Computing Challenges

Ian Fisk

October 12, 2012

The Real Challenge

- Computing at the scale we're currently work is rarely limited by a hard technology issue
  - Limited by what people are willing to spend
- Computing is often the limit of a physics program
  - We make choices and priorities about the events we can collect and the analyses we can do based on how many computing resources we have
- Expectations of the community track with improvements in computing technology
  - Expectations scale with what we see in other parts of life (laptops, commercial services, other communities)

Ian Fisk
FNAL/CD

LHC vs Tevatron

Looking at Tevatron a few years into Run2 and LHC on year 3

	Tevatron	LHC
Trigger	50Hz	ATLAS 500Hz CMS 350Hz LHCb 2kHz
RAW Event Size	150k	ATLAS 1.5MB CMS 0.5MB
RECO Event Size	150k	ATLAS 2MB CMS 1MB
Reco Speed	1-2 seconds on CPU of the time	10s on CPU of the time

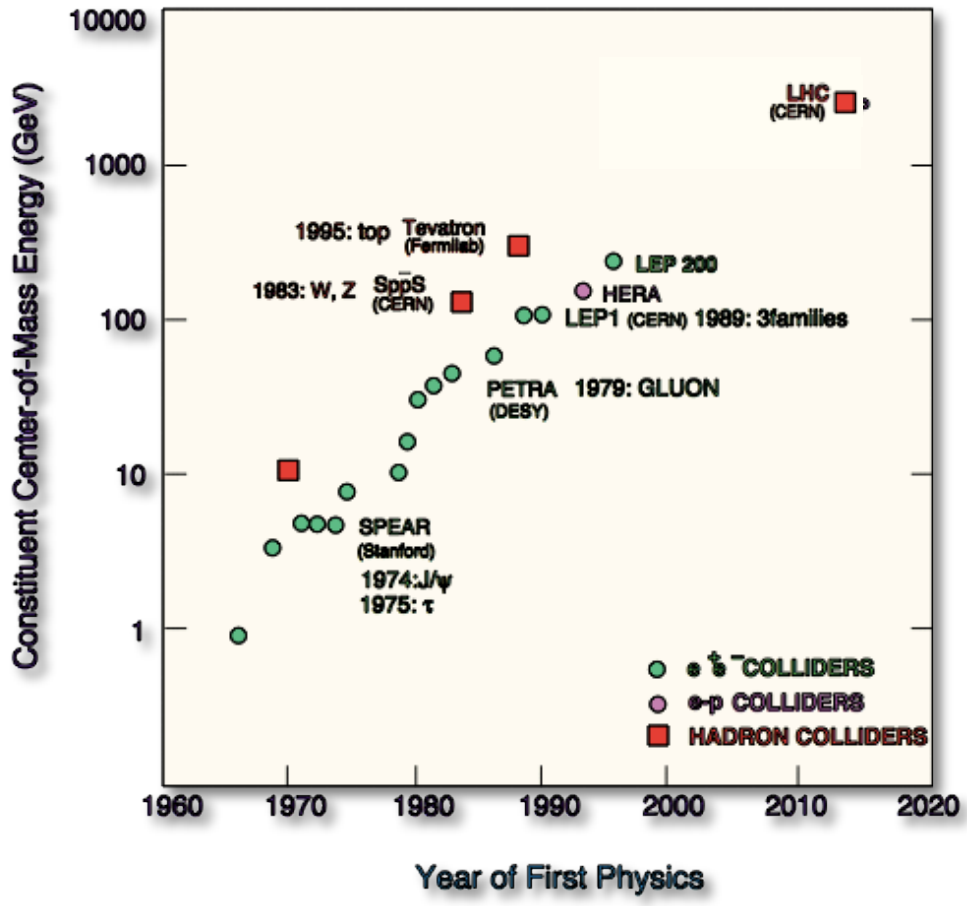

Roughly a factor of 10 in the relevant quantities

Predicting Improvements

- We have not generally been good about predicting what we will need
  - Tend to underestimate computing improvements
    - Also underestimate computing needs
    - Hard to think about factors of 10 or more

