HPC in Particle Physics simulations

David Grellscheid

ICTP, Trieste /

Durham University

2014-11-

The Abdus Salam International Centre for Theoretical Physics

What are the fundamental building blocks of Nature?

Quarks

u up	***c*** charm	***t*** top
d down	***s*** strange	***b*** bottom

Leptons

ν_e electron neutrino	ν_μ muon neutrino	ν_τ tau neutrino
e electron	μ muon	τ tau

Force carriers

γ photon	
Z Z boson	
W W boson	
g gluon	

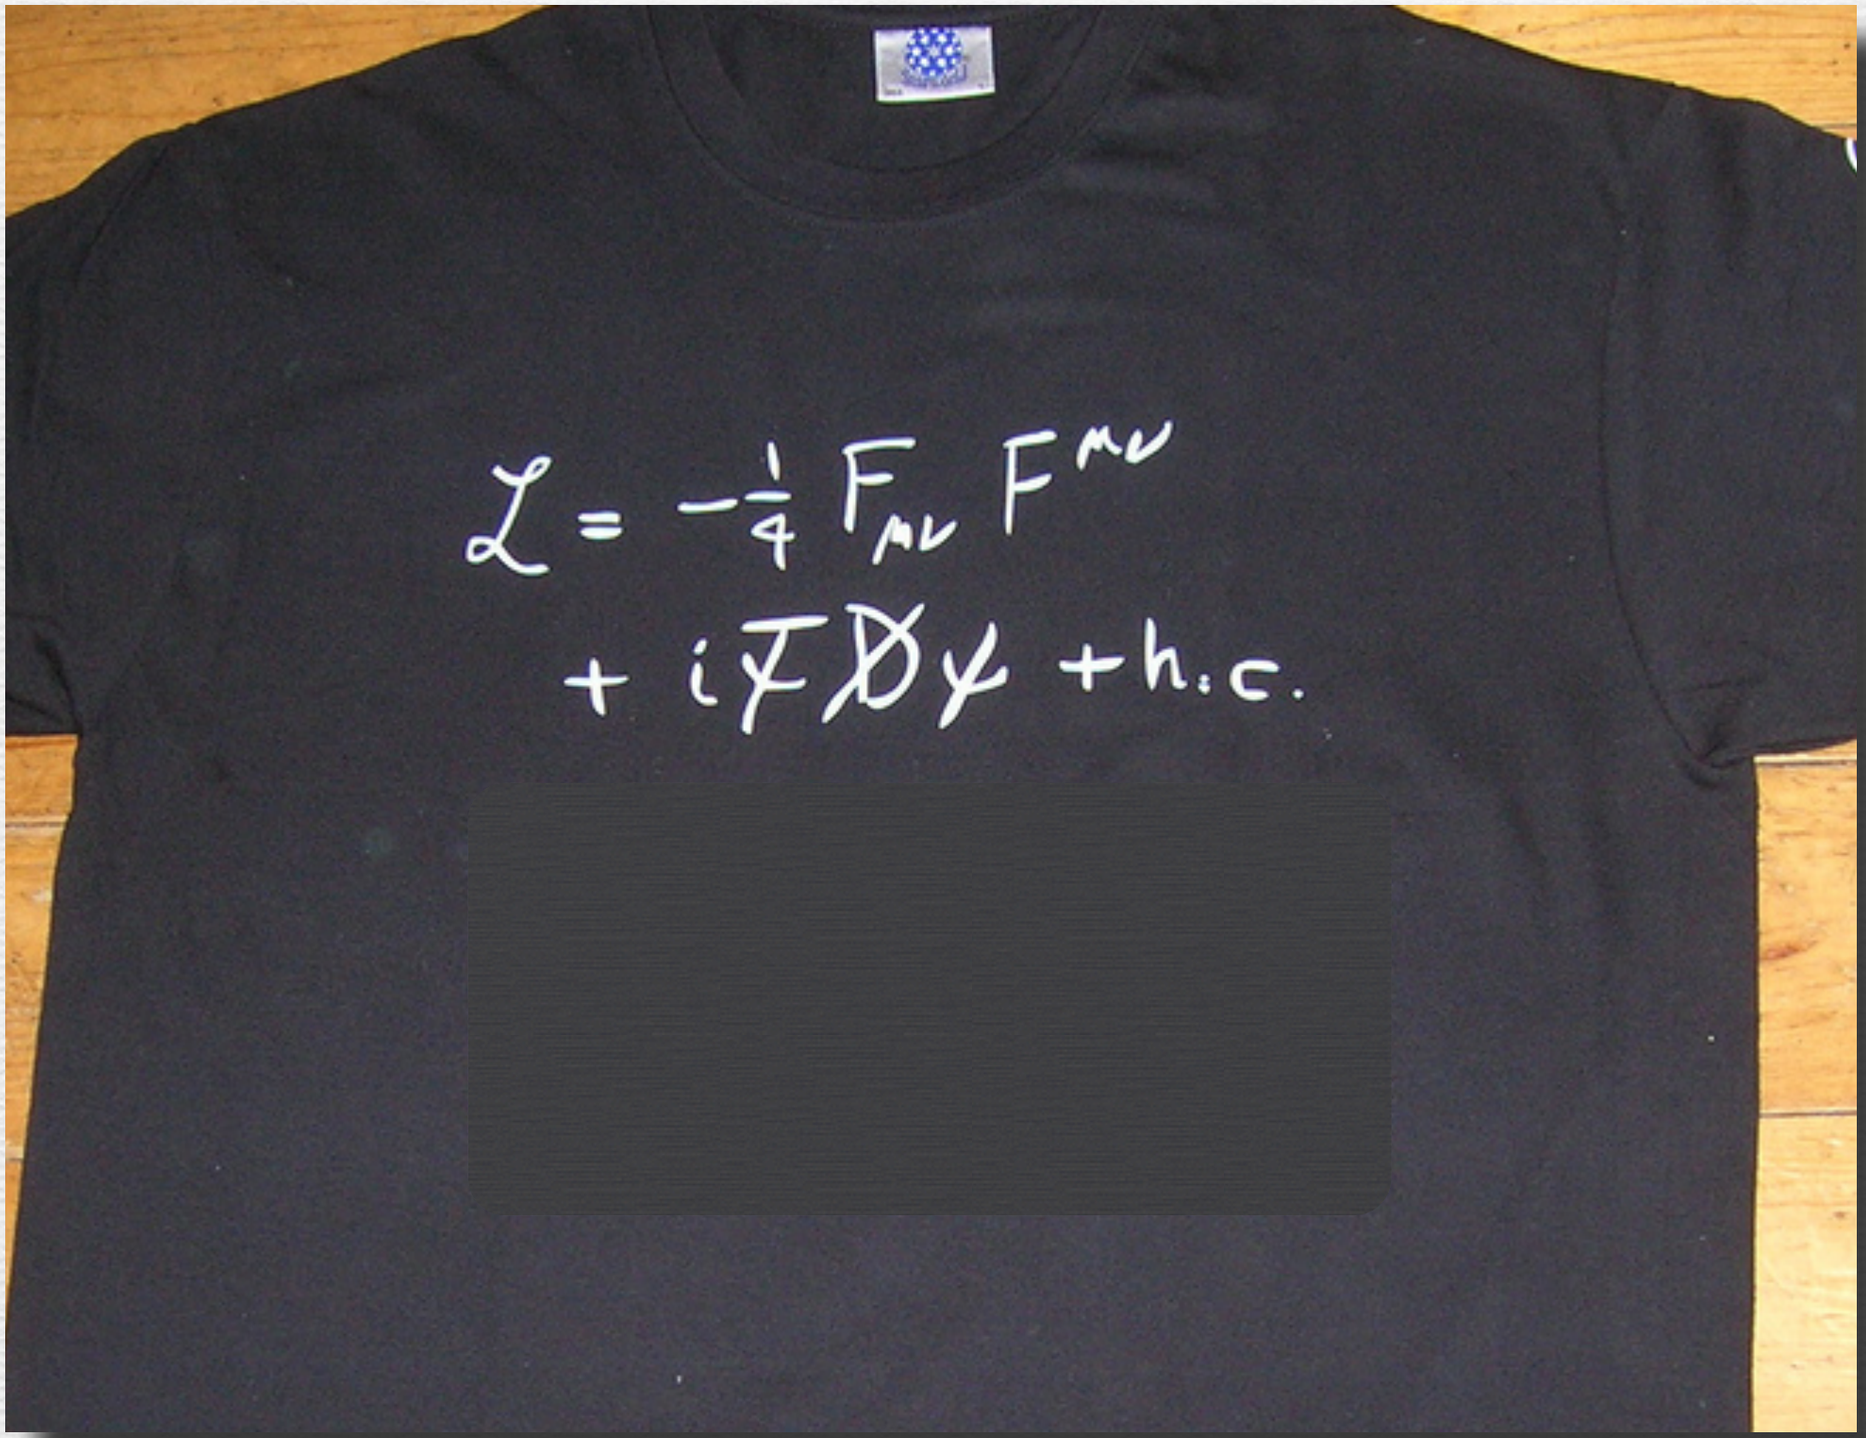
$\mathcal{L} = -\frac{1}{4} F_{\mu\nu} F^{\mu\nu}$
$+ i \not{F} \not{D} \psi + h.c.$

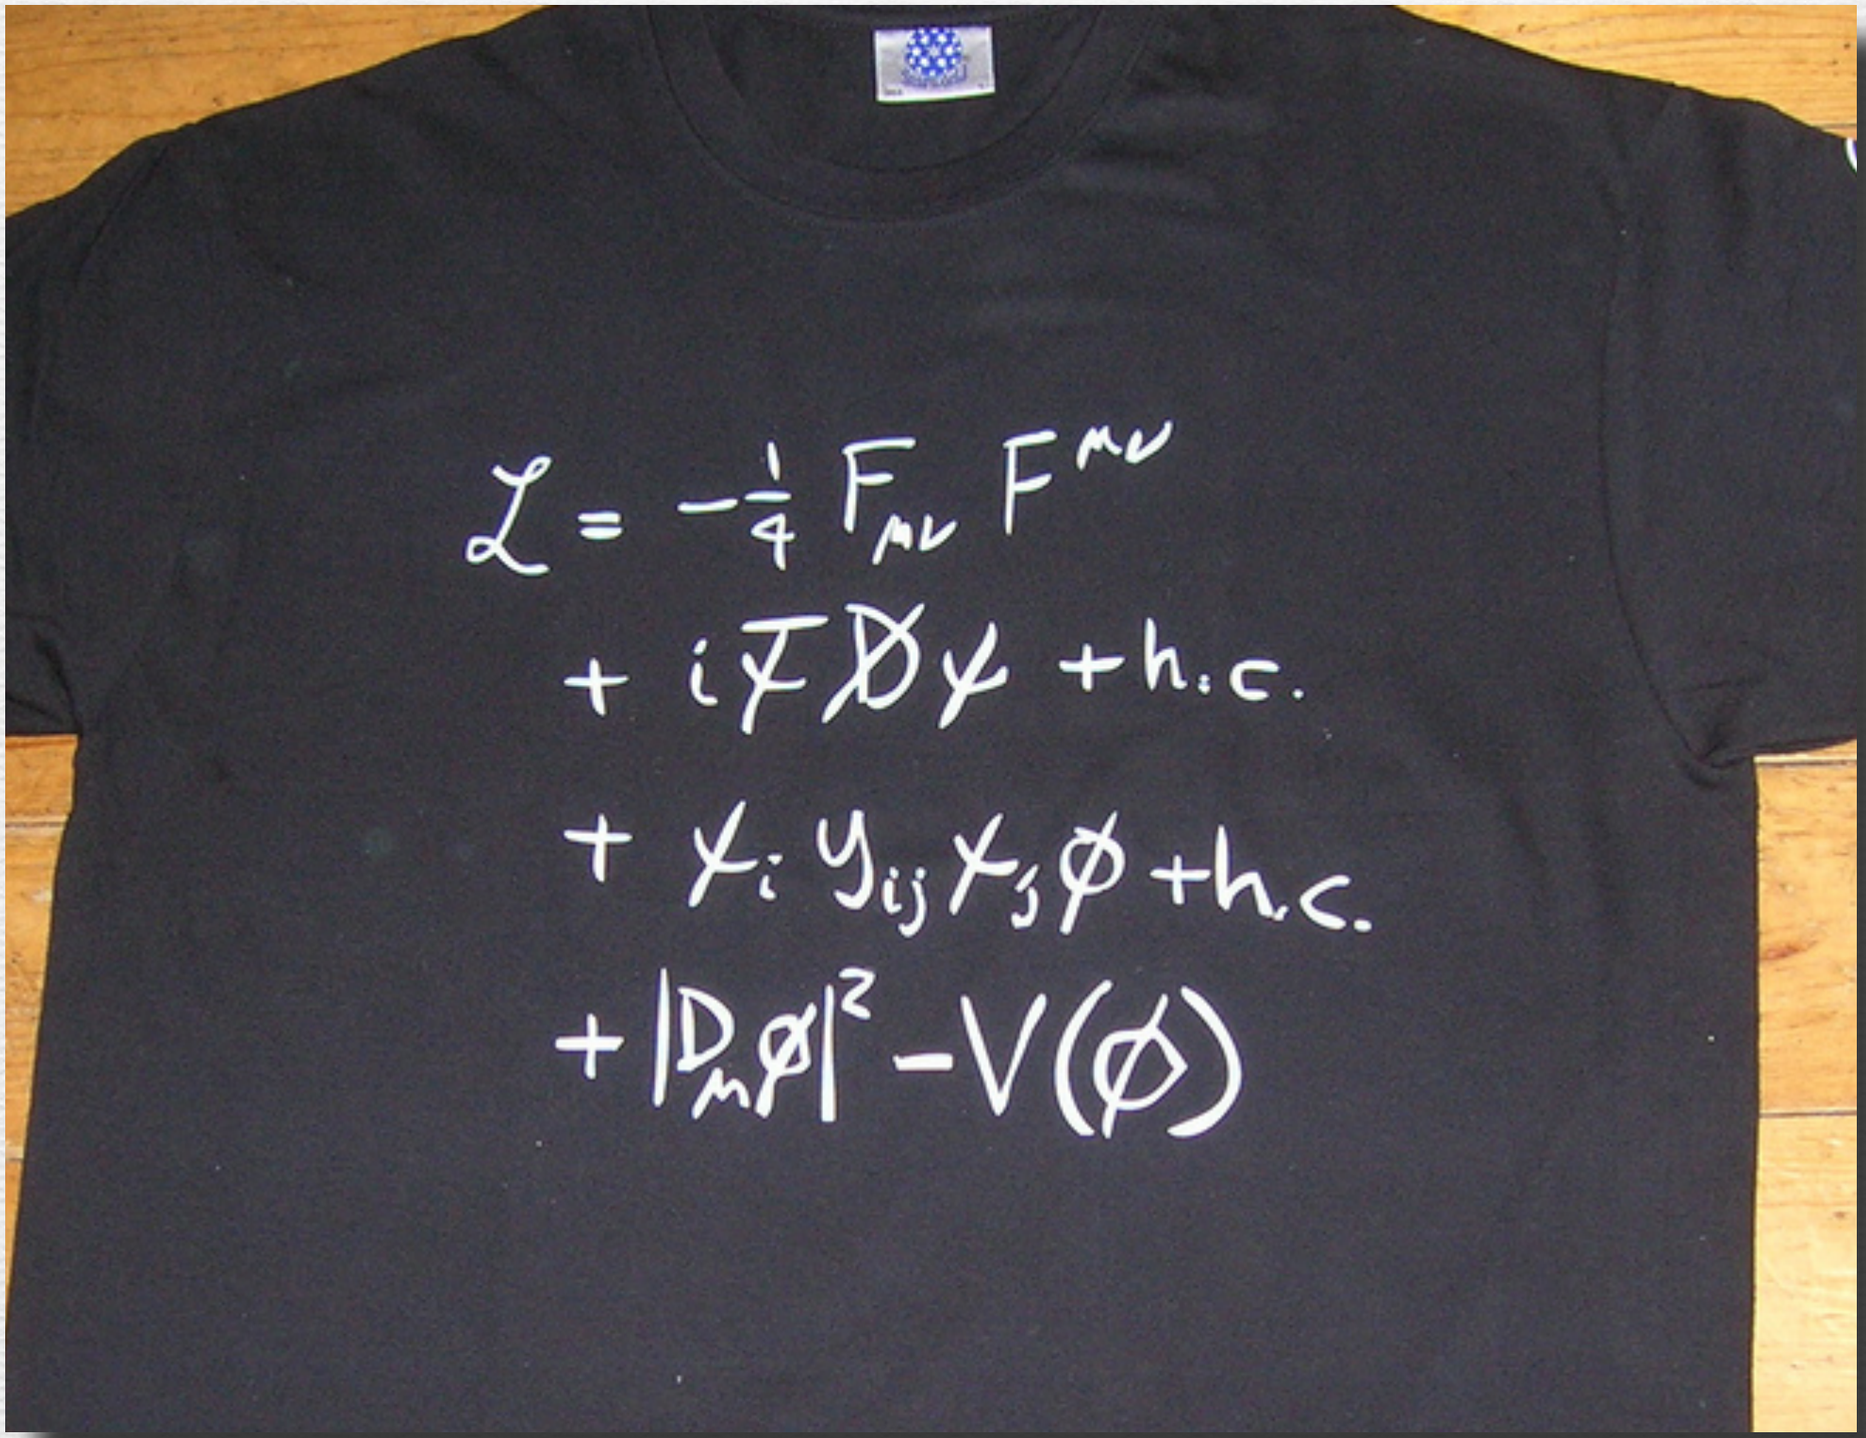
$\mathcal{L} = -\frac{1}{4} F_{\mu\nu} F^{\mu\nu}$
$+ i\bar{\psi} \not{D} \psi + h.c.$
$+ \chi_i y_{ij} \chi_j \phi + h.c.$
$+ |D_\mu \phi|^2 - V(\phi)$

