Fermilab

CMS: Challenges in Advanced Computing Techniques (Big Data, Data reduction, Data Analytics)

With input from: Daniele Bonacorsi, Ian Fisk, Valentin Kuznetsov, David Lange

Oliver Gutsche

CERN openlab technical workshop

6. November 2015

Introduction

- CMS is handling a tremendous volume of data to extract physics from LHC collisions:
  - Recording data with the detector
  - Transferring files between the world-wide distributed CMS computing sites
  - Processing data and simulations
  - Analyzing data and simulations with thousands of physicists in parallel
- Challenge: data volume and complexity will increase in the next stages of the LHC significantly ➜ Evolution of the computing systems necessary to enable successful physics harvest:
  - Optimize latencies
  - Reduce resource needs
  - Increase selection and reduction efficiencies
- We would like to investigate Advanced Computing Techniques in the areas of Big Data and Data Analytics to achieve this goal

Challenge:
Improve latency of completing high priority production requests

- An automated workflow system is essential for CMS Run 2 operations
  - **1000s of requests managed per week**: Each with many jobs to complete. Example workflows: simulate 100k MC events, re-reconstruct data from October
  - **GRID Computing is a complex ecosystem**: Distributed computing resources with varying capacities and capabilities
  - **Managing input data**: Most requests have sizable input samples that must be in place at an appropriate site before jobs can run.
- We face 2 challenges to improve latency where our vast monitoring data can make a big impact through data analytics:
  1. **Finish what we start**: Improved job scheduling using actual performance statistics and capabilities
  3. **Eliminate tails**: Identifying and fixing slow or stuck requests is a manual process

