Present and Future Computing Requirements for CMS and ATLAS

Liz Sexton-Kennedy
Fermilab

NERSC BER Requirements for 2017
September 11-12, 2012
Rockville, MD

Special Thanks:

Torre Wenaus (ATLAS computing)
Oliver Gutsche (CMS computing)
Peter Elmer (CMS Offline Software)

1. Project Description

There are x institutions from y countries... too many to list

- Summarize your project(s) and its scientific objectives through 2017

- Our present focus is Higgs and SUSY physics

- By 2017 we expect to have enough data to

  ★ measure many properties of the Higgs like resonance

  ★be able to rule out many SUSY theories or discover its existence

  ★ make many precise Standard Model measurements

2. Computational Strategies

- LHC computing is High Throughput Computing
  - ★ goal is to process as many events as possible in the shortest period of time.
  - ★ it's an embarrassingly parallel problem
- There are 4 different problem domains each with their own set of codes.
  - ★ Event generation - from hard scatter particles to hits in the detector
  - ★ Digitization and Reconstruction - model the detector response including extra interactions and reconstruct the simulated data as if it were collected detector data
  - ★ Reconstruction of collected data
  - ★ Group and Individual User Analysis

2. Computational Strategies - 2

- These codes are characterized by a series of combinatoric algorithms
- Our biggest computational challenges are the volume (there is a large needle in the haystack problem) and complexity of the data to be processed
- Our parallel scaling is limited by memory and memory bandwidth
- We expect our computational approach and codes to change to support more fine grained parallelism by 2017. This will require a reengineering of the current code base.