CERN Accelerator Complex

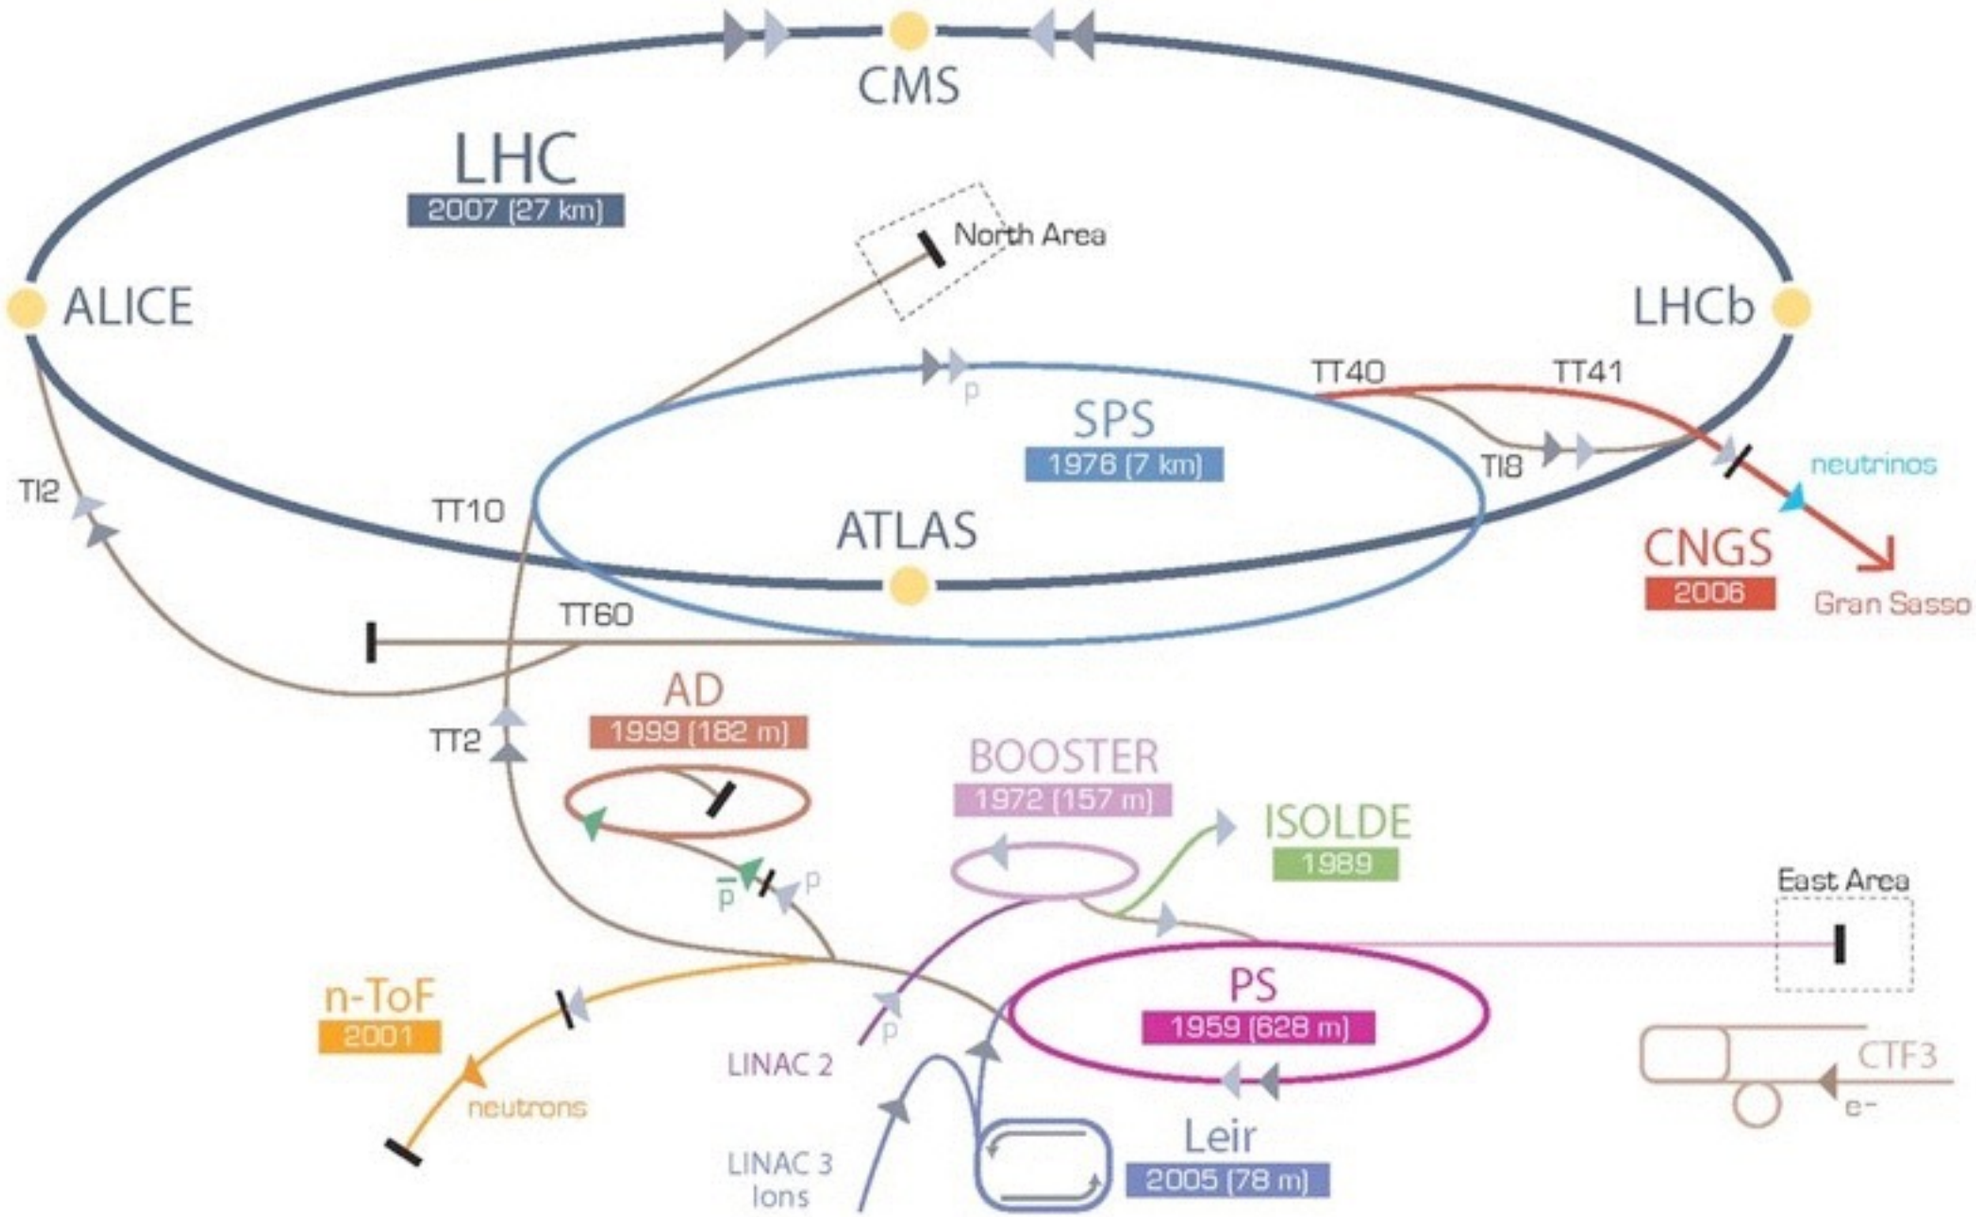

p (proton) ion neutrons p̄ (antiproton) proton/antiproton conversion neutrinos electron

LHC Large Hadron Collider SPS Super Proton Synchrotron PS Proton Synchrotron

AD Antiproton Decelerator CTF3 Clic Test Facility CNGS Cern Neutrinos to Gran Sasso ISOLDE Isotope Separator OnLine DEvice

LEIR Low Energy Ion Ring LINAC LINear ACcelerator n-ToF Neutrons Time Of Flight

CERN Accelerator Complex

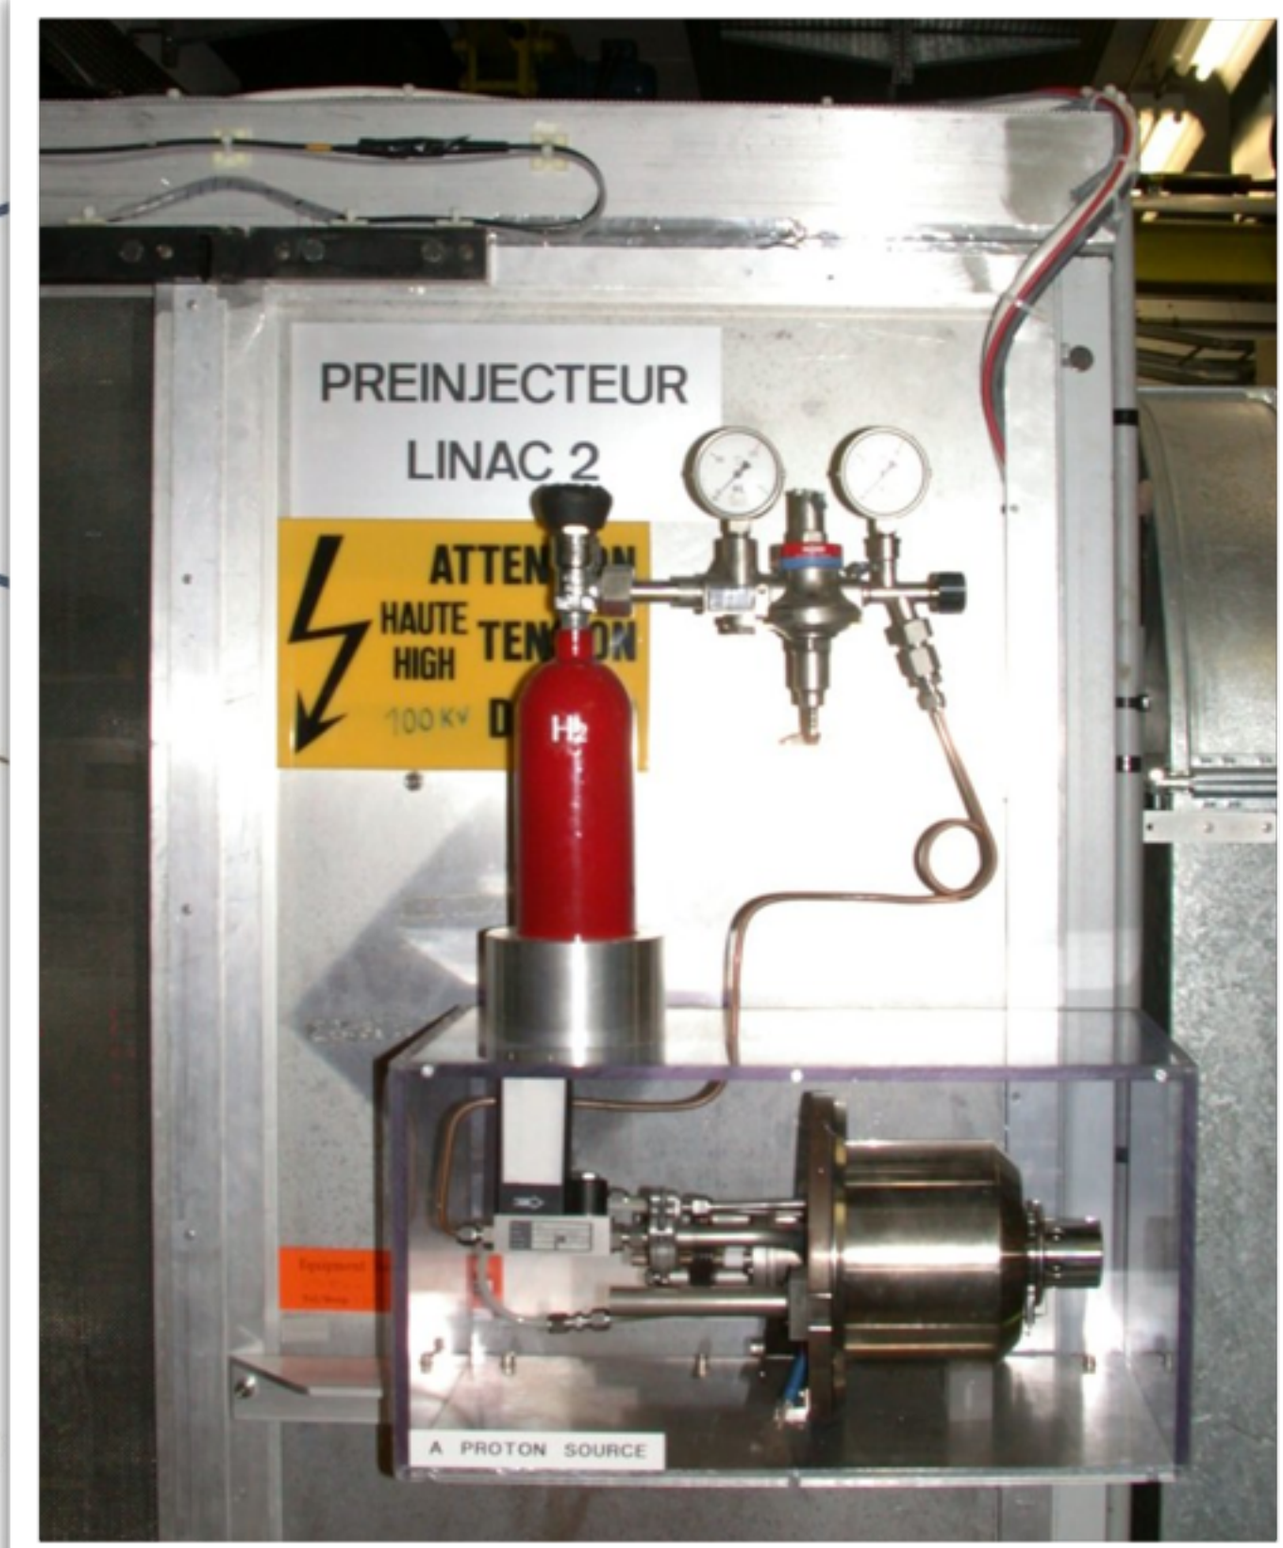

p (proton)