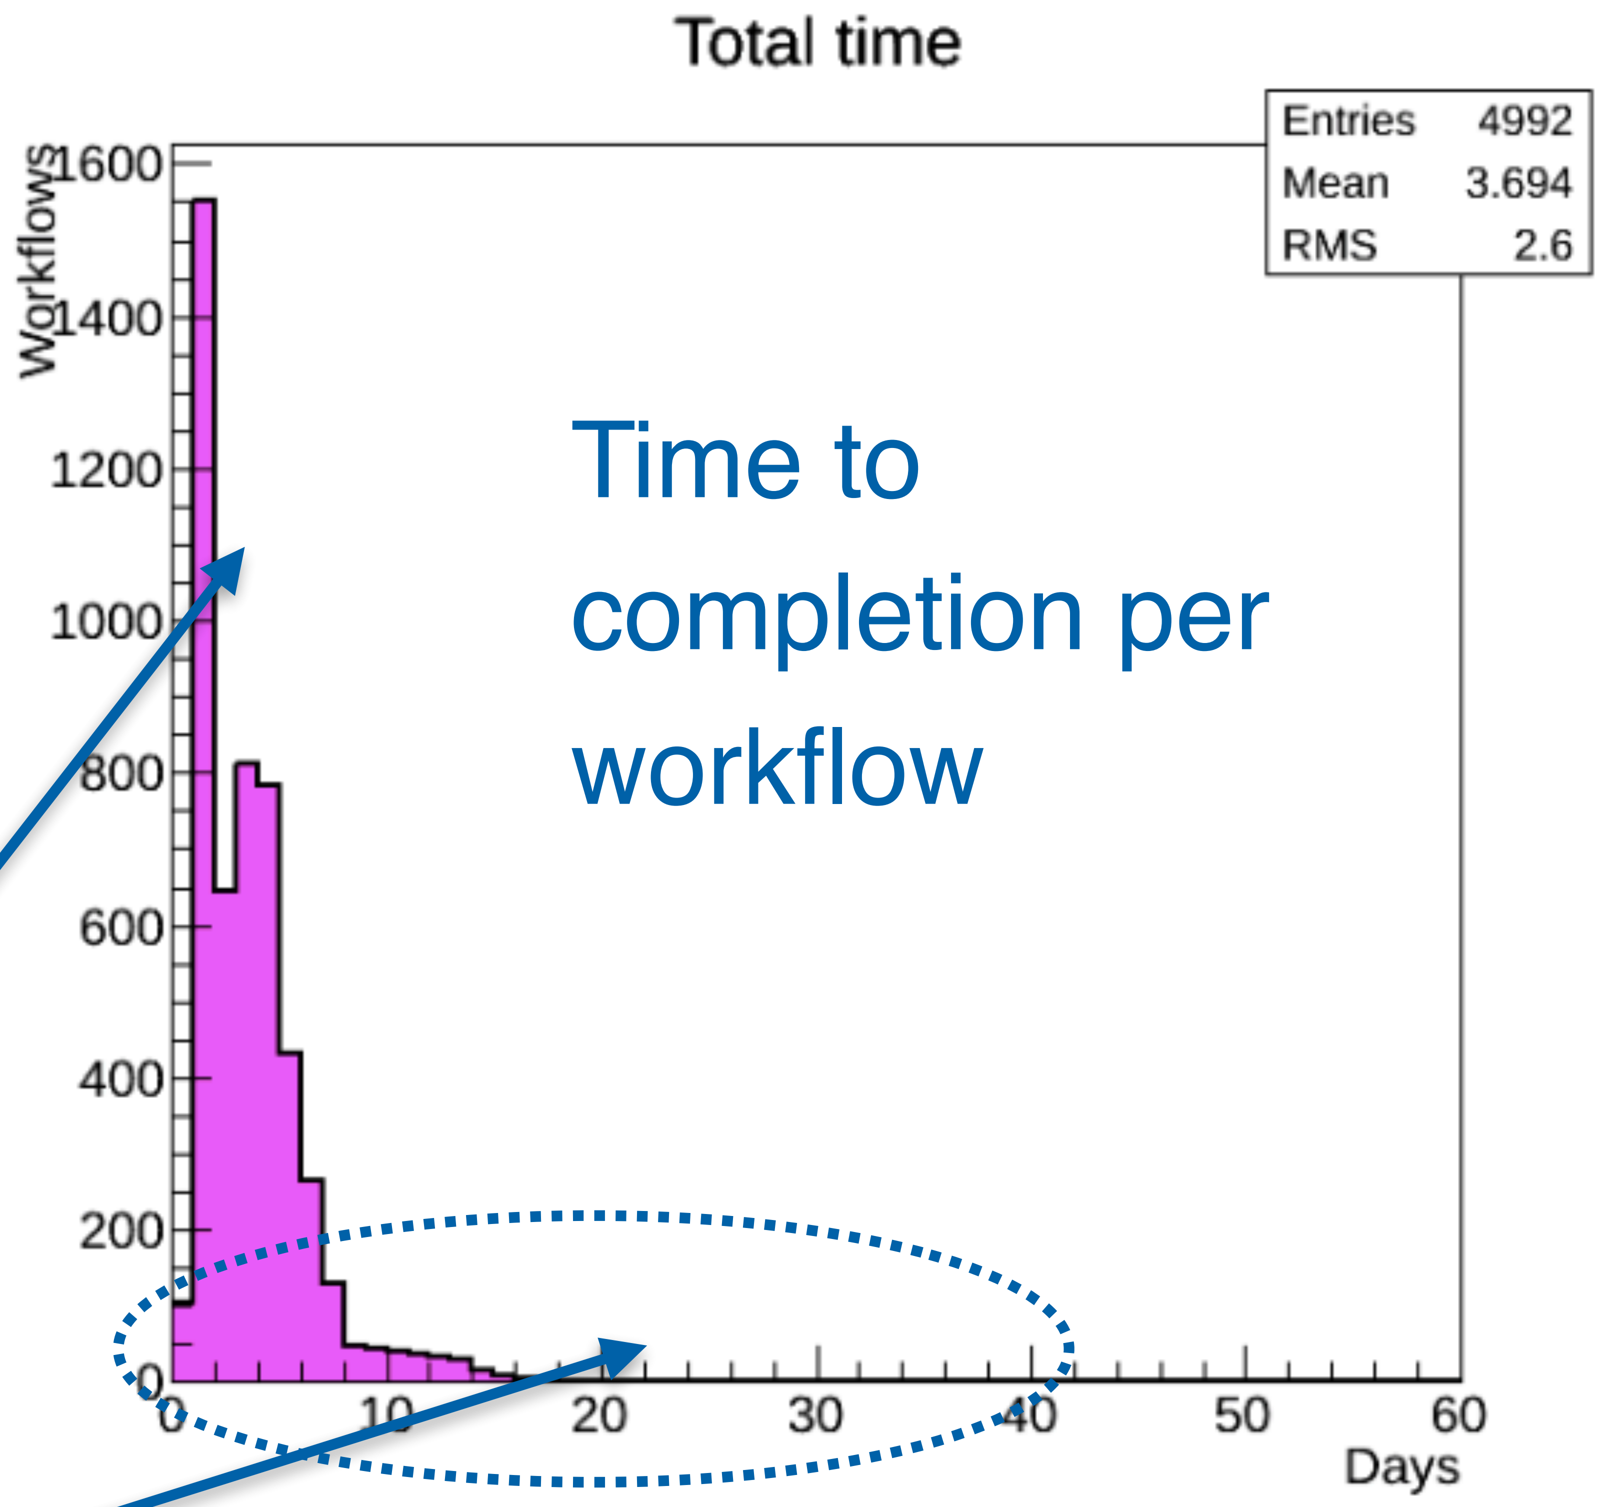

- We propose to use machine learning techniques to try to predict how CMS data should be placed and processing resources scheduled in order to achieve a dramatic reduction in latency for delivering data samples to CMS analysts.

Challenge:
Event selection and categorization

- CMS recordes collisions at 40MHz
  - Hardware triggers identify 100 kHz of potential interesting events
  - In ~160ms the higher level trigger further reduces this to **1kHz events to collect**

- Is there information we could get out of the 99kHz of events we currently throw out?
  - We have a model that explains most of what we see ➜ Standard Model of Particle Physics, we have simulations to train with.
  - The things we can't explain in the model are the most interesting, but hard to design triggers for

- This would allow two interesting changes:

  1. The first is to extract information from events that would otherwise be rejected.
     - Even unselected events, if identified, could be used for background distributions, cross section measurements, and trigger efficiency measurements.

  2. The second idea would be the more ambitious idea that uncategorized events might potentially be the most interesting.
     - Leading to more efficiency trigger design and lower risk of missing an unexpected discovery.

- Trained systems for event classification are now being used in a variety of fields. Goal of this project is to identify the feasibility of near real time event classification of events at the HLT level in CMS using supervised machine learning.

Challenge:
Exploiting data popularity to model the evolution of data management

- CMS uses a Dynamic Data Placement system (DDP) to transfer production and user-based data to specific sites for local access.
  - The information needed to perform these choices is based on acquired statistics on dataset accesses (popularity).

- We propose to extend this system and use Machine Learning techniques to
  - **predict popularity of newly created datasets** and
  - **their intelligent placements** across CMS sites

- Such prediction can complement DDP
  - **as initial data seeding** and
  - **reduce operational cost** both in terms of data transfer and data residency at sites.
    - can be studied separately for both long-term resident datasets as well as newly created ones

- In a further stage, such studies can then be extended and combined into analysis of job scheduling.

Challenge:
Reduce latency of end user analysis in HEP using Big Data Technologies