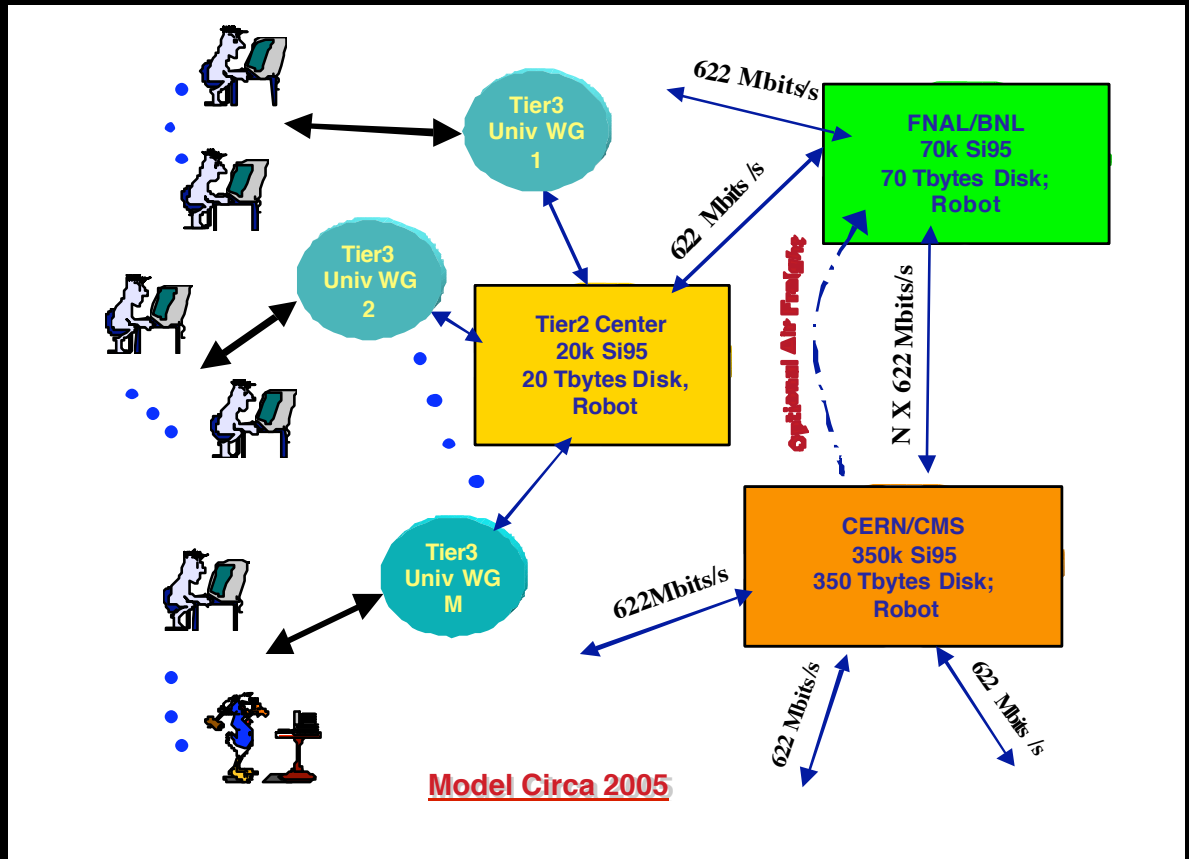

Ian Fisk
FNAL/CD

Increases

- LHC Computing has been increasing at less than Moore's Law for some time
  - Since 2010 we've gone up by 50%
    - Would be ~200% larger by Moores Law
    - Trigger rate also increased from 300 to ~500
- We recently predicted the need of a increase of 100% between 2012 and 2015
  - This is half of Moore's Law, but has already hit resistance from agencies

Ian Fisk
FNAL/CD

Increases Processing

- If we say there are 3-5 doubling cycles between 2013 and 2021
  - Somewhere between factor of 8 and factor of 30 increase in capacity
  - In 2013 we expect to be capable of 1kHZ.
    - We assume improvements in code performance and are assured to have increases in event complexity
    - Let us assume 3 times longer reconstruction time
    - This would correspond to 10kHz sustainable data taking
      - Not clear it's justifiable from a physics perspective
- Currently we can sustain IO of roughly 2GB/s
  - Again assume 8-30 increase
    - 16GB to 60GB
    - Again assume a factor of 3 increase in size
      - 10k to 40k events

Ian Fisk
FNAL/CD

Examples

- Unless there is a significant changes in the ways we work
  - Someplace in the neighborhood of 10-20k Hz sustainable rate
- A significant change would be to have take a higher rate and perform more of the complete chain of processing and analysis and write of synthesized output
  - This is an extension of the data scouting technology
  - Would be a big change of mentality

Ian Fisk
FNAL/CD

Looking forward

- Intrinsically, there is nothing really unscalable about the way we work now
  - We need to solve the multi-core problem, so we submit more requests that use multiple-cores or entire boxes simultaneously
    - Short term problem and this will ensure the workflow management scales as we increase cores.
- Network, CPU, and Storage are improving together
  - As the volumes of data increase, storage becomes an issue
    - More dynamic use of the storage can improve this
- Big question is whether we can afford it moving forward and if there are more economical ways of working