electron

AD Antiproton Decelerator CTF3 Clic Test Facility CNGS Cern Neutrinos to Gran Sasso ISOLDE Isotope Separator OnLine DEvice
LEIR Low Energy Ion Ring LINAC LINear ACcelerator n-ToF Neutrons Time Of Flight

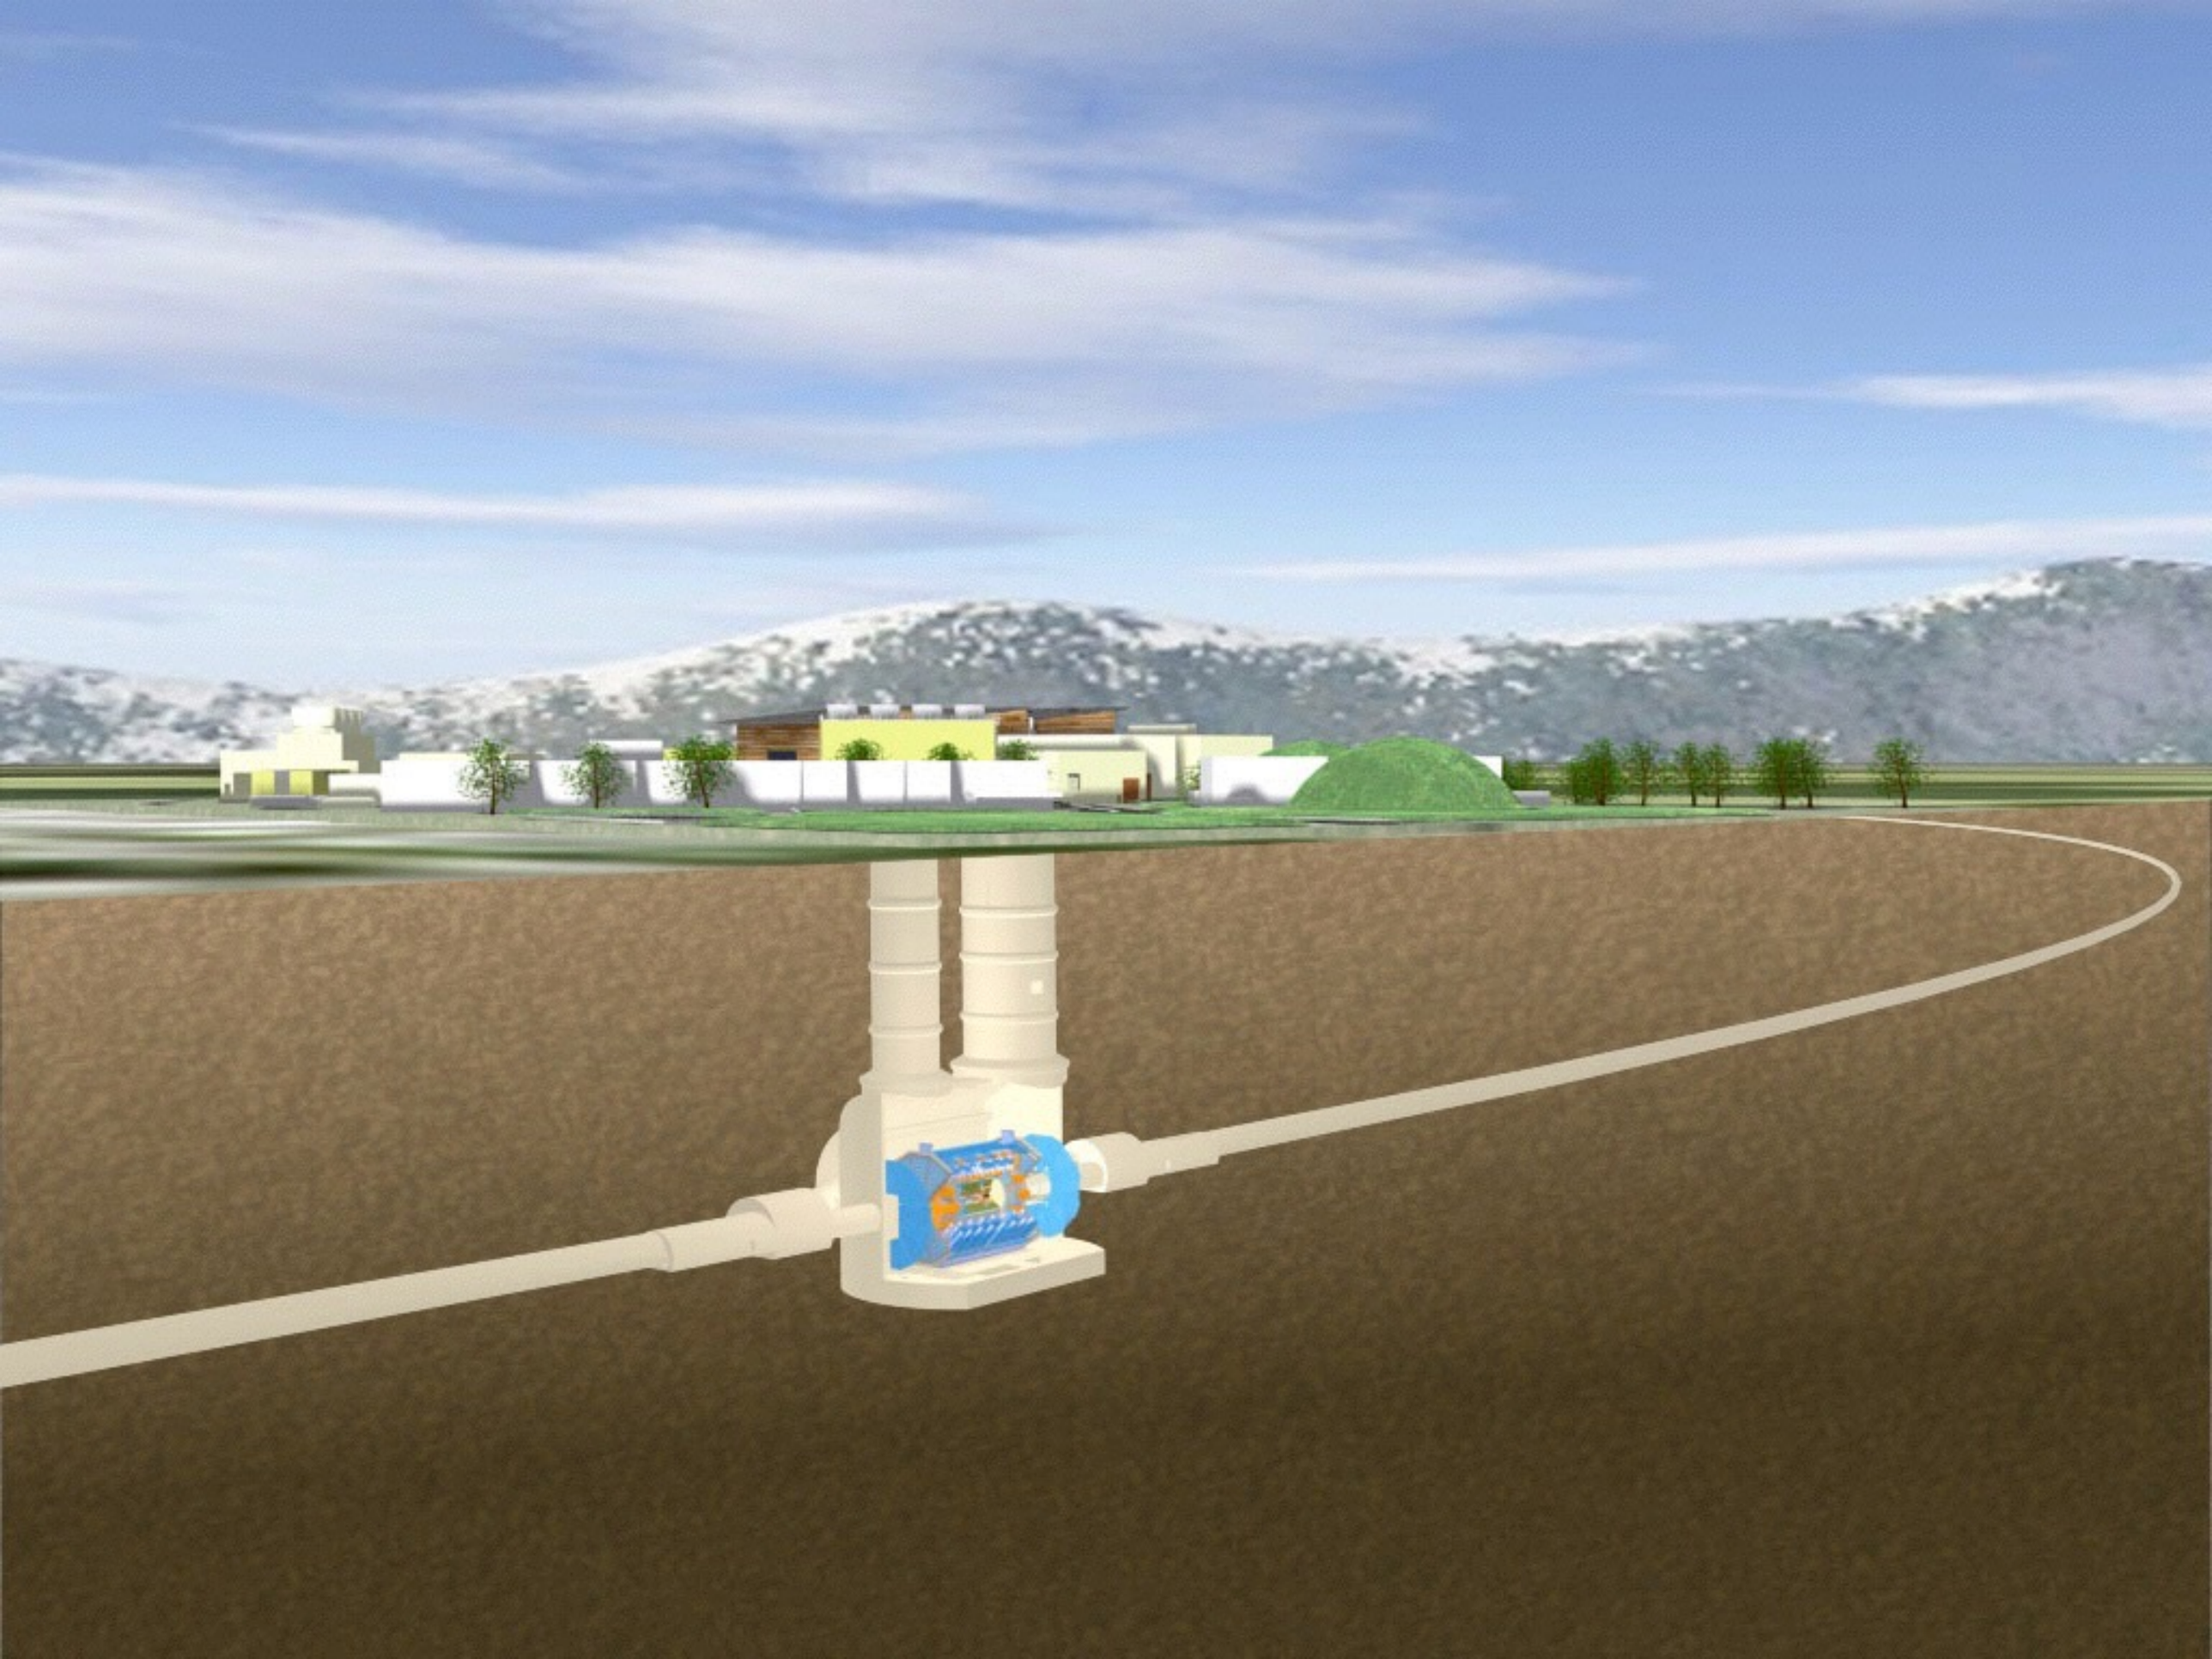

4μ candidate with $m_{4\mu}$= 125.1 GeV

p_T (muons)= 36.1, 47.5, 26.4, 71 .7GeV m_{12}= 86.3 GeV, m_{34}= 31.6 GeV
15 reconstructed vertices