2.Finding the Needle in the Hay Stack

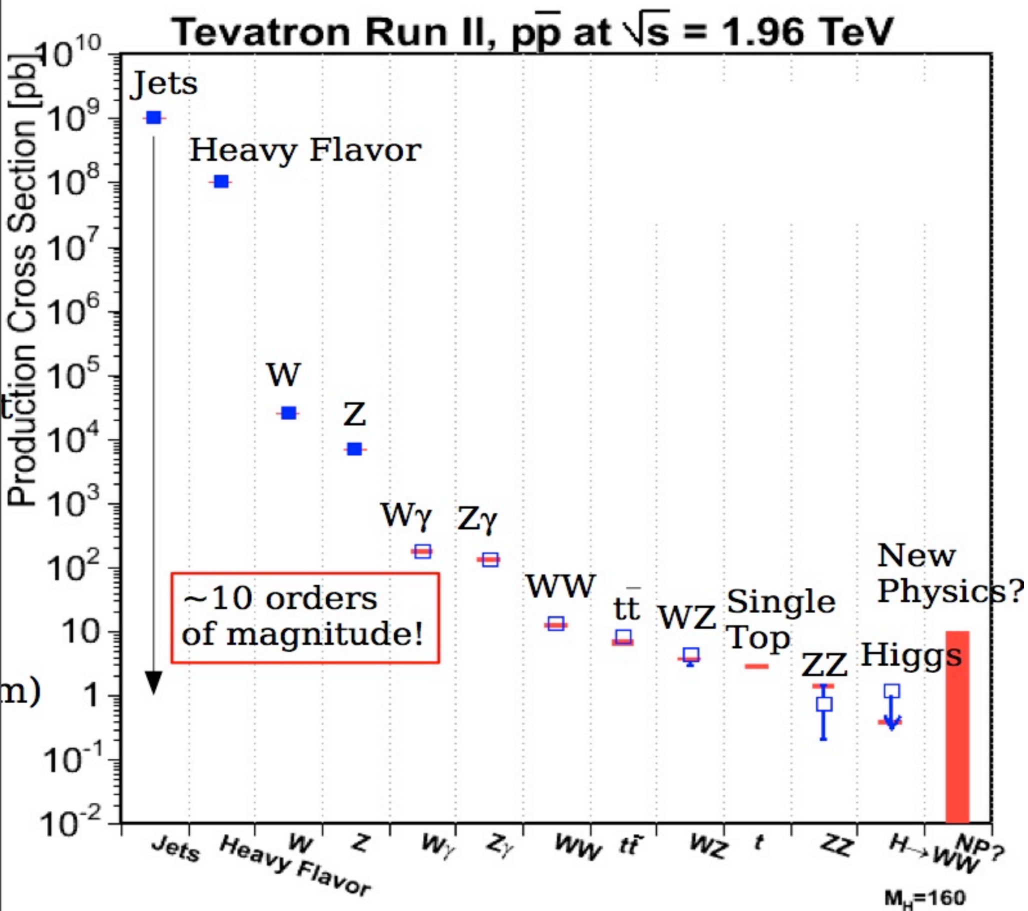

