EF Computing Resource Needs

Peter Onyisi

31 Aug 2021

The University of Texas at Austin

ATLAS EXPERIMENT

On Projections and Feasibility…

- M. Delfino, “Computing at LEP,” CHEP 1991:

> disk buffer, by the Online computers. Note that, using a 50% machine duty cycle, this corresponds to a data accumulation rate of over 2 GB/day, which is substantial. The raw

> The usual data set that has been used for analysis in a LEP detector so far is in the order of dozens of GB, which is fairly large.

- FCC-hh (with prior FCC-ee) is further from today than we are from 1991

Introduction

- Resource requirements set by detector readout channels, luminosity, trigger system, compression …
  - i.e., more speculative the less on-shell the design of a detector is
  - very developed projections from ILC; FCC-hh is speculative
- Will focus on disk requirements for raw data
  - CPU needs for reconstruction & simulation are closely tied to specific choices of the collaborations
  - LHC experience suggests analysis formats may be an extremely important part of resource needs (e.g. for even aggressive ATLAS HL-LHC projections, analysis formats ~ raw formats)
- Was asked to talk about MC but this is quite tricky: strongly depends on experiment computing model & final analysis format

Storage Technologies

- Relevant players: magnetic tape, hard drives, solid state drives
- Tape is cheapest, but not really random-access: very powerful for archiving
  - significant improvement in recording density possible relative to today
  - loss of vendors: now basically IBM
  - total LTO tape shipped in 2019: 114 EB
- Hard drives are the traditional “random-access” technology
  - hitting physical limits on recording density of current technology
  - 1000 EB shipped in 2020
- Solid state drives are true random-access
  - most expensive per TB currently, but may converge with hard drives at some point
  - 200 EB shipped in 2020

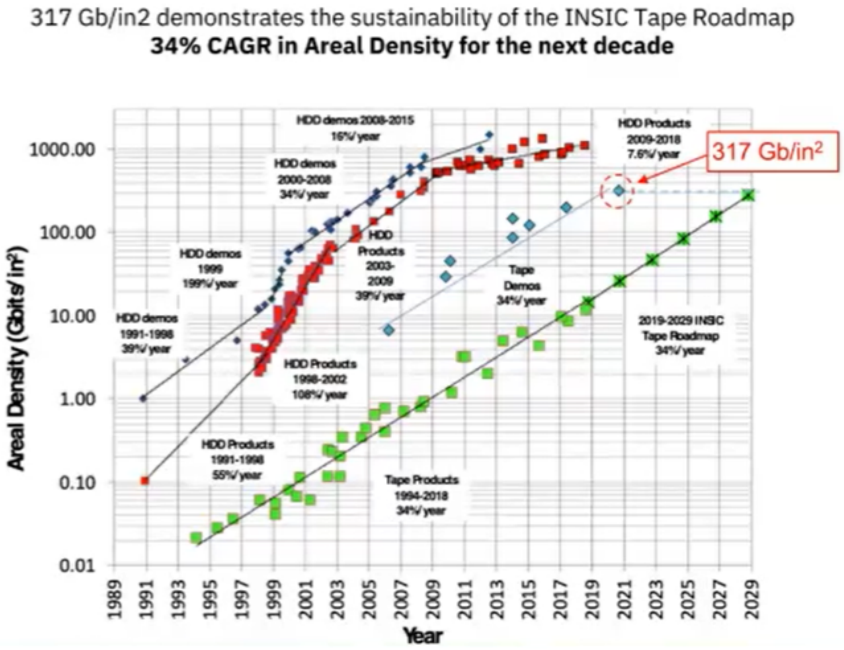

Some Characteristics

Lepton colliders:

- Electroweak processes: relatively small & democratic cross sections (modulo t-channel, γγ processes)
- Very loose triggers in general
- ~ order of magnitude more total MC events than data
- Operation on Z peak presents unique challenges