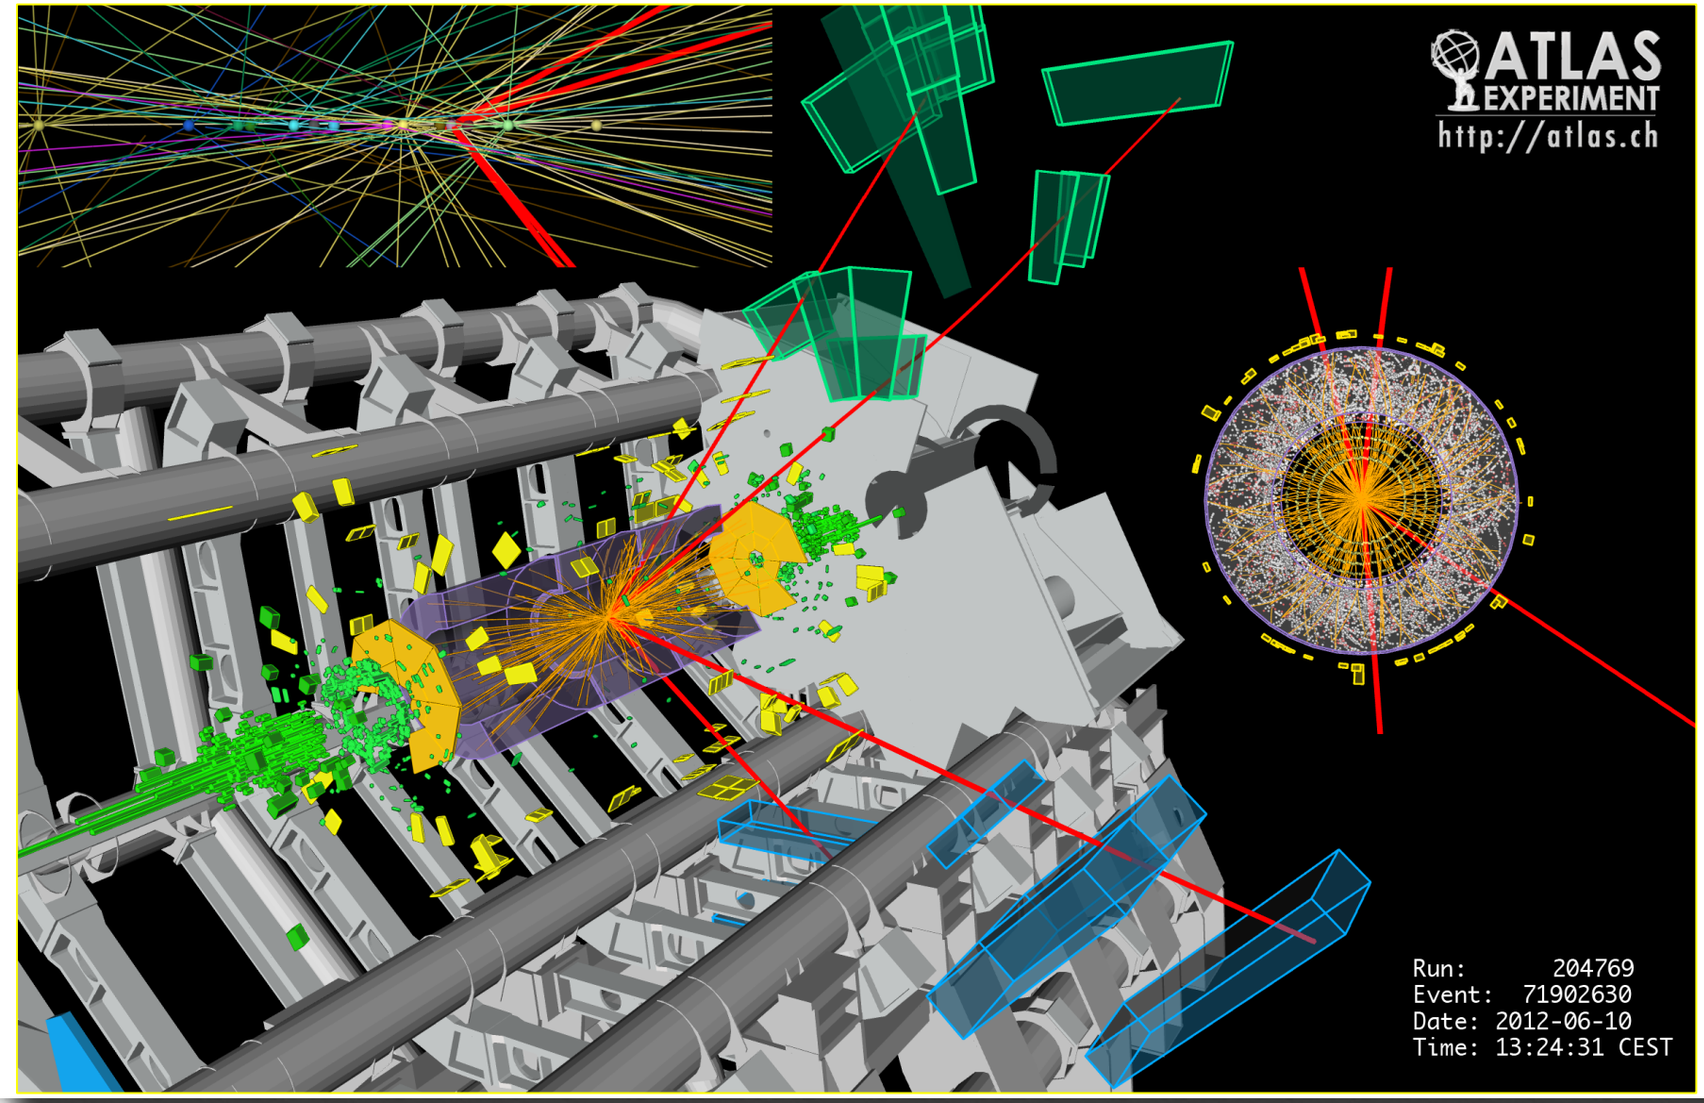

4e candidate with m_{4e}= 124.6 GeV

p_T (electrons)= 24.9, 53.9, 61.9, 17.8 GeV m_{12}= 70.6 GeV, m_{34}= 44.7 GeV
12 reconstructed vertices