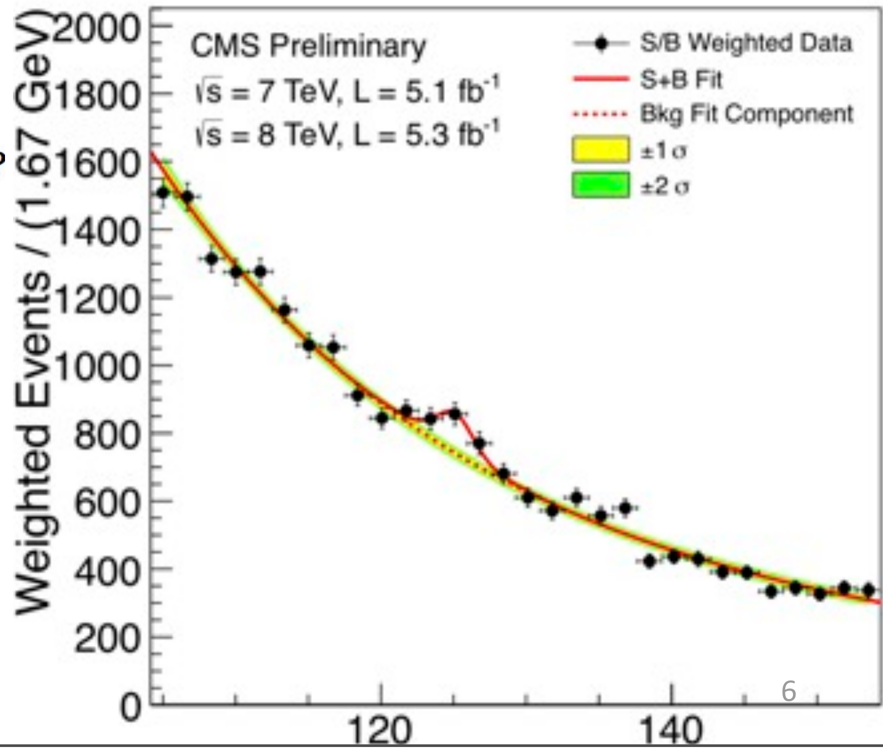

2.Event Complexity and combinatorics

