power for data storage, data management, shipping and processing
- Massively parallel data (high throughput, not high performance)
- Economies of scale

Disadvantages

- Need more disk space than we have (VERY big data)
- Need more computing than we have (VERY big computing)
- Need to move the data around the world for analysis, for 1000’s of scientists at hundreds of institutions
  - Political and security issues: Are we going to allow the data to be sent to China? Are we going to allow our Indian collaborators to have accounts on the computers at our National Labs? Are we going to allow Grid jobs to come in from Iran?
- Lots of scientists will have to use lots of software they didn’t write, and SysAdmins will have to suffer with helping jobs run with software that “can’t” be changed
  - “I just installed upgraded the software and now it doesn’t run anymore. Help?”
- Often we don’t know what problem we are solving
  - If we knew what we were doing, it wouldn’t be called Research
- Each computing location has it’s own particular set of quirks (advantages and disadvantages)

Big Computing, Big Data, High Performance and High Throughput

It is worth describing some of the words people throw around in big science

- **Big Computing vs. Big Data (algorithms)**
  - **Big Computing is about the machines and infrastructure to collect huge amounts of data around the world and run millions of CPU hours to analyze it**
  - **Big Data is about writing better algorithms to analyze twice the data in half the time**
- **High Performance Supercomputing vs. High Throughput**
  - **High Performance is often about jobs that need to integrate large amounts of data, like weather simulations. Can do this with memory or threaded processing where the CPU's talk to each other**
  - **High Throughput is when I have lots of pieces of data that can be analyzed individually, and then just added at the end. Especially useful for particle physics where each collision is separate**

Many of the things I'm describing MAY be helpful to how you think, even if they aren't directly useful

- **The problems we are trying to solve may not map well for people who are trying to map weather patterns, or analyze huge MRI images**

Navigating the Waters: Money, Important People and Politics

- **Getting our computing done is VERY expensive, and most "higher ups" want it to be cheap and easy**
- **Countries (and smaller) want to be Leaders-in-the-field, and to be able to demonstrate this in easy-to-measure ways**
- **Often the things that are most helpful to the scientists aren't sexy, and are things no one wants to pay for**
- **We appear to have LOTS of money, and everyone wants a piece**
- **Science often gets done locally**

What's a Scientist to do?

In many ways there are two separate problems that both need to be solved

- **How should the collaboration do its computing?**
- **How should an individual scientist get his computing done in that context?**

If you want to be successful, you need to solve both problems

Structured Software and Servant-Based Leadership/Management

Need to solve a number of problems:

- How can we write software that thousands of scientists can use?
- How do we make it so EVERYONE can get to ALL of the data?
- How do we deal with distributed Computing?
- How do we get people, with huge cultural differences, to agree on a plan (and follow it)?

Lots of Leadership AND Management (those are separate)

- Herding cats? That's too easy… herding birds?
- Lots of planning… LOTS
- Need both vision AND the ability to execute
- Need real professionals (who can make useful code) AND scientists (who make sure the code does something useful)
- The goal is to enable the scientists and then leave them alone, not to tell them what to do

Personal Experience at My Home Institution

- My home institution has GREAT Supercomputing because our local leadership decided that would "put us on the map", but we don't use it
  - Spent lots of money to buy CPU and disk
  - Doesn't want the risk of outside users: Not useful for High Energy Physics
- We actually use a much smaller "cluster" by special agreement
  - Allows jobs and data from around the world (Open Science Grid)
  - Automated data distribution and local jobs accessing remote databases
  - High THROUGHPUT, not high PERFORMANCE
  - Just run LOTS of independent jobs on multi-Tb datasets
    - Priority over idle!!!
  - Luckily, we sit on the high speed link from Dallas to Houston: Easily get ~1Tb/hour from Chicago and 0.75Tb/hr from San Francisco
  - Links to places around the world (CERN-Switzerland, DESY-Germany, CNAF-Spain, UK, FNAL-US, Pisa, CalTech, Korea, France Etc.)