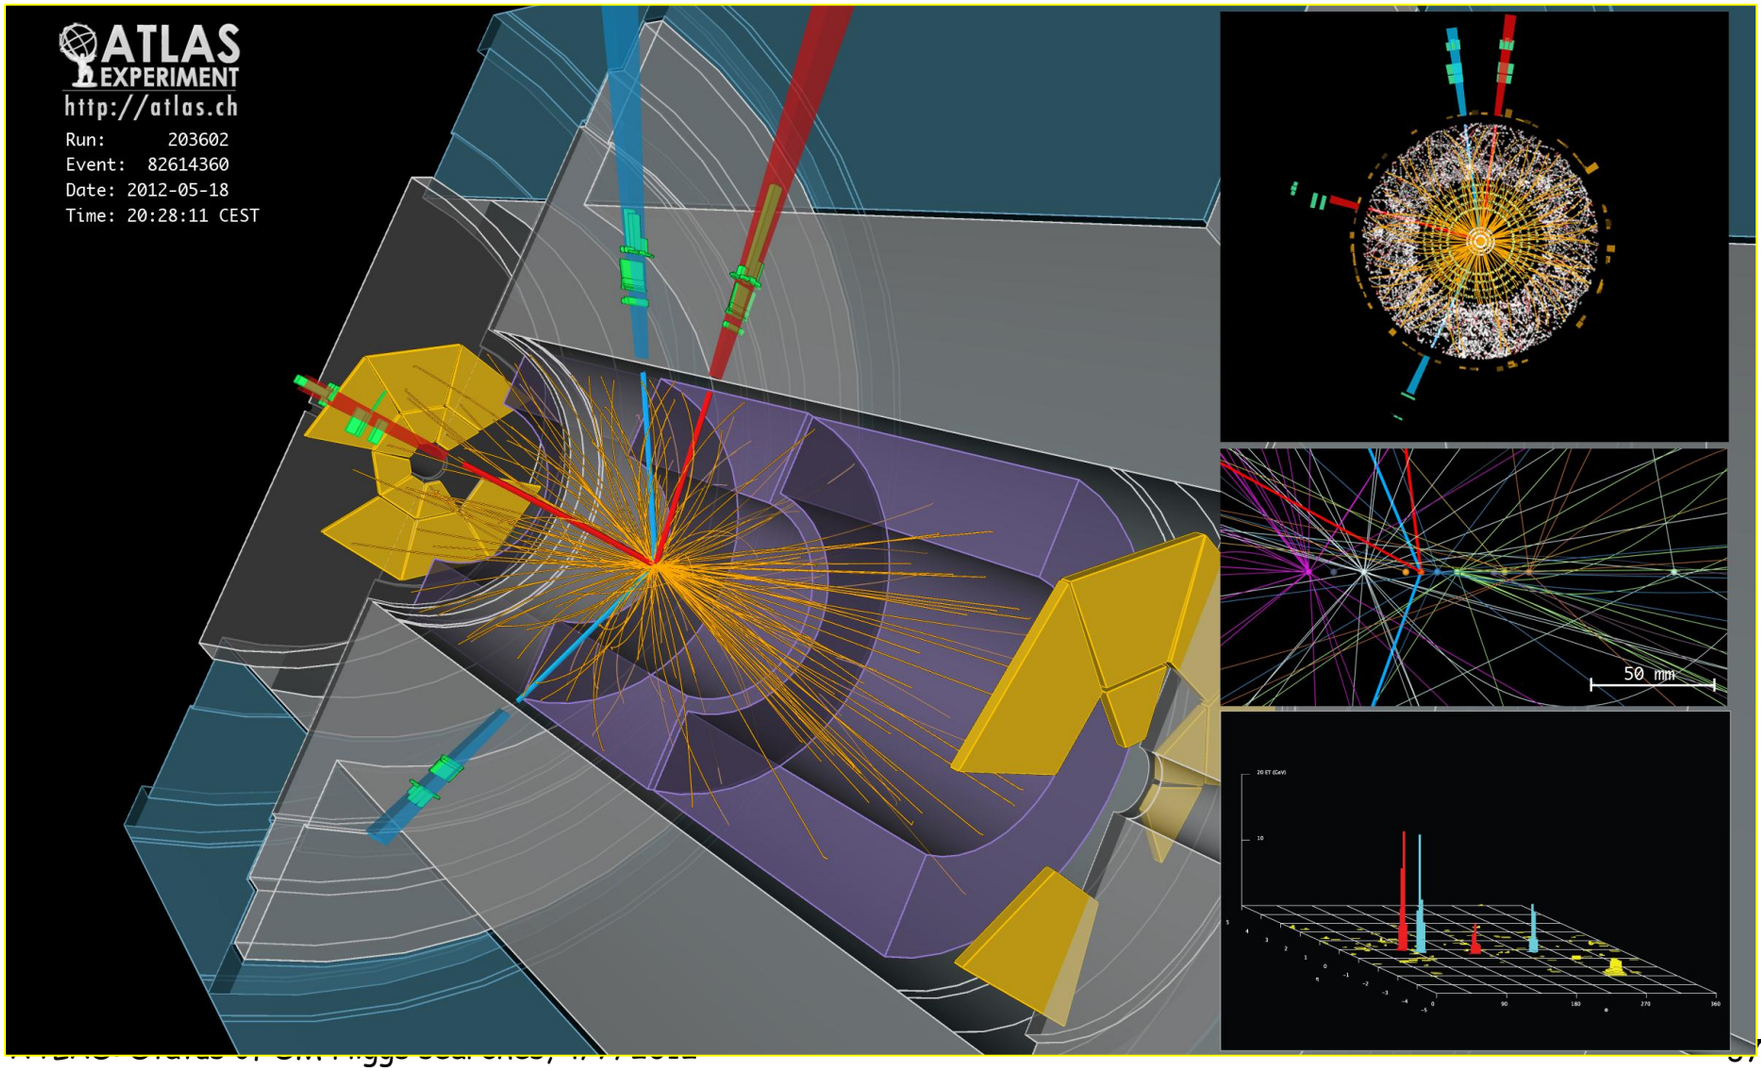

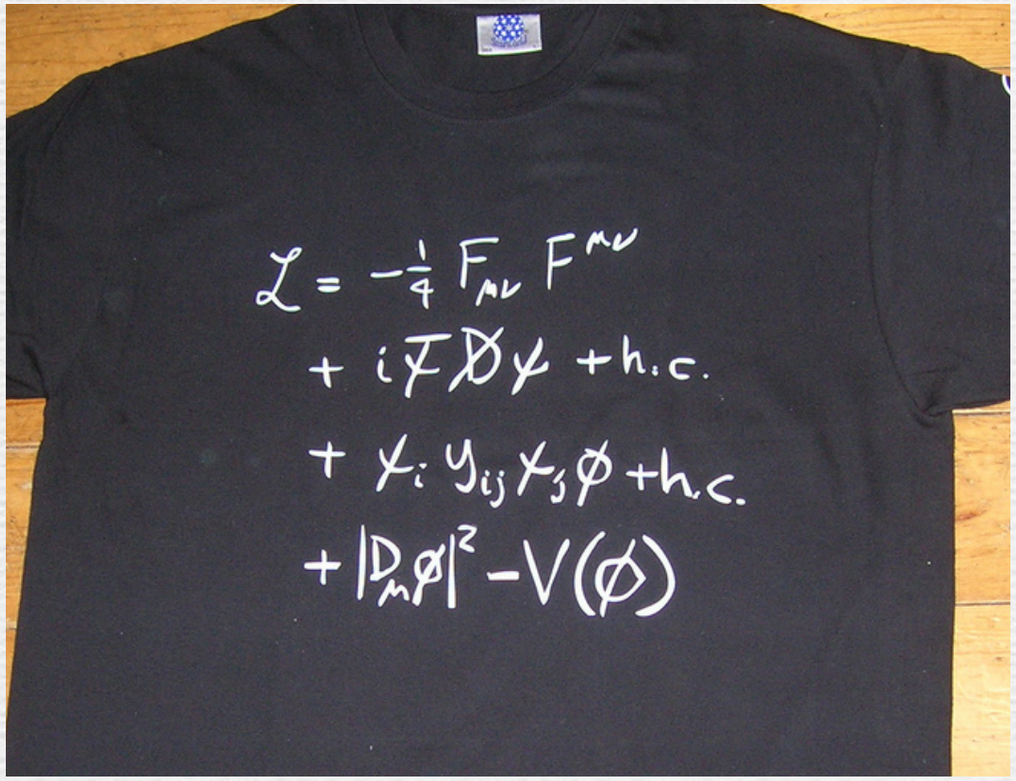

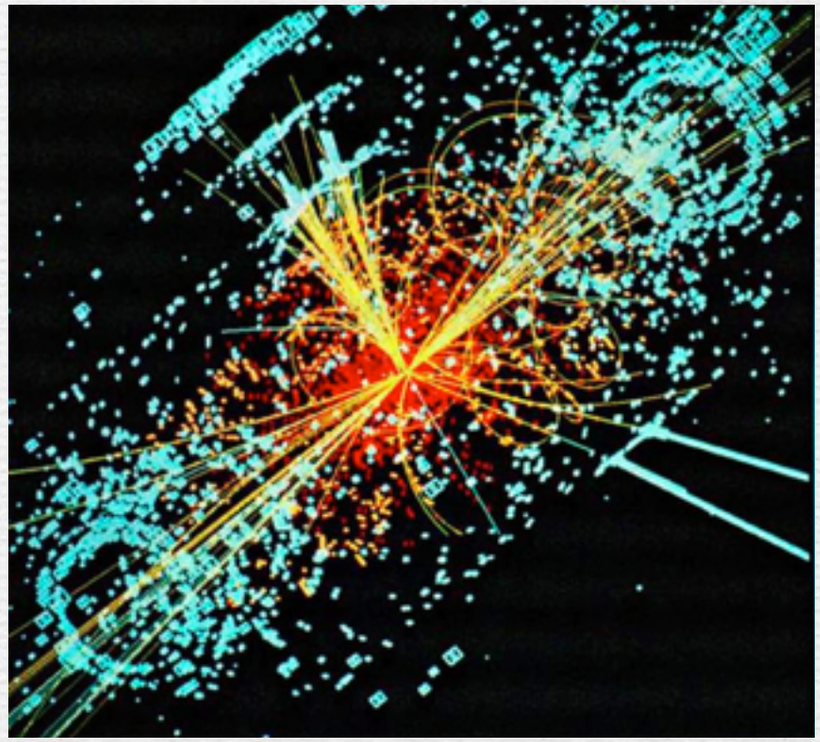

Theory

Experiment

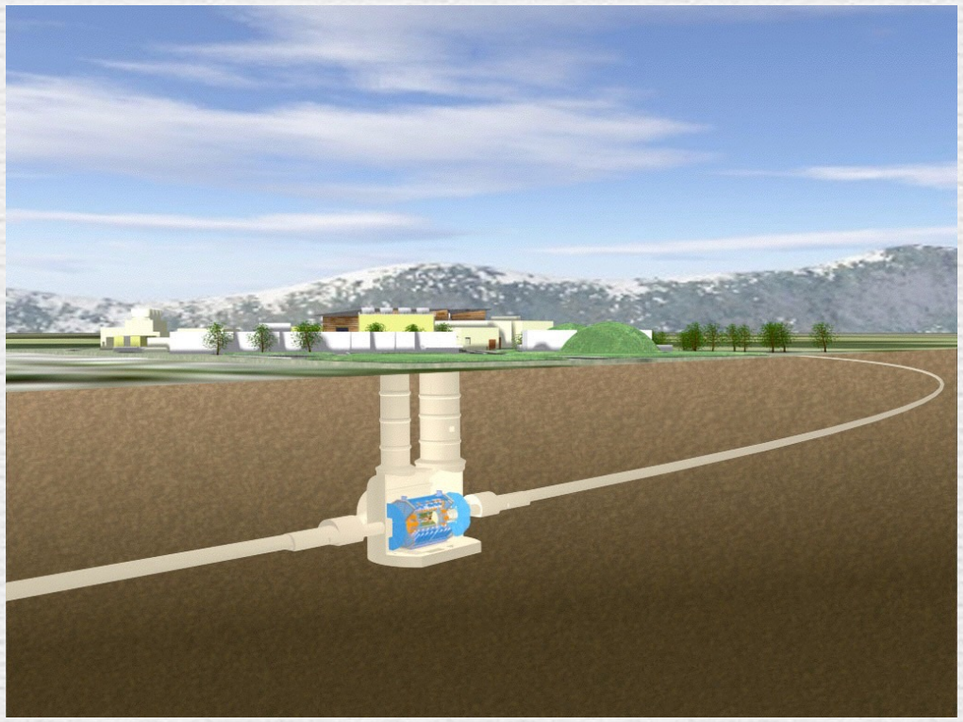

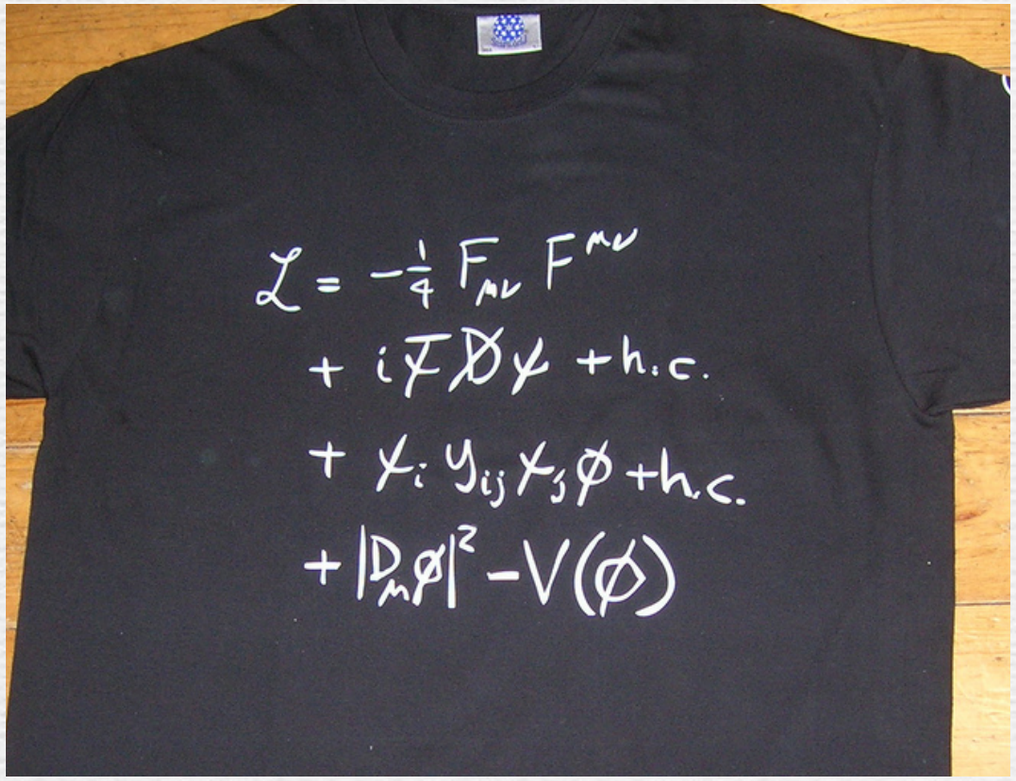

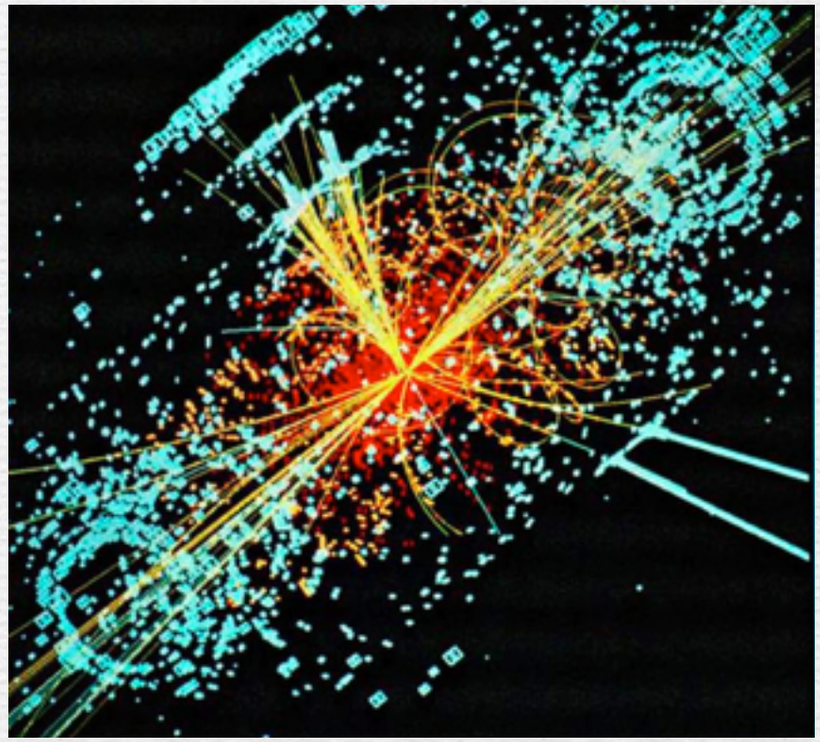

?