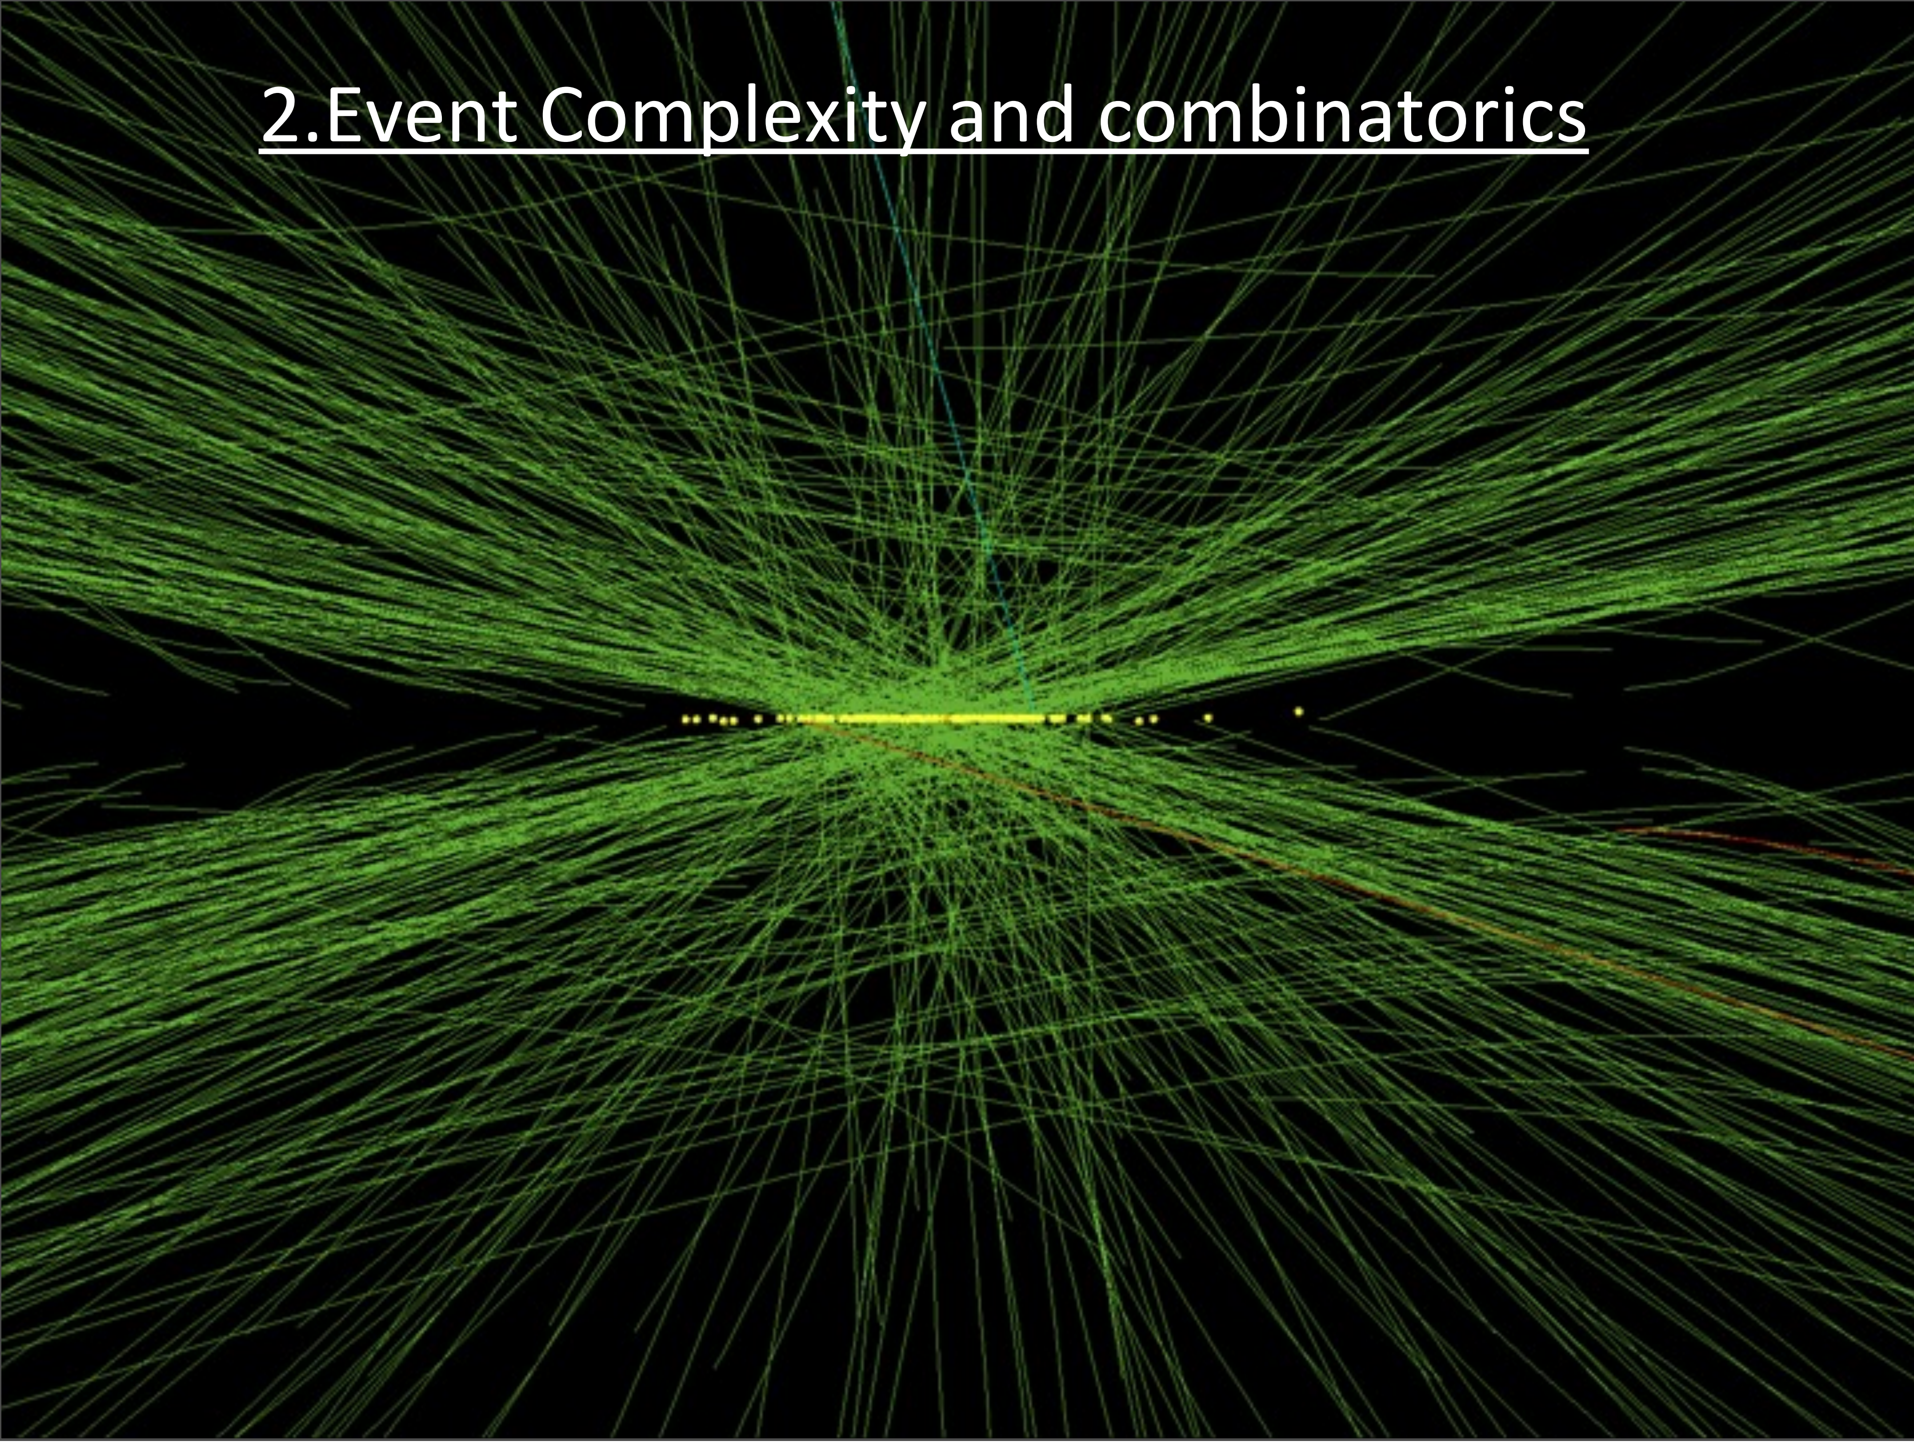

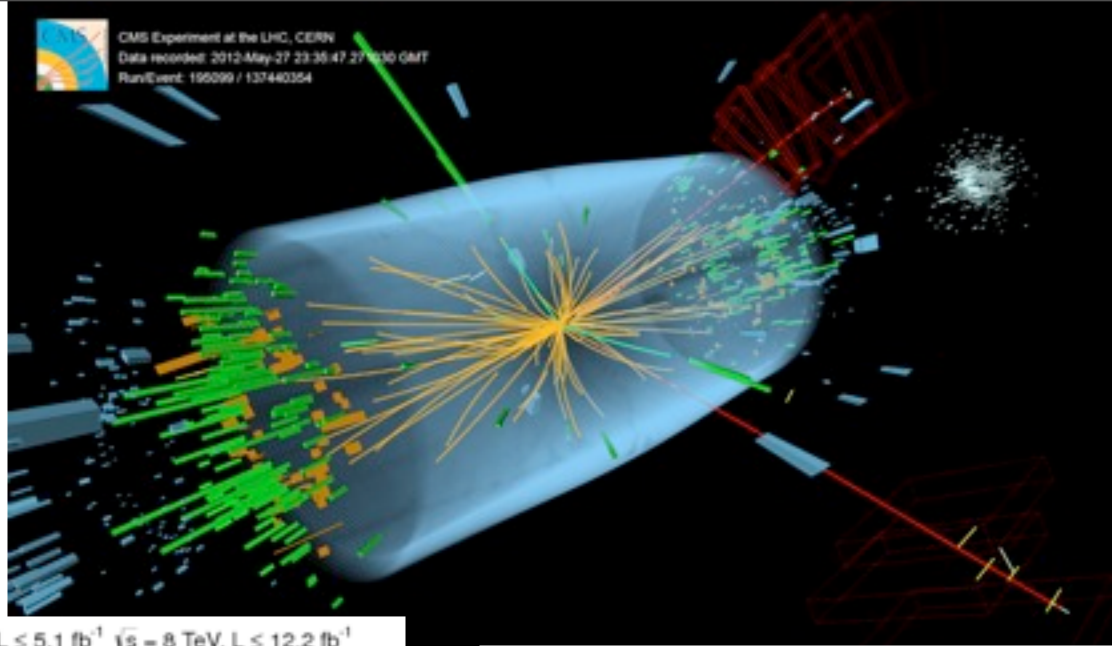
CMS Experiment at the LHC, CERN
Data recorded: 2012-May-27 23:35:47.271030 GMT
Run/Event: 195099 / 137440354