Ian Fisk
FNAL/CD

Looking Forward

- Computing is at something of a cross roads
  - In one direction are clouds
    - Generic computing services that are bought, shared, or contributed
    - Computing as a service
  - In the other direction are very specialized systems
    - High performance, low power
      - Massively multi-core
      - GPUs
- Most likely we will use both depending on the needs

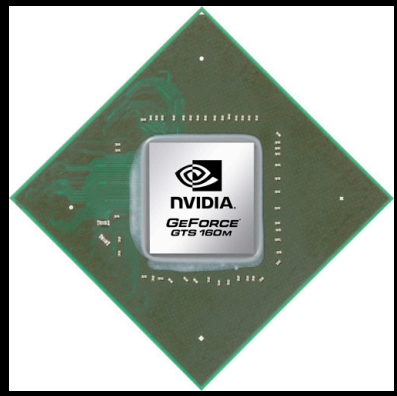

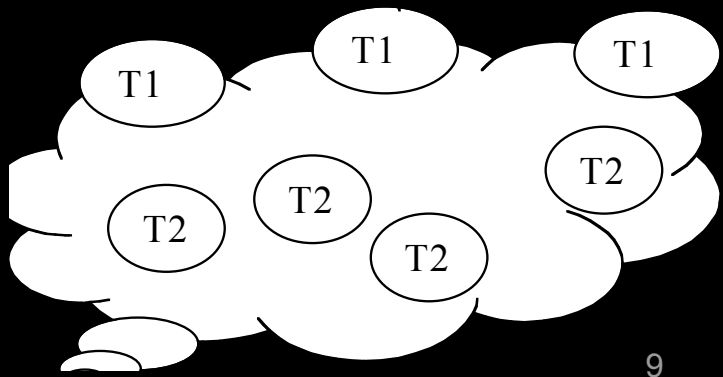

Ian Fisk
FNAL/CD

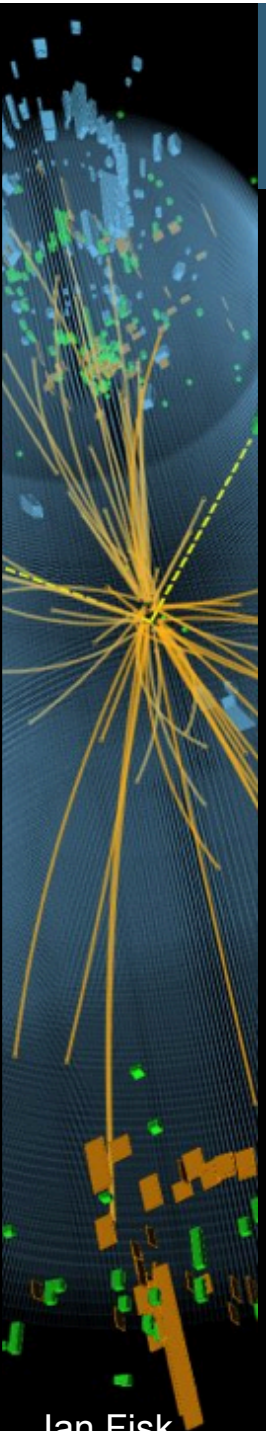

Looking Forward

- Cloud provided computing tends to be factors more expensive than providing the resources in house for resources that are heavily used
  - The company needs to make money and you have to assume they have a huge efficiency gain over you associated with scale to make the service competitive
  - There are a variety of examples of non-commercial clouds which are interesting
  - Costs for commercial facilities is coming down
    - Interesting to cover peak periods
- Need to prepare for a time when this could be the norm
  - Unclear if Clouds follows a utility model or a rental car model

When is the crossover?

- What would convince us we had bought our last farm?
  - Economic models
  - Fall back solutions and risks
  -

Ian Fisk
FNAL/CD

Changes in the model

- Up to now we have typically worked in a model where events are selected, stored, and processed
  - Triggers are as loose and simple as possible up to what we can afford for offline computing
  - Events are treated as precious and we spend a lot of resources protecting them
- HEP is likely always do be driven by atom units which are collision events
  - Drives how we compute

Ian Fisk
FNAL/CD

Changes

- A significant change would be to have take a higher rate and perform more of the complete chain of processing and analysis and write of synthesized output
  - This is an extension of the data scouting technology
  - Would be a big change of mentality
- Something between trigger and offline

Ian Fisk
FNAL/CD

Outlook

- Future is uncertain
  - Scale needed and resources to meet them are a bit in flux
  - In preparing the Computing Frontier we need to understand the needs, so we try to meet them

Ian Fisk
FNAL/CD