Theory ⟷ Experiment

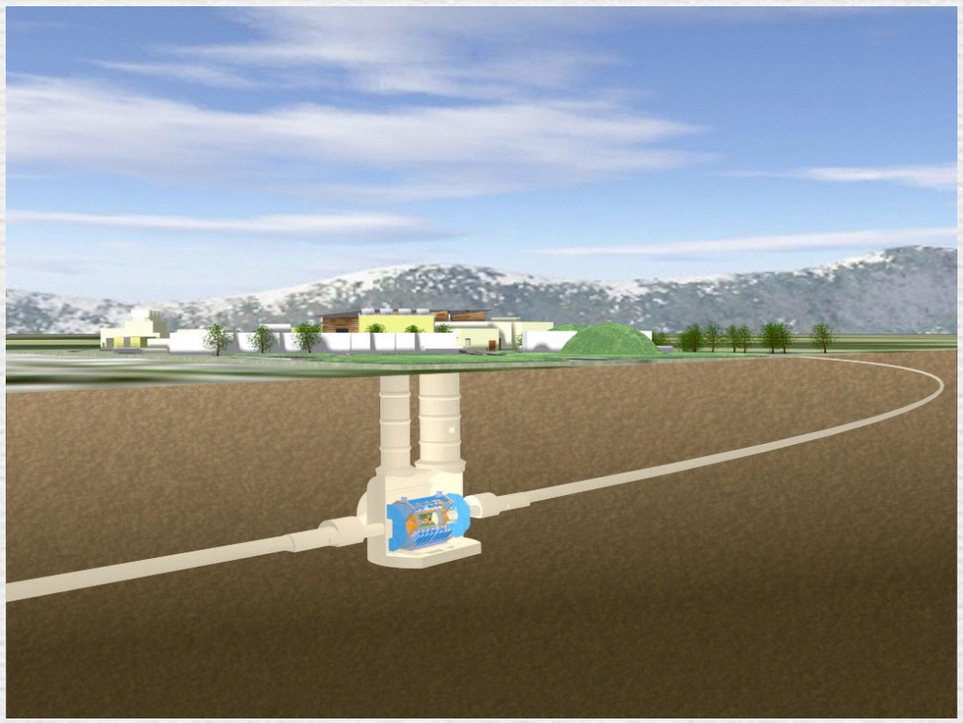

Proton bunches 25 ns apart

10^{11} protons per bunch

~25 collisions per crossing

~100 million readout channels, every 25 ns

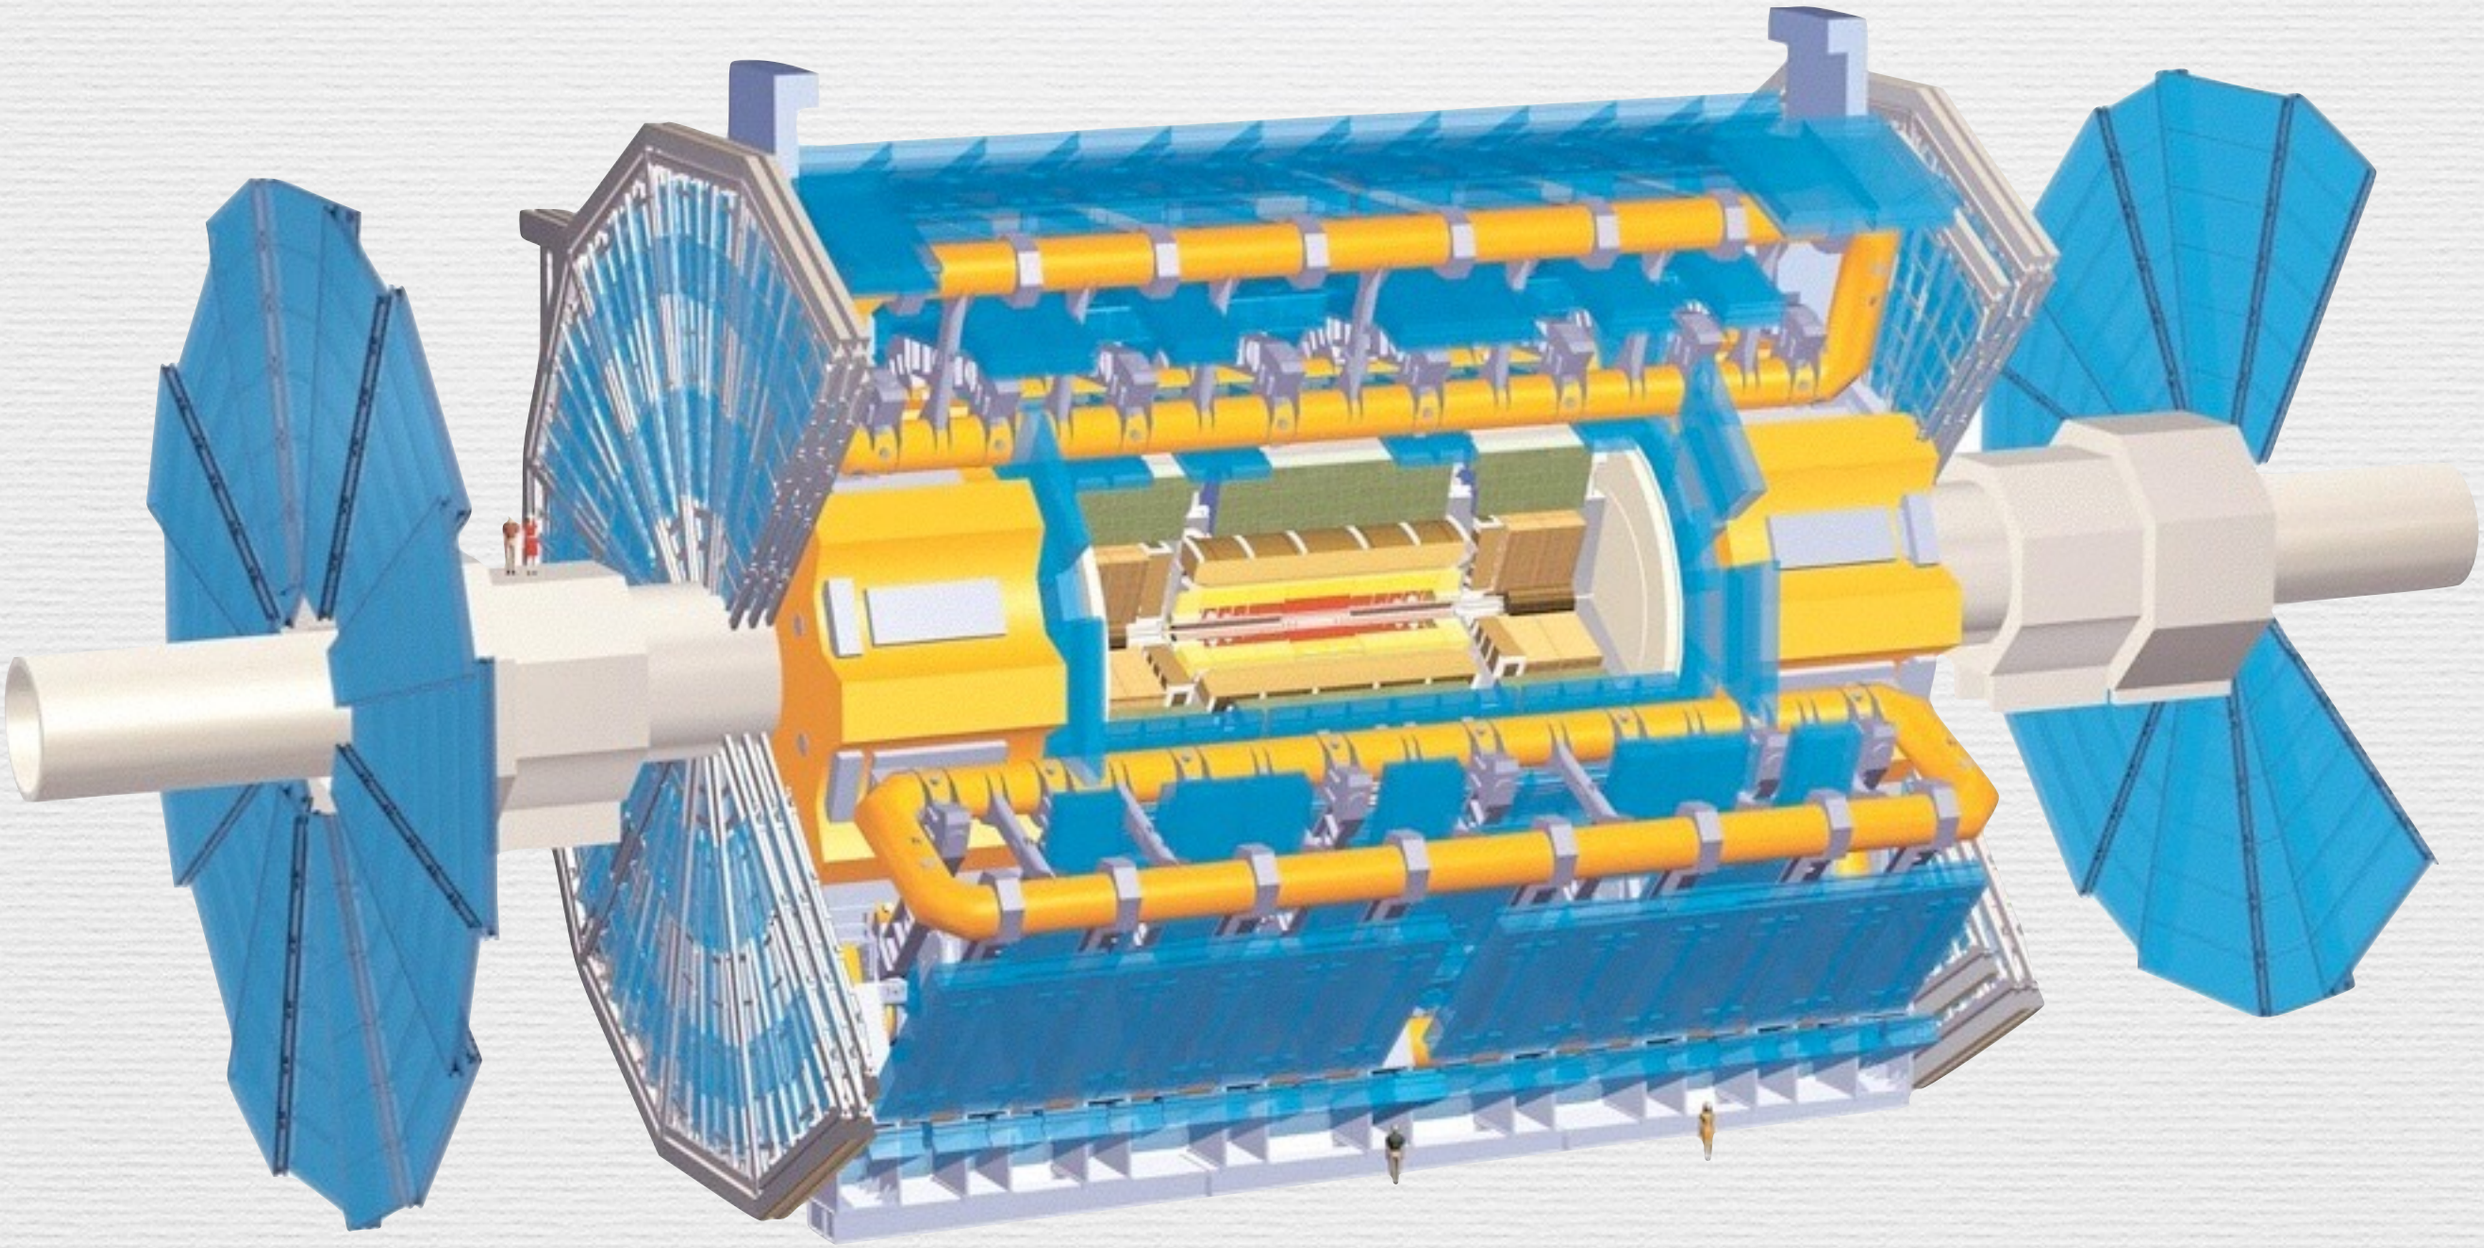

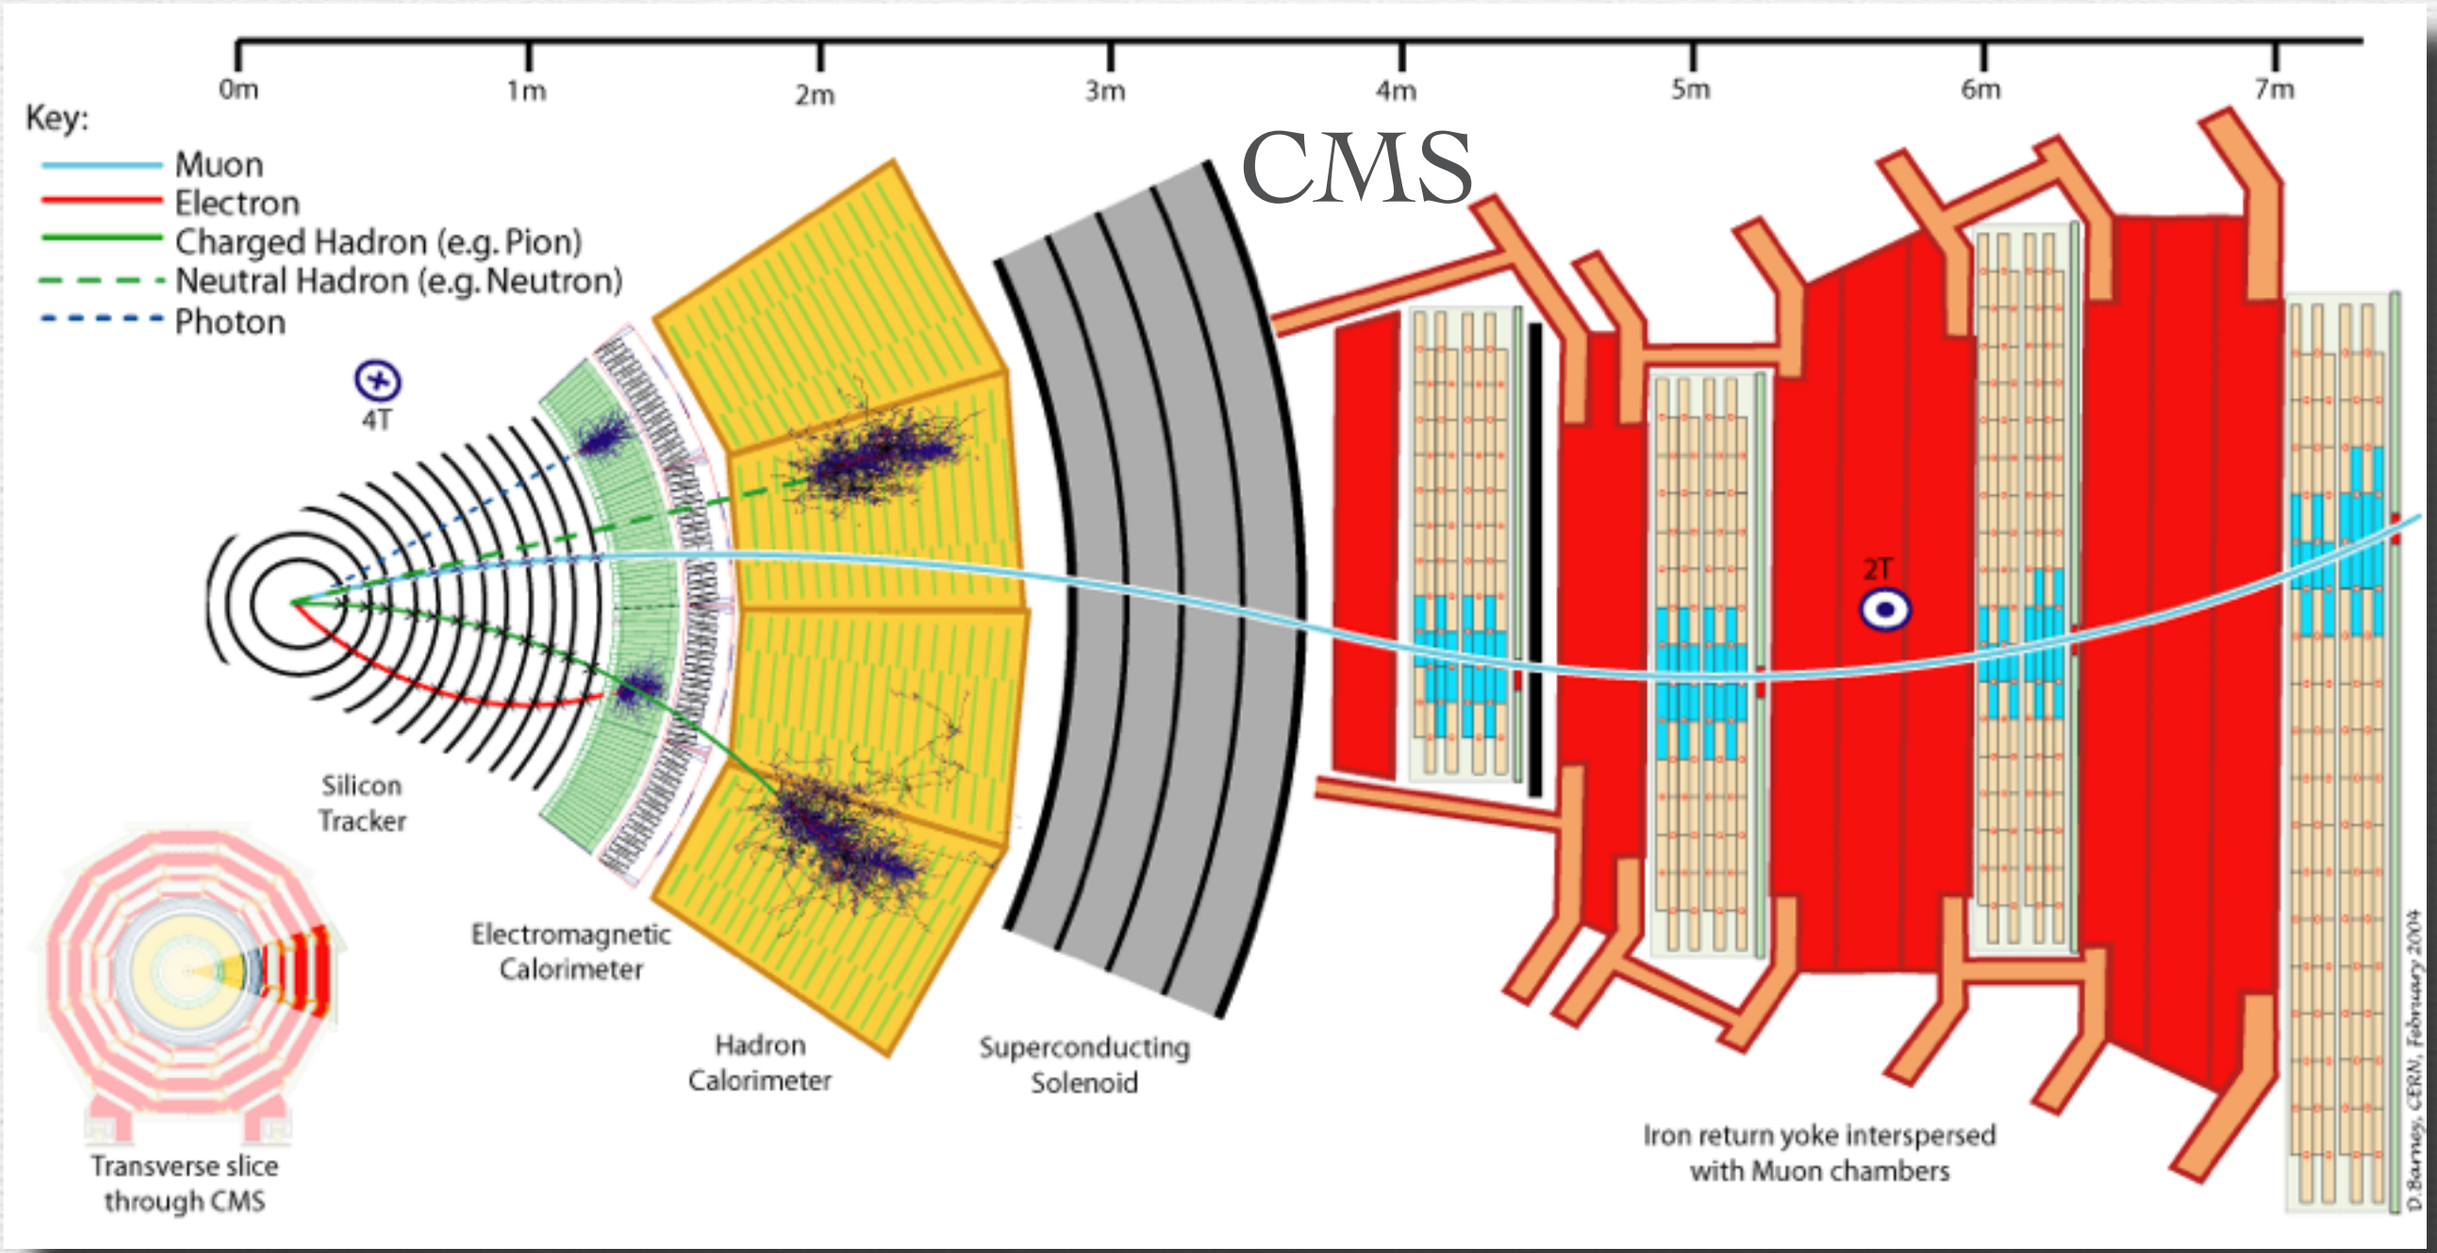
0m
1m
2m
3m
4m
5m
6m
7m
Key:
Muon
Electron
Charged Hadron (e.g. Pion)
Neutral Hadron (e.g. Neutron)
Photon
CMS
4T
2T
Silicon
Tracker
Electromagnetic
Calorimeter
Hadron
Calorimeter
Superconducting
Solenoid
Iron return yoke interspersed
with Muon chambers
Transverse slice
through CMS
D.Barney, CERN, February 2004