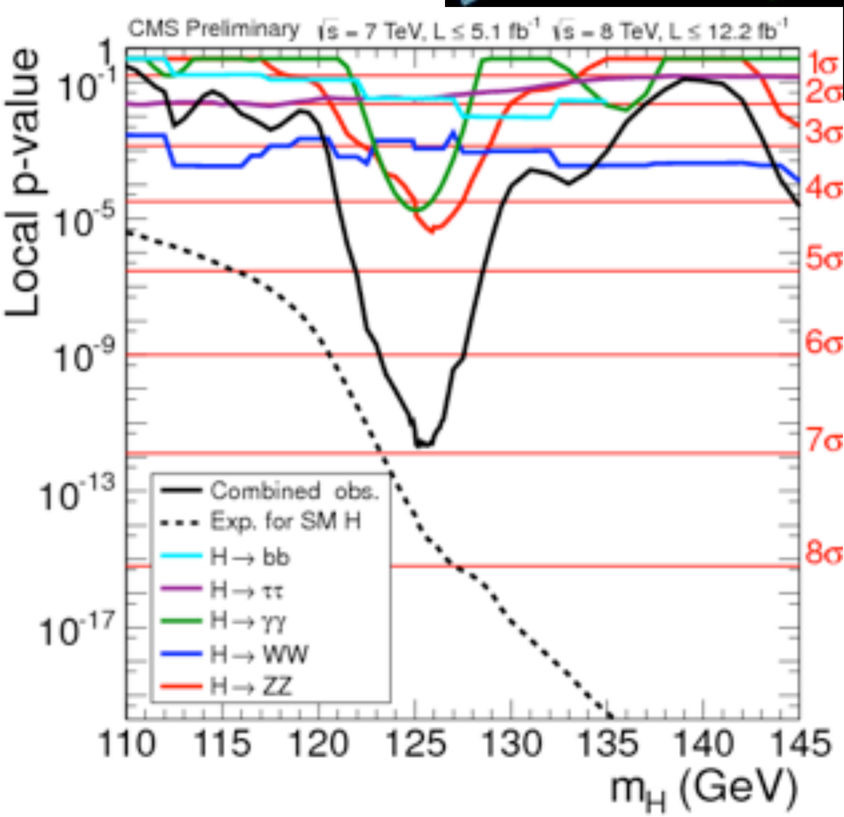
CMS Preliminary √s = 7 TeV, L ≤ 5.1 fb⁻¹ √s = 8 TeV, L ≤ 12.2 fb⁻¹
Local p-value
1
10⁻¹
10⁻⁵
10⁻⁹
10⁻¹³
10⁻¹⁷
1σ
2σ
3σ
4σ
5σ
6σ
7σ
8σ
Combined obs.
Exp. for SM H
H → bb
H → ττ
H → γγ
H → WW
H → ZZ
110
115
120
125
130
135
140
145
m_H (GeV)

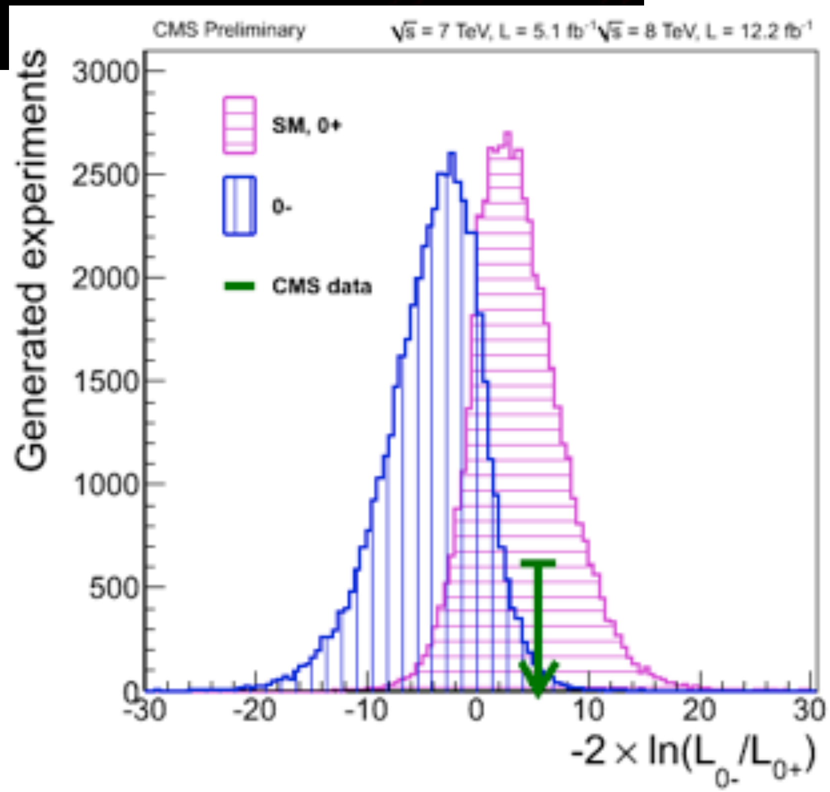
CMS Preliminary
√s = 7 TeV, L = 5.1 fb⁻¹ √s = 8 TeV, L = 12.2 fb⁻¹
Generated experiments
SM, 0+
0-
CMS data
0
500
1000
1500
2000
2500
3000
-30
-20
-10
0
10
20
30
-2 × ln(L_0-/L_0+)