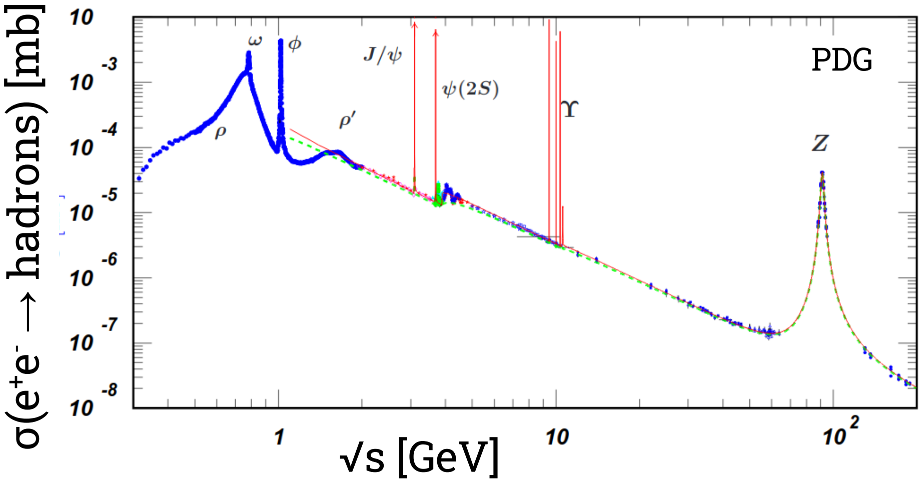

Hadron colliders:

- Extremely broad range of process cross sections
- Very tight triggers
- Simulations biased towards electroweak/rare processes; total MC events ~ a few times data

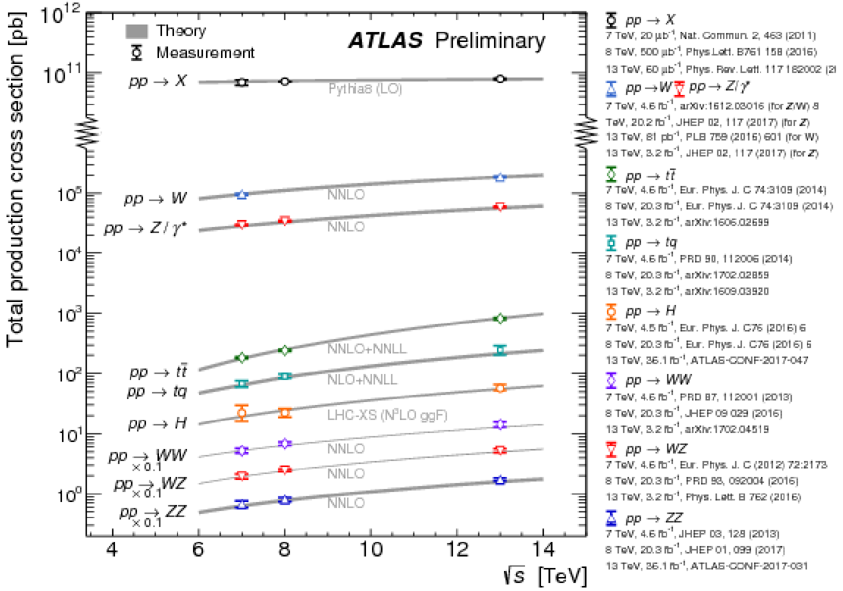

Broader Questions

- Raw data sizes per event depend on the data acquisition model
  - not every channel is written out!
  - "smarter" hardware may be able to do more aggressive reduction of data before it is recorded: but unknown new physics motivates keeping more raw data
- The number (and format) of events written depends on trigger model
  - writing partial raw data, only keeping high-level reconstructed information, …
- For redundancy raw data are replicated
  - 2× factor on storage needs is included here
- Use of disk vs "colder" storage mechanisms (i.e. tape) is a knob that can be optimized
  - future of tape technology is uncertain, cold stores in 2035 may look very different from now
- Monte Carlo "raw data" may not be kept: depends on expense of simulation & tradeoff with storage
  - already the case for ATLAS & CMS
  - fast simulation makes the value of "raw MC" even lower
- Storage needs scale by number of interaction points in design