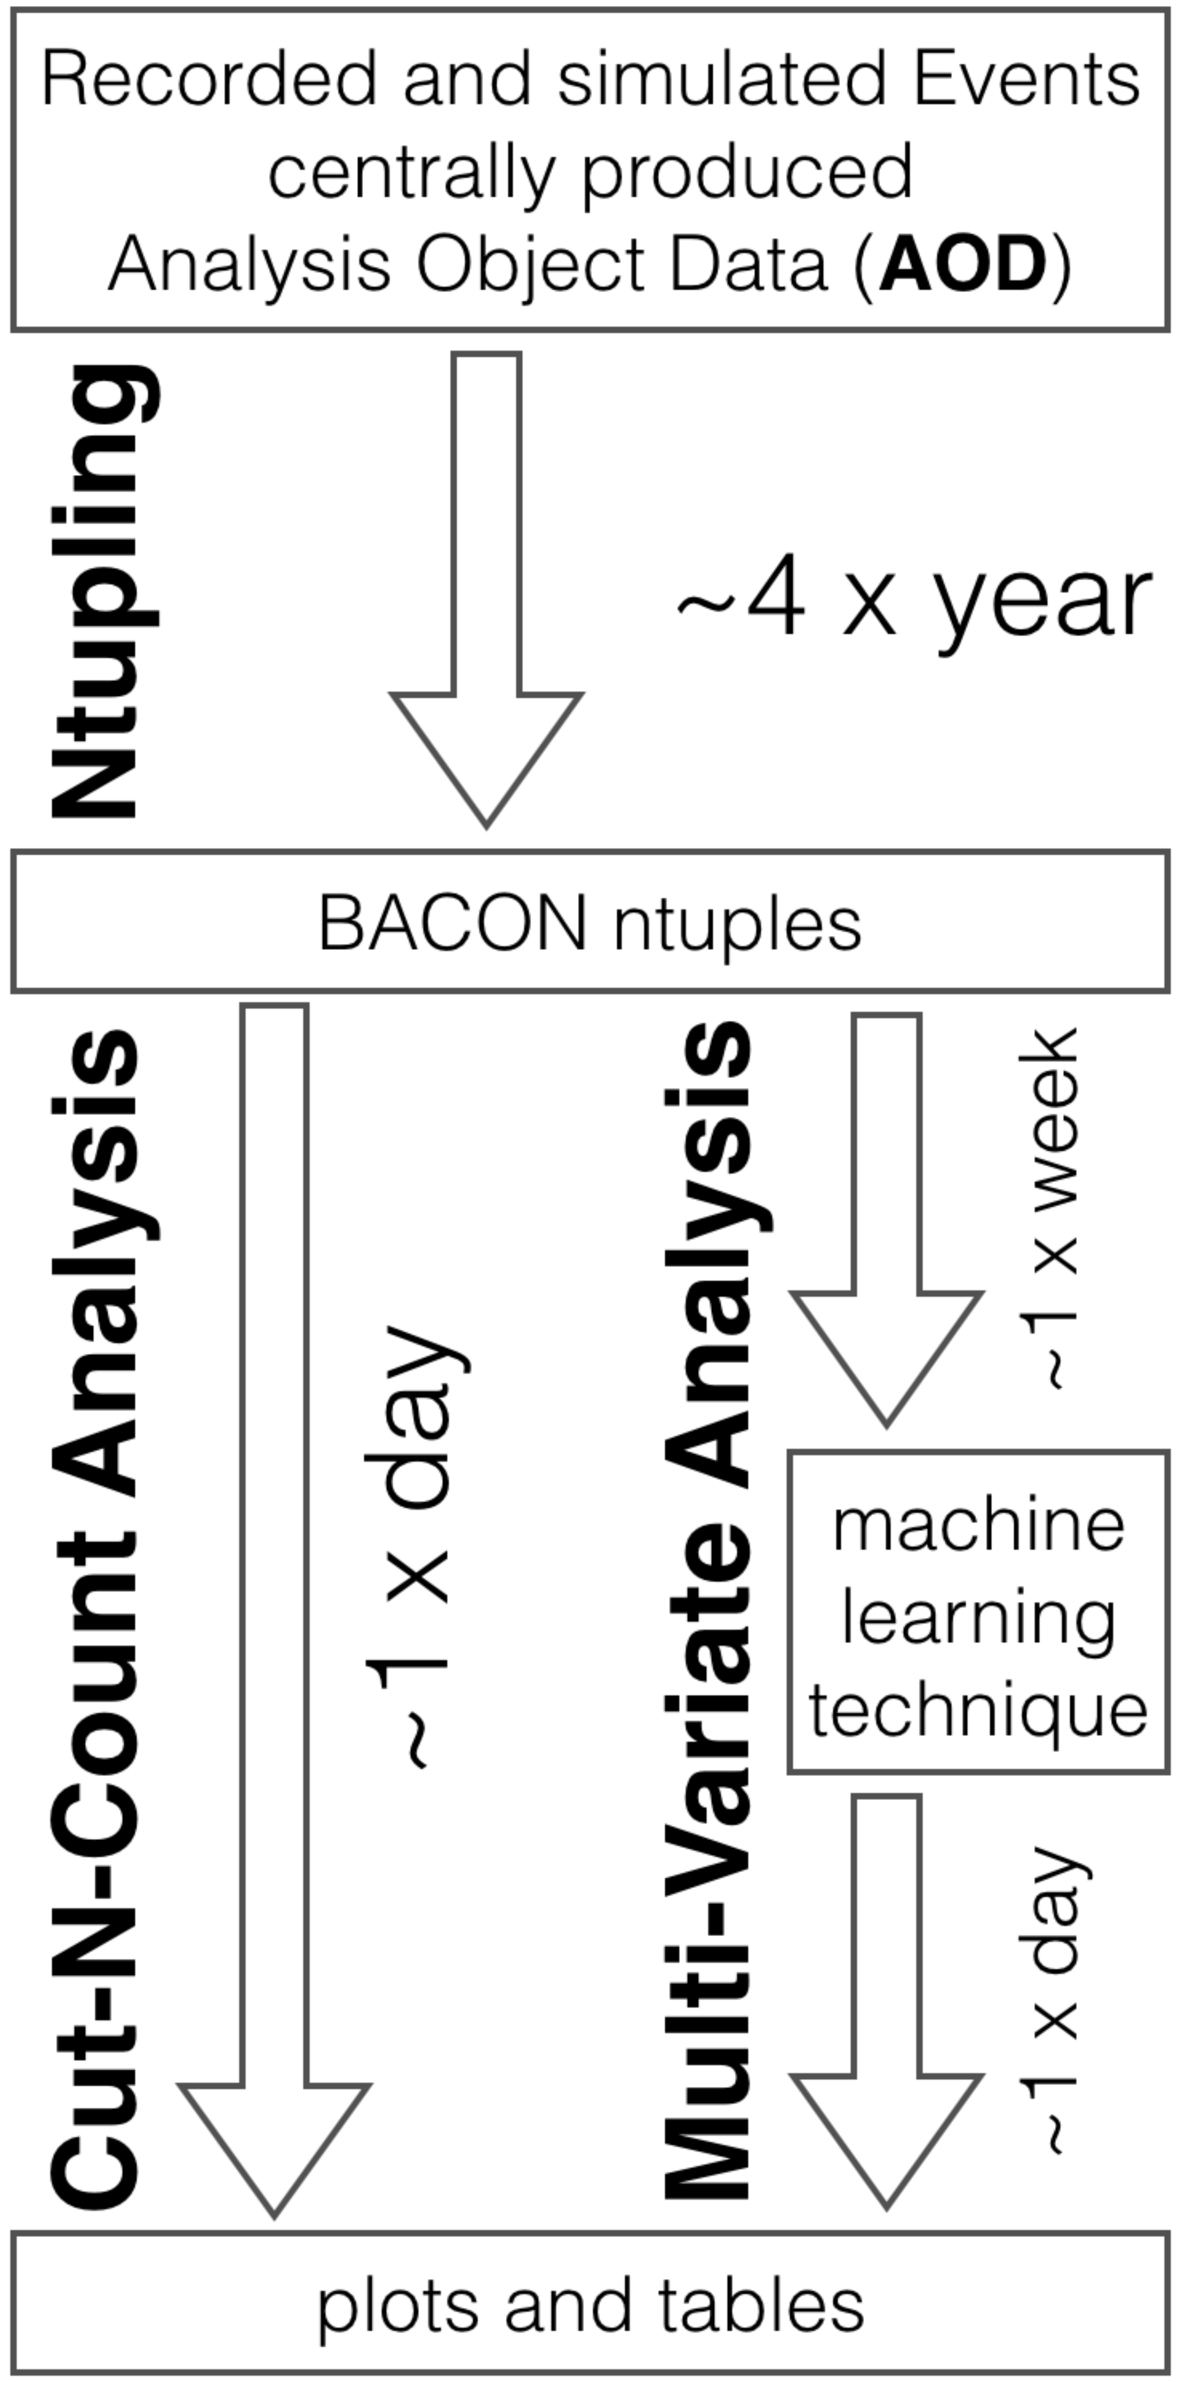

- High Energy Physics (HEP) analysis ➜ finding "the needle in the haystack"
  - HEP has been successfully filtering the vast data of the past and current accelerators for many years
  - Being one of the first disciplines handling Big Data, HEP developed own techniques and software approaches to solve these problems.
    - HEP has stayed mostly isolated from the Big Data developments in industry that have big success these days
  - Analysis of particle collisions is an experimental art based on the rapid investigation of hypotheses.
- Analysis currently is a multi-step process that takes scientists weeks to months to complete.
  - In the future HEP data volumes are increasing even further, and analysts adapt but stay within the paradigms of the current solutions.
- We would like to investigate areas where industry technologies can potentially revolutionize HEP analysis using concrete real-life use cases. The goal is to
  - Reduce the time from an idea to a physics result and to increase the flexibility and interactivity of HEP analysis.
  - Analyze petabyte-size data volumes with low turn-around and enable many scientists to explore the data in different ways using the same system in parallel or batch mode.