Some Lessons Learned

- Monitoring how quickly data gets transferred can tell you if there are bad spots in the network locally as well as around the world
  - Found multiple bad/flaky boxes in Dallas using PerfSonar
- Monitoring how many jobs each user has running tells you how well the batch system is doing fair-share and load balancing
  - Much harder than it looks, especially since some users are very “bursty”: They don’t know exactly when they need to run, but when they do they have big needs NOW (telling them to plan doesn’t help)
- Things didn’t really start working well until we had a new type of Experts: know both the software *and* the SysAdmins
  - This is the person you turn to when things break. Software problem here? At CERN? Cluster problem here? At CERN?
  - Useful to have users interface with local software experts (my students) as the first line of defense before bugging Admins at A&M and elsewhere around the world
- Hard to compete with national labs as they are better set up for “collaborative work” since they trust collaborations, but we can’t rely on them alone
  - Upside to working at the lab: Much more collective disk and CPU, important data stored locally
  - Downside: No one gets much of the disk or CPU (most of our users could use both, but choose to work locally if they can). Need something now? Usually too bad
  - Different balancing of security with ability to get work done is difficult

Real-Time Monitoring is More Helpful than we'd like to Believe

- **Constantly interrogate the system**
  - **Disks up? Jobs running? Small data transfers working?**
- **Run short dummy-jobs for various test cases**
  - **Both run local jobs as well as accept automated jobs from outside, or run jobs that write data off-site**
- **Automated alarms for "first line of defense" team, but can be sent to the Admin team**
  - **Send email as well as make the monitoring page Red**

More detail about our monitoring at http://hepx.brazos.tamu.edu/all.html

I wish more things ran this way. Fun example for "If I ran the world": Every time I started driving "off road" in a big city I wish GoogleMaps would use that as a trigger to check to see if it's maps are out of date (usually, there is construction or a new road). Or, ten people in an hour all make the same "wrong turn" and GoogleMaps needs to "recalculate"

Looking Forward

There will be continued struggles as computing gets to be a bigger and bigger part of our world (we have "good" problems)

- **Need to make it easier to write software that is useful for longer**
- **I think that Grid Computing (where CPU cycles are distributed, like power is now) could well be the future.**
  - **Many groups are simply learning to wrap up their jobs and ship to Google for processing at cost-per-CPU hour price**
  - **Instead of an Excel file on my machine, going to GoogleSheets (Cloud computing) may be the future as CPU and data transfer rates get better (better sharing of data, like the change the WWW brought)**
- **We will have to come to terms with the fact that the people with the dollars to spend see science differently than the scientists**
  - **Efficiency vs. lavishing thinking time to understand what the heck is going on**
  - **While engineers are comfortable being managed, scientists aren't (both administrators and scientists will have to learn)**
    - **Every scientist needs an MBA? Physics is the liberal arts of the 22nd century?**

Want to get ahead? Some suggestions

- **Remember what engineers know: There are three things you can have: You can have it done quick, you can have it done cheap and you can have it done right. Pick two!**
- **Teach administrators the difference between Big Computing and Big Data**
  - **While medicine can be really enhanced by analyzing twice the data in half the time, discovery doesn't really work that way**
- **Realize that if you have custom software, the "biggest bang for the buck" is spent on the people running making the software work on the computing, and not on buying more machines**
  - **Can't get this money from the funding agencies (they just want to pay for scientists)**
  - **SysAdmins can't really help here, and it's a misuse of their time and expertise**
  - **If you build it, they MAY come. A solution in search of a problem?**
- **Invest in networking, and integrated infrastructure, not just in big machines**
  - **Computing rooms and cooling will be a bigger fraction of the cost, or you need a plan to use computing somewhere else**

Conclusions

- High Energy Particle Physicists are (and have been) Leaders in Big Computing for decades
- There are lots of important lessons for scientists, but stealing knowledge/experience/tools from scientists has almost always paid off
- The future is bright for Data Harnessing, but it's important to realize that while CPU, disk and fast algorithms can speed up the analysis it can't replace the hard part: Figuring out what the heck is going on, and what is worth doing
- With wise leaders, managers and individuals with computing expertise the sky is the limit for discovery…

Abstract

High Energy Particle Physicists are (and have been) leaders in Big Data/Big Computing for decades. In this talk we will focus on the Big Collaborations (including the Large Hadron Collider that recently discovered the Higgs boson) and their needs, as well as how we work with the rest of our collaborators doing dark matter searches, astronomy and large scale theoretical calculations/simulations. We will discuss our use of the *Brazos cluster* for the bulk of our computing needs because it has both allowed us to cope with our High Throughput requirements as well as our issues with working with collaborators, data and software from around the world in a grid computing environment. Finally, we will present some results on how well things have worked, as well as some comments about what has worked and what would be helpful in the future.