A Baseline: ATLAS @ HL-LHC

- ATLAS forecasts an increase of ~3× in event size and ~8× in event rate at μ=200 compared to 2018
- MC/data ratio similar to now or slightly lower
- Main MC storage format is initial analysis format (AOD) or even more processed (and smaller)
  - MC does not exceed raw data, though it does significantly exceed data AOD
- LHC has two general purpose detectors + LHCb & ALICE : total complex needs are > 2× more than below

ATLAS HL-LHC Computing CDR

Event size	Event rate	Time/year	Data/year	MC factor	MC/year
4.4 MB	10 kHz	0.7×10^7 s	0.6 EB	2-2.5x	0.2 EB

Circular e^+e^- Collider

- Estimates from CEPC & FCC-ee
- Very different data rates for Z factory and high energy operation
  - high energy is a minor perturbation on Z factory operation; any system sized for the latter will have no problem with the former
  - FCC-ee mentions value of triggerless readout in reducing Z systematics
  - handling of Bhabha scattering events for luminosity is critical
- Is 10× data statistics for MC enough for precision Z physics?
  - systematic variations etc. - or can everything be constrained in a data-driven way?

Scenario	Event size	Event rate (non-Bhabha)	Time/year	Data/year
FCC-ee Z-pole	1 MB	~ 100 kHz	10^7 s	2 EB
CEPC 240 GeV	20 MB	2 Hz	1.3×10^7 s	260 PB

FCC-ee CDR
CEPC CDR

Linear Colliders

- Low rates & relatively low backgrounds at ILC/CLIC permit extremely relaxed setups
  - just record entire bunch trains & pull out interesting events later
- ILD has detailed projections of MC & additional formats:
  - because "interesting" events are < 1% of bunch crossings, simulation is small (1 order of magnitude less) compared to data even with 10× statistics
  - the non-raw data (incl. MC) roughly double the total requirements

Computing for the ILD experiment
Detector Technologies for CLIC

Scenario	Train size	Train rate	Time/year	Data/year
ILD 500 GeV	178 MB	5 Hz	0.8×10^7 s	14 PB
CLIC 3 TeV	88 MB	50 Hz	1.2×10^7 s	110 PB

Muon Collider

- Quite speculative estimates
- Backgrounds from beam muon decays (~10^8 particles/event!) pose a major challenge
  - major challenges for simulation & tracking
  - CPU required is extremely important **now** for detector & machine interface studies – full simulation ~ 1 day/event
- Detector design has significant requirements from the suppression of these backgrounds (timing, granularity, radiation hardness …)
  - hence data rates very speculative
- Some early estimates: ~ 50 MB/event, write O(kHz) to tape
- MC needs probably more ILC-like than LHC-like (assume 10x)
- Numbers assume ~ 50× reduction in rate from trigger

N. Bartosik, CHEP 2021

Event size	Event rate	Time/year	Data/year
50 MB	2 kHz	10^7 s	2.0 EB

Hadron Collider

- FCC-hh numbers are extremely speculative
- Very different scale of challenge in DAQ compared to other proposals: e.g. full tracker data rate 1-2 PB/s
  - hardware triggerless readout much more tricky than for lepton colliders
  - limitation is data links between hardware and trigger farm, as well as power/cooling constraints
- Very high physics process rates, e.g. 42 kHz for $W \to \ell\nu$: potentially important decisions to make on prescales
- @ 50 MB/event, keeping raw output to 10 EB/year (with backup) limits high level trigger rate to ~ 10 kHz (!)

FCC-hh CDR

Event size	Event rate	Time/year	Data/year
50 MB	10 kHz	10^7 s	10 EB

Summary

- Most extreme requirements are probably set by FCC-hh
- Otherwise the requirements for future experiments are not wildly beyond those for HL-LHC
  - with exceptions of e^+e^- Z-pole and possibly muon collider
  - opportunity to be more aggressive? (people trying to design to LHC limits?) or just a natural limit from accelerator/instrumentation/DAQ constraints?
- One might expect a similar situation for CPU requirements: detector layouts quite similar to now
  - these needs in any case are probably much more dependent on chosen physics tradeoffs + future algorithm developments

Thanks: J. Strube, Y. Cheng, S. Pagan Griso, S. Jindariani, N. Bartosik, D. Lucchesi, A. Sailer, C. Helsens, G. Ganis