3. Current HPC Usage

- We currently do not use NERSC resources
- In 2012 CMS and ATLAS have used 160K cores continuously on the World Wide LHC Computing Grid
- There is no limit to the number of cores that can be used in parallel as long as each can be allocated 2 Gbytes of memory
- A typical job reads an input file of order 10G and writes out twice that in a reconstruction job and much less in an analysis job
- Our software is developed within the experiment due to the unique detectors, however common tools like root and Geant4 are used. Many computing services and infrastructure components are shared.
- We currently do not use NERSC storage resources, should we? advances in network technologies could enable this.

4. Projecting Requirements

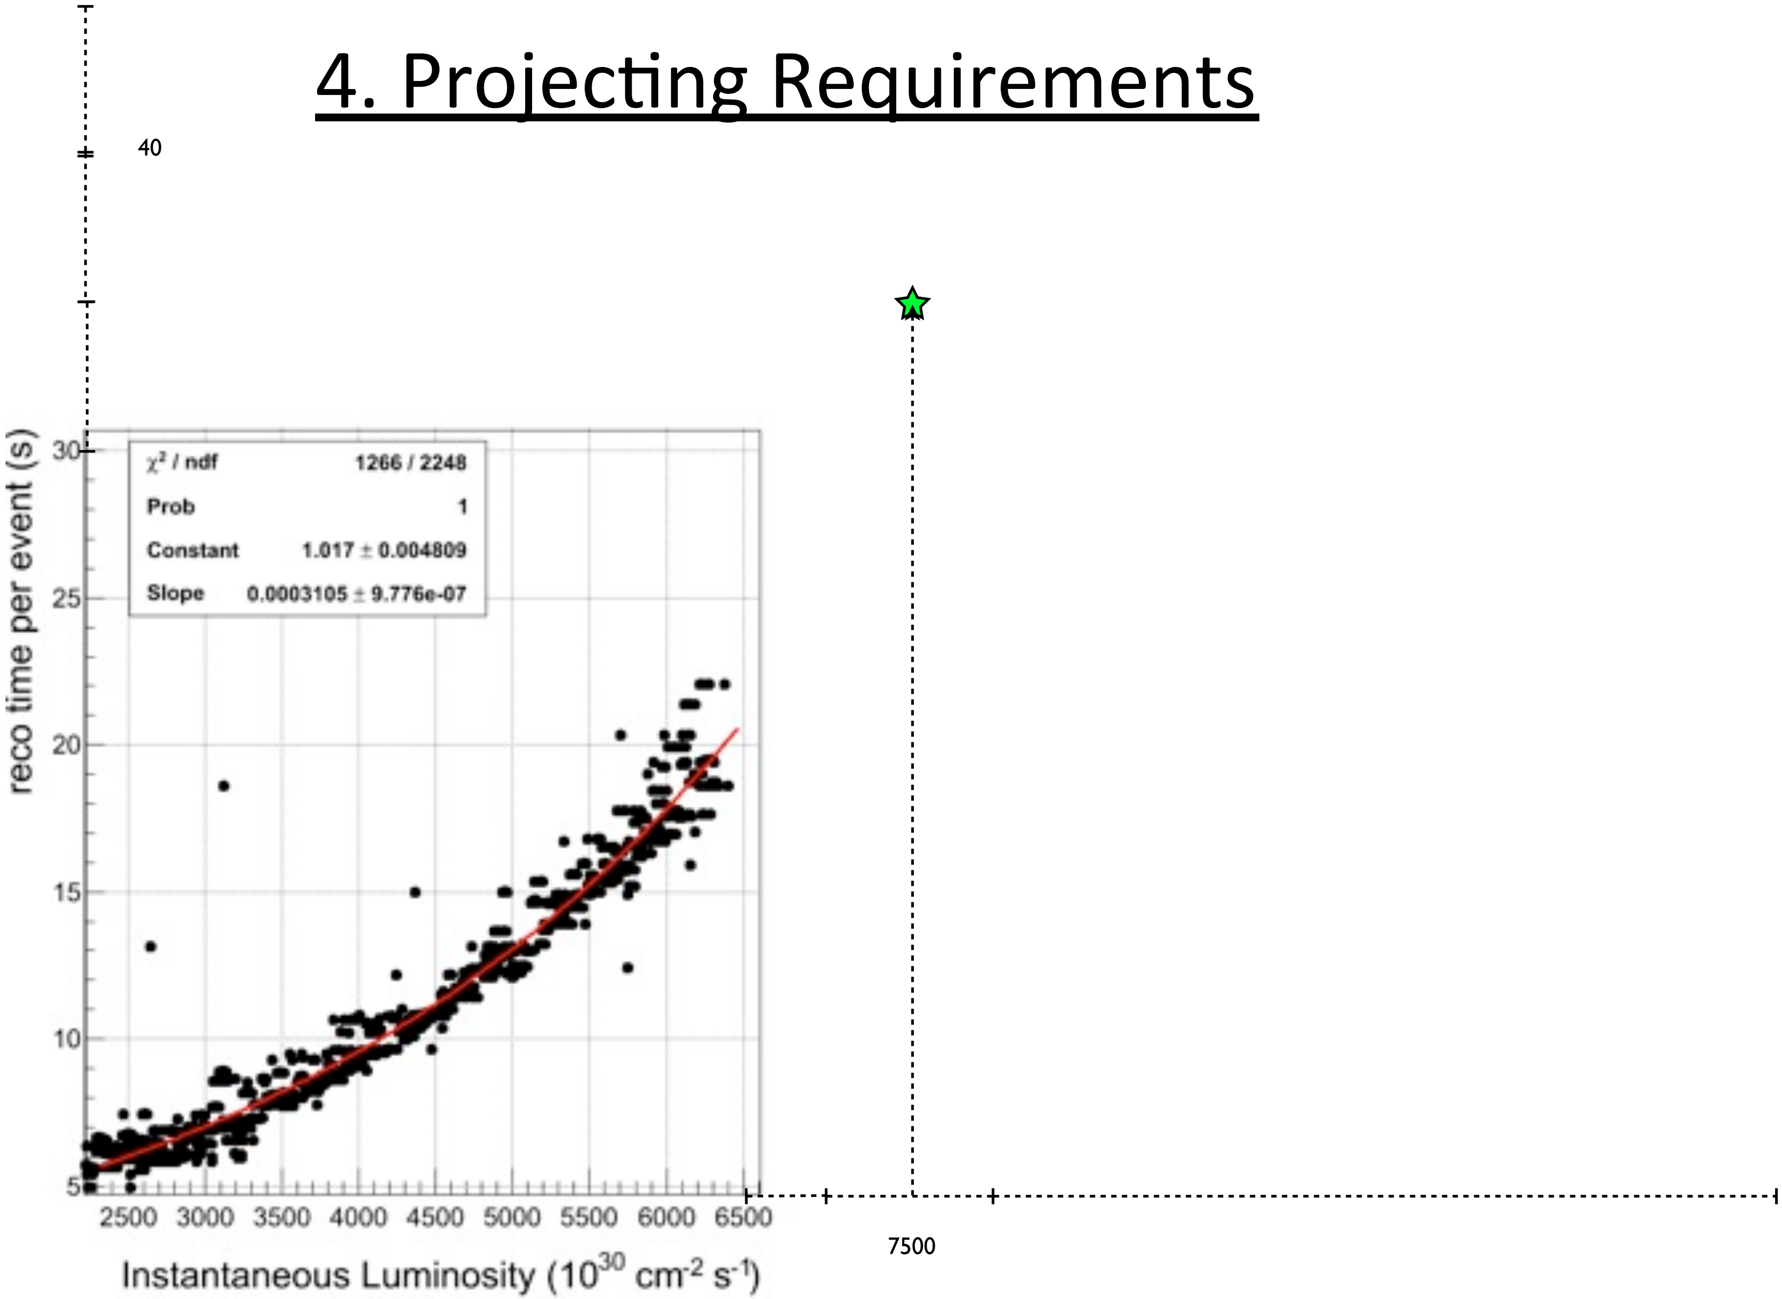

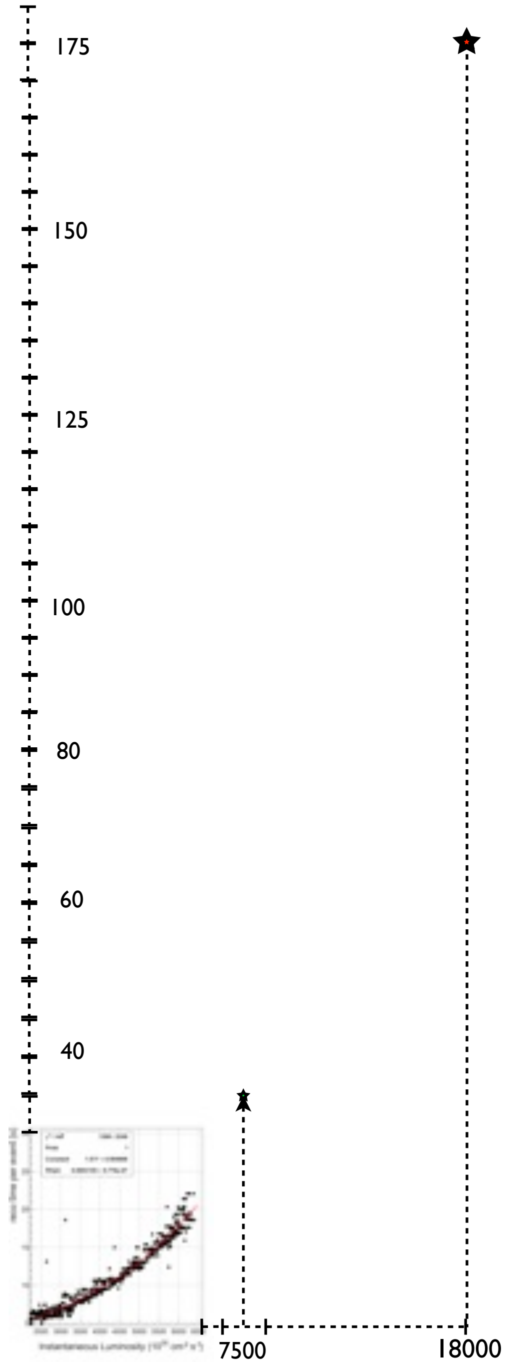

- This point is measured from simulation
- Simulation performance measurements are always an underestimate. What is not clear is
- Our owned resources will only increase as we replace older technology with new in the next couple of years
- Decisions about how to run the LHC collider may be determined by software and computing considerations
- A factor of 2.5 x number of events is assured by increased trigger rates

4. HPC Requirements for 2017

- With the big caveat that we only have to scale by 2.5 times due to the increase in number of collected events, we will need 200K-800K core-hours

- CMS and ATLAS will create 190PB of data by 2017

- Memory per core needs will decrease as we improve our code however it will increase with the complexity of the events. We could require anything from 1-4GB per core in 2017

- In the area of services or infrastructure we plan to move from a dependence on owned resources to a composition of owned resources for the highest throughput needs (prompt reconstruction) and opportunistic resources for our other problems

5. Strategies for New Architectures

- Our strategy for running on new many-core architectures is to develop a fine grained parallelism - task based framework using Thread-Building-Blocks, TBB

- Other than the memory needs we don't see any specific issues using more lightweight cores. Together with the above, we expect to do the optimization work necessary to run them
- There are no clearly identifiable "kernels" to run on GPUs however we do make many calculations that can exploit simpler SIMD-style vectorization

- We could really use help with the transition of the physics codes to concurrent algorithms.

5. Summary

- I have no doubt that computing resources will be much more tightly constrained in the second collider run. We will prioritize our owned resources to deliver the most cutting edge science. What will suffer is the bread a butter standard model physics that we could be doing.

- The LHC could use as many cycles as you can spare. If NERSC would like to help meet the needs for LHC computing, my recommendation is to help us commission the infrastructure needed to treat NERSC facilities as another distributed computing site.

- 32X0 = 0 so this question is not for us, however I could point out that in the period between 2010 and 2012 our resources increased by 50% far from the 200% Moore's law scaling.

- General discussion