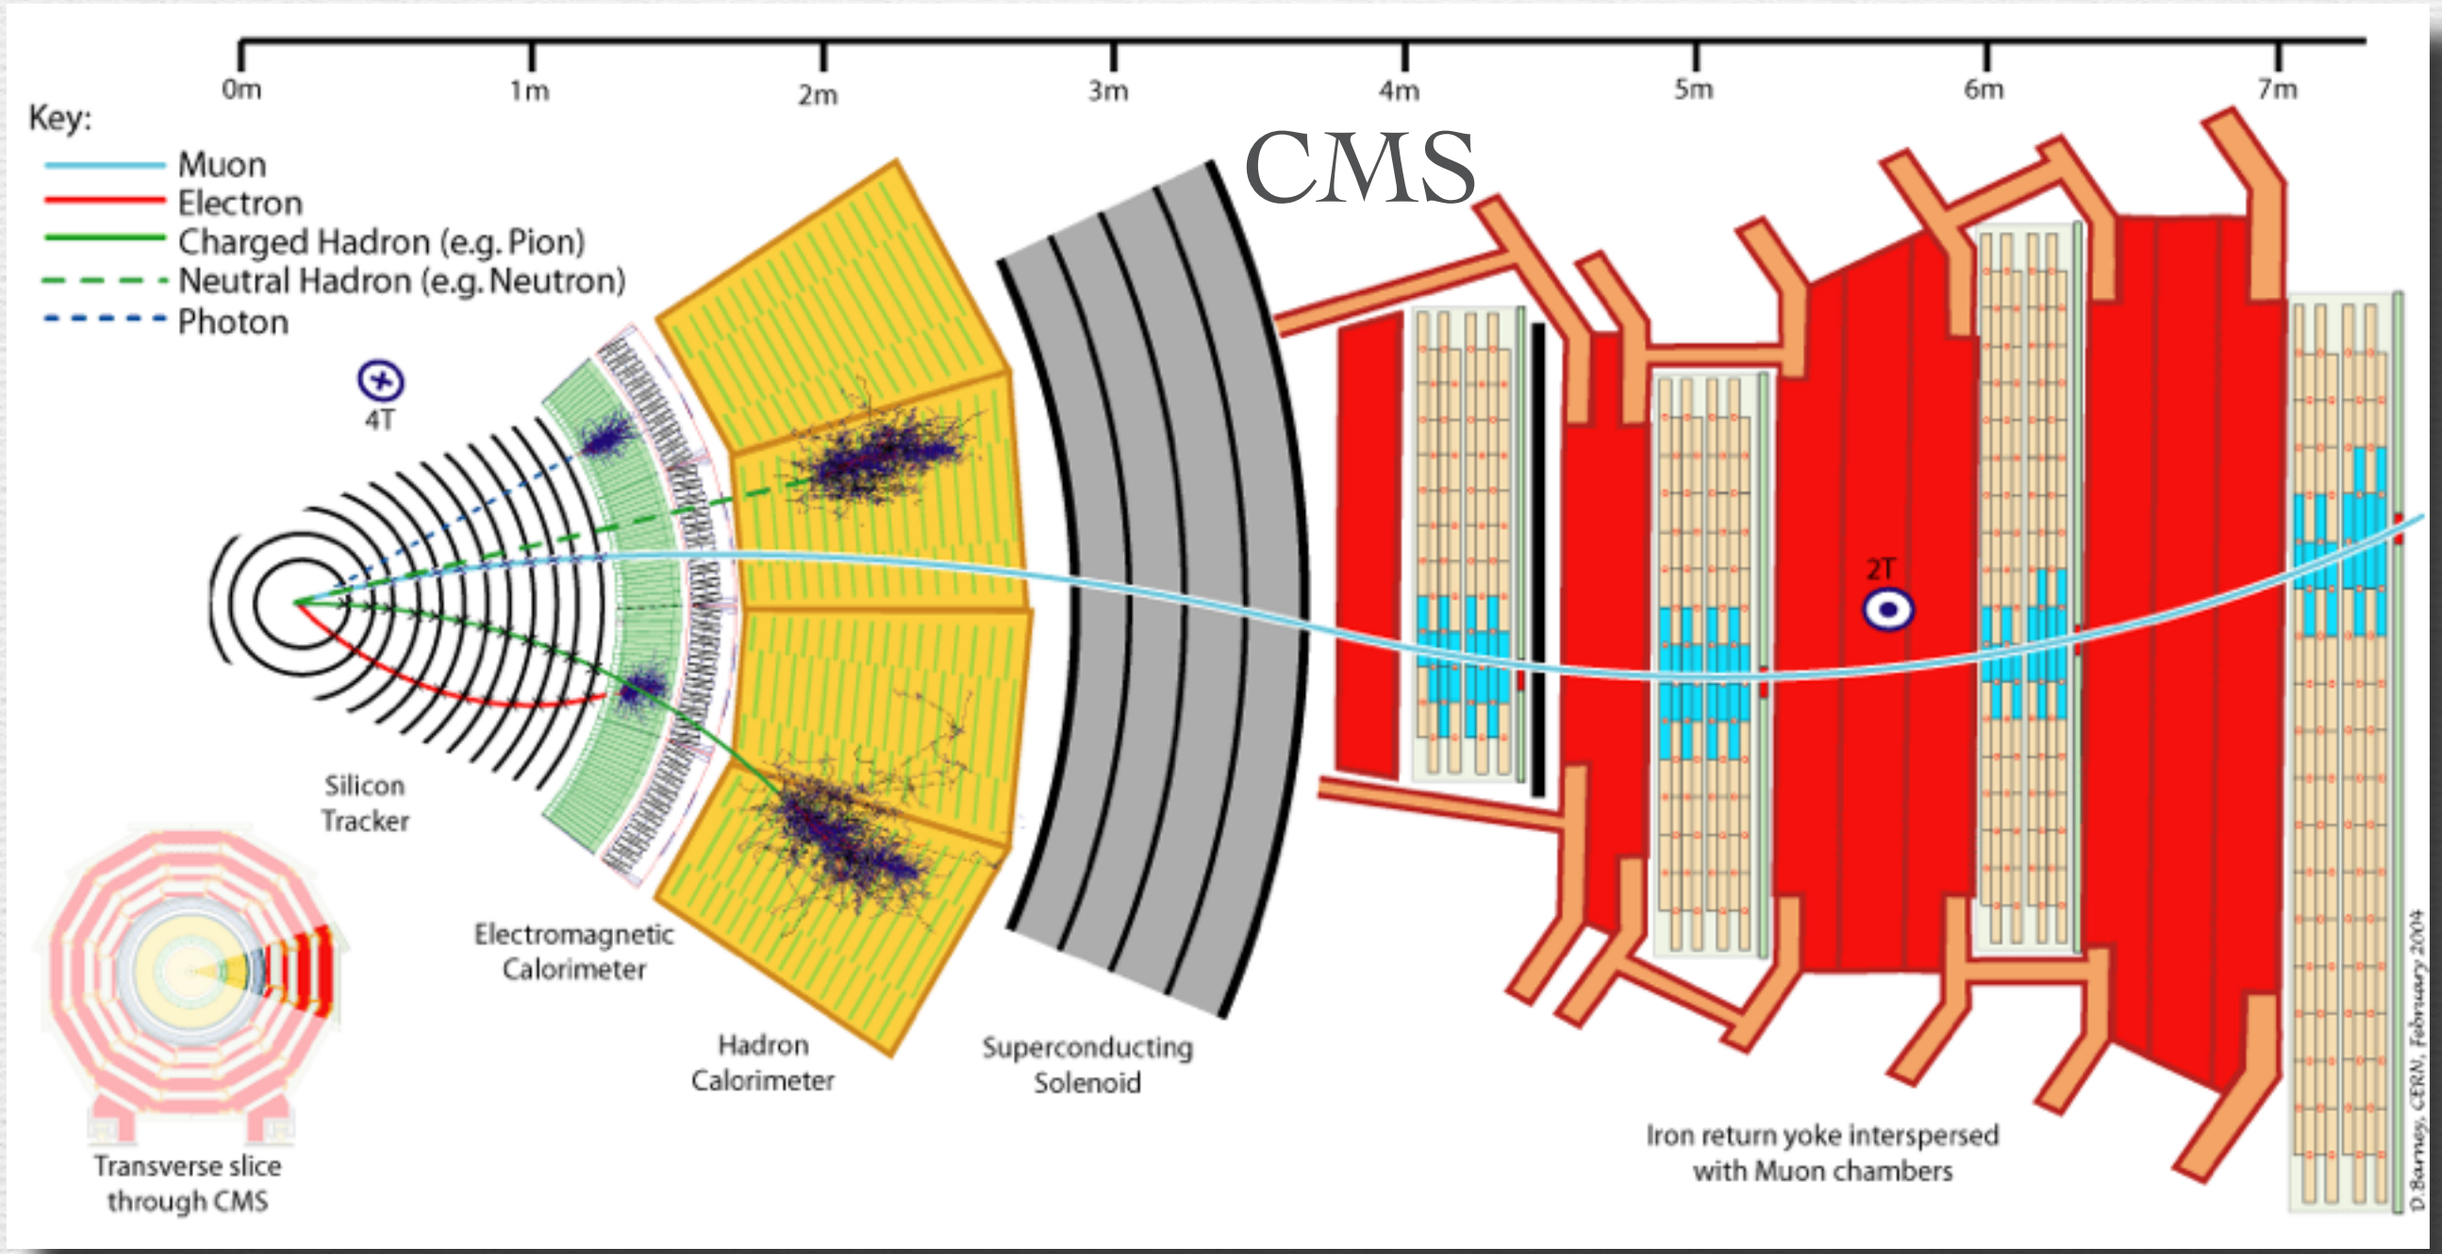

After zeroes removed, 1.6 MB / event

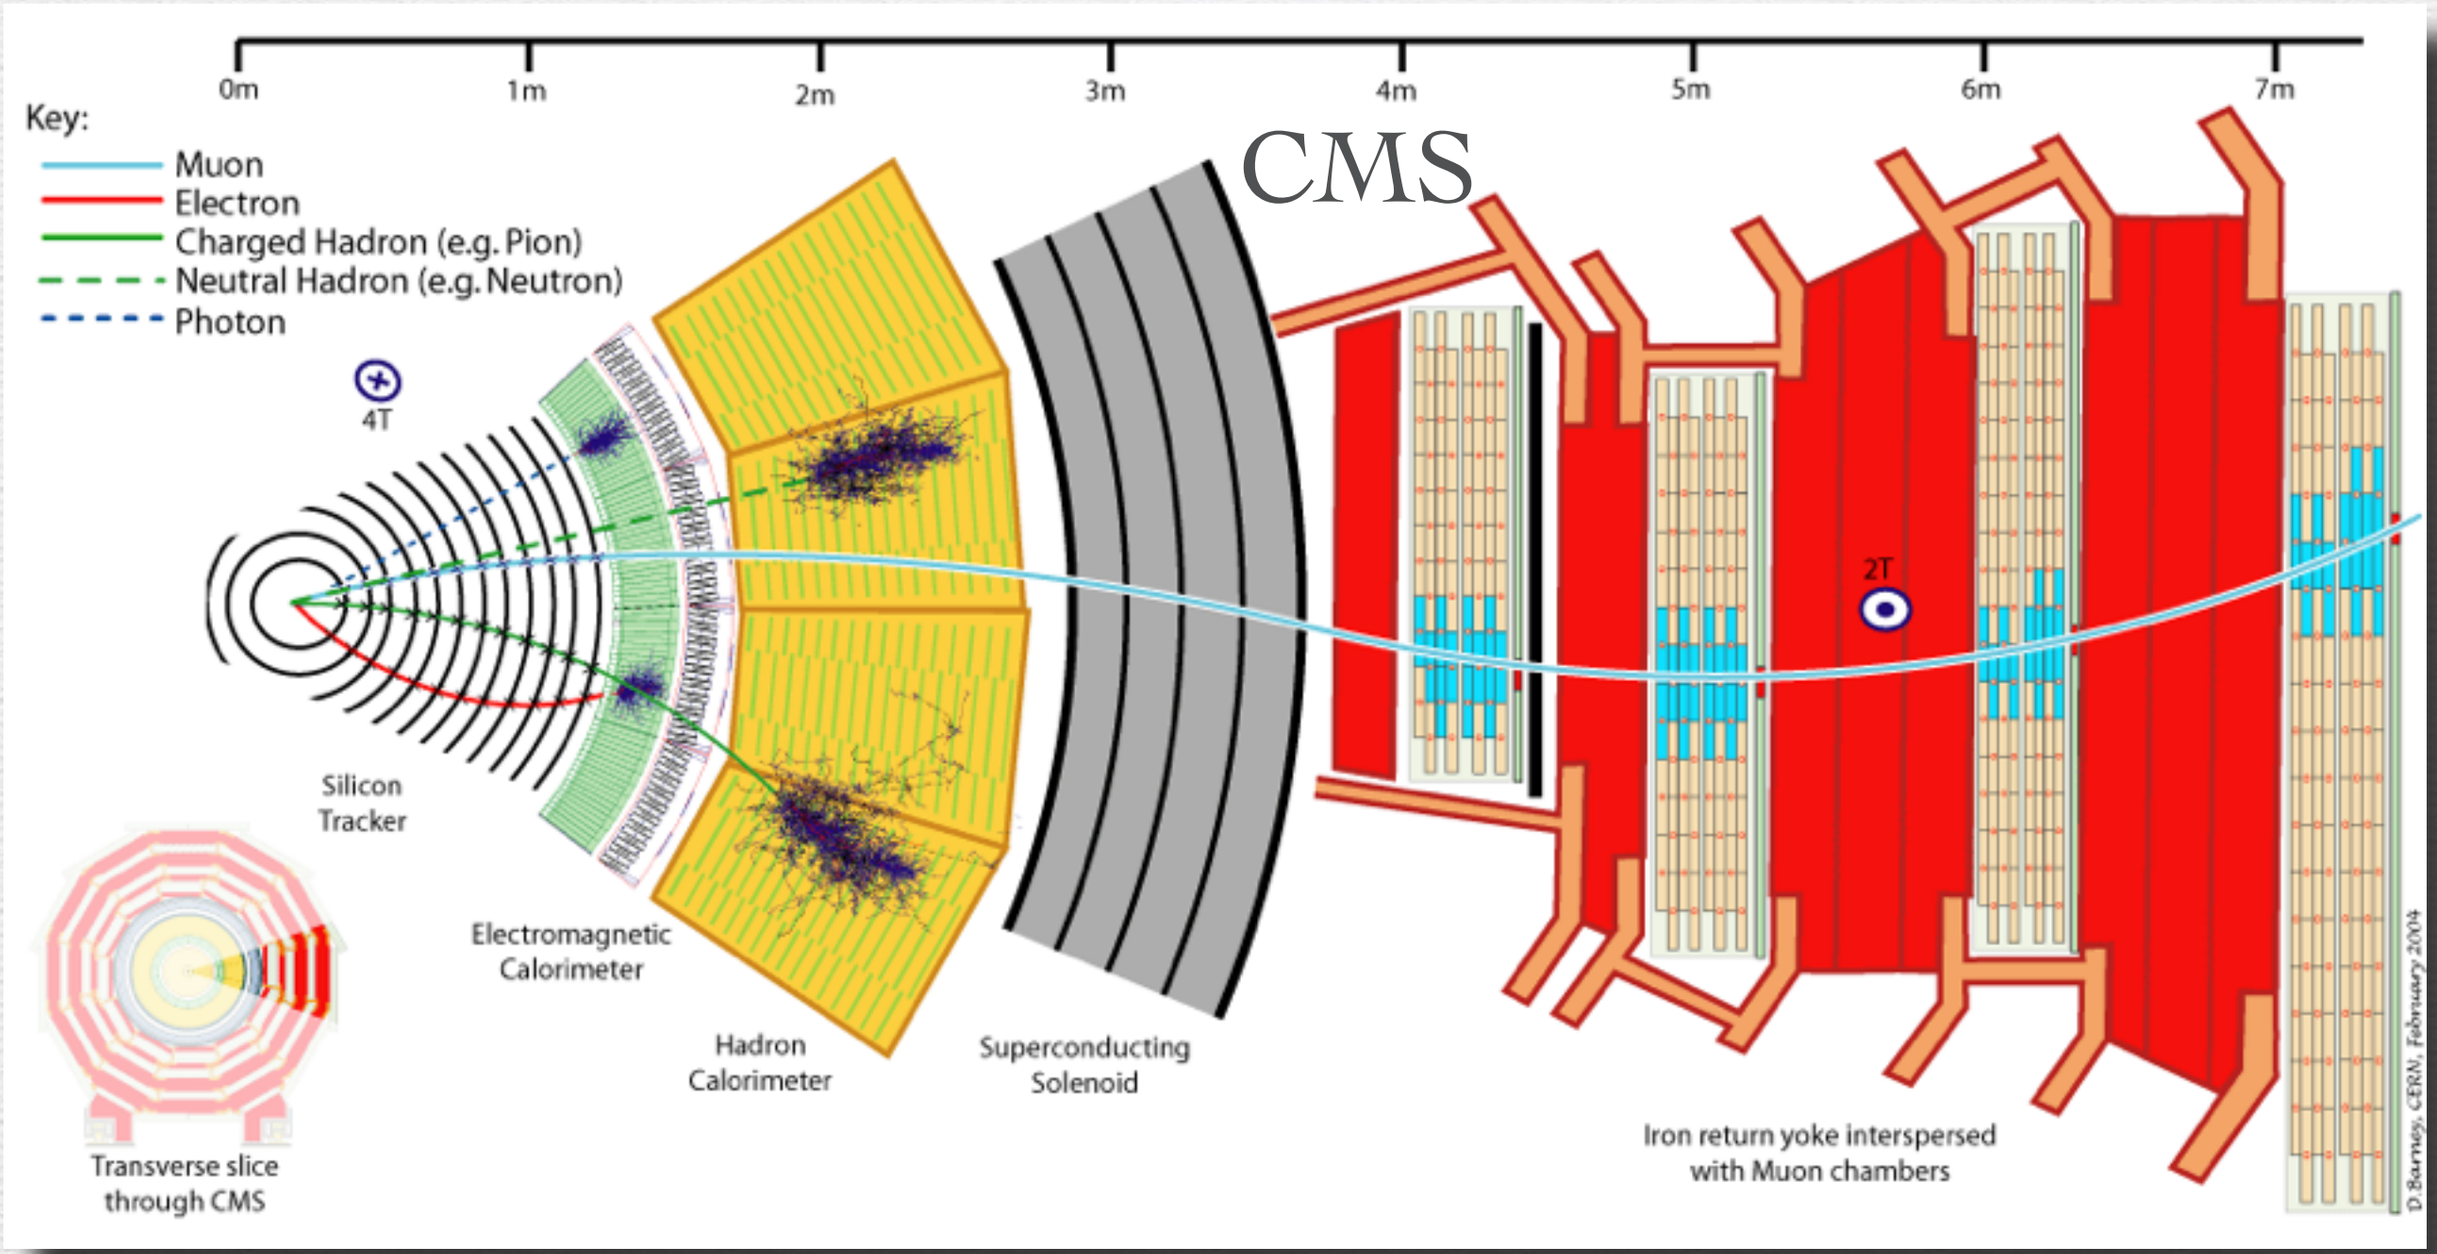

After zeroes removed, 1.6 MB / event

* 40 M events / s = 64 TB / s ?

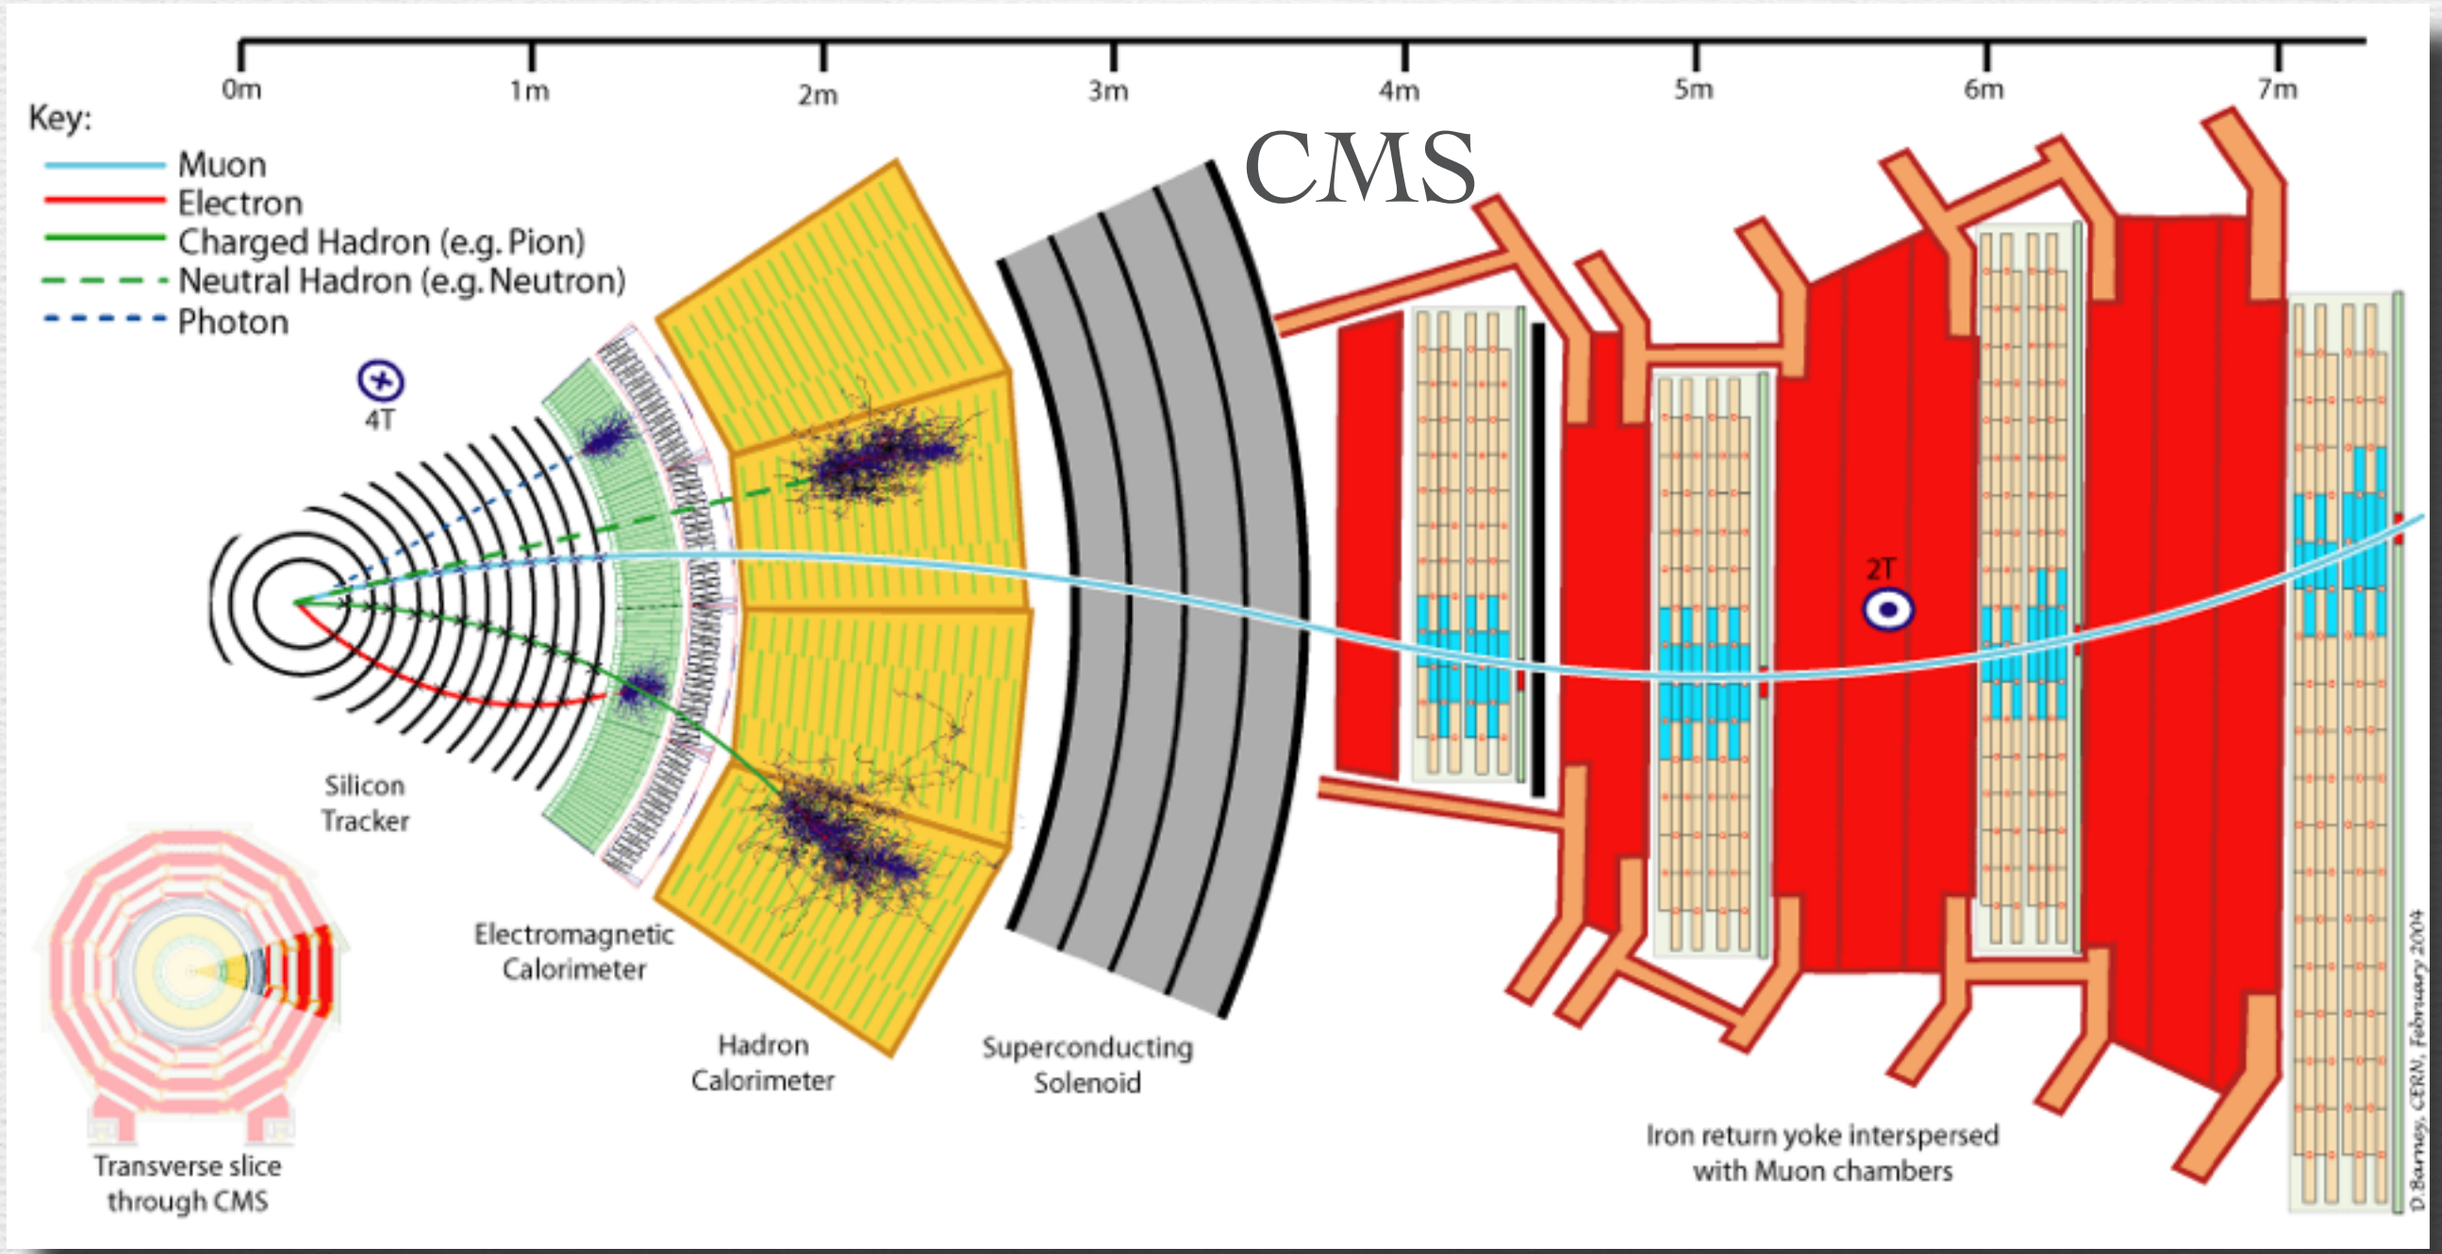

After zeroes removed, 1.6 MB / event

* 40 M events / s = 64 TB / s ?

Can't save everything

Trigger system to keep only interesting events

	Incoming event rate per second	Outgoing event rate per second	Reduction factor
Level 1	40 000 000	100 000	400
Level 2	100 000	3 000	30
Level 3	3 000	200	15

Trigger system to keep only interesting events

	Incoming event rate per second	Outgoing event rate per second	Reduction factor
Level 1	40 000 000	100 000	400
Level 2	100 000	3 000	30
Level 3	3 000	200	15

200 events / s * 1.6 MB / event = 320 MB / s

Trigger system to keep only interesting events

	Incoming event rate per second	Outgoing event rate per second	Reduction factor
Level 1	40 000 000	100 000	400
Level 2	100 000	3 000	30
Level 3	3 000	200	15

200 events / s * 1.6 MB / event = 320 MB / s

= ~ 3200 TB / year raw data

Trigger system to keep only interesting events

	Incoming event rate per second	Outgoing event rate per second	Reduction factor
Level 1	40 000 000	100 000	400
Level 2	100 000	3 000	30
Level 3	3 000	200	15

200 events / s * 1.6 MB / event = 320 MB / s

= ~ 3200 TB / year raw data

Analysis is done offline,
~3000 collaboration members should have
equal access to data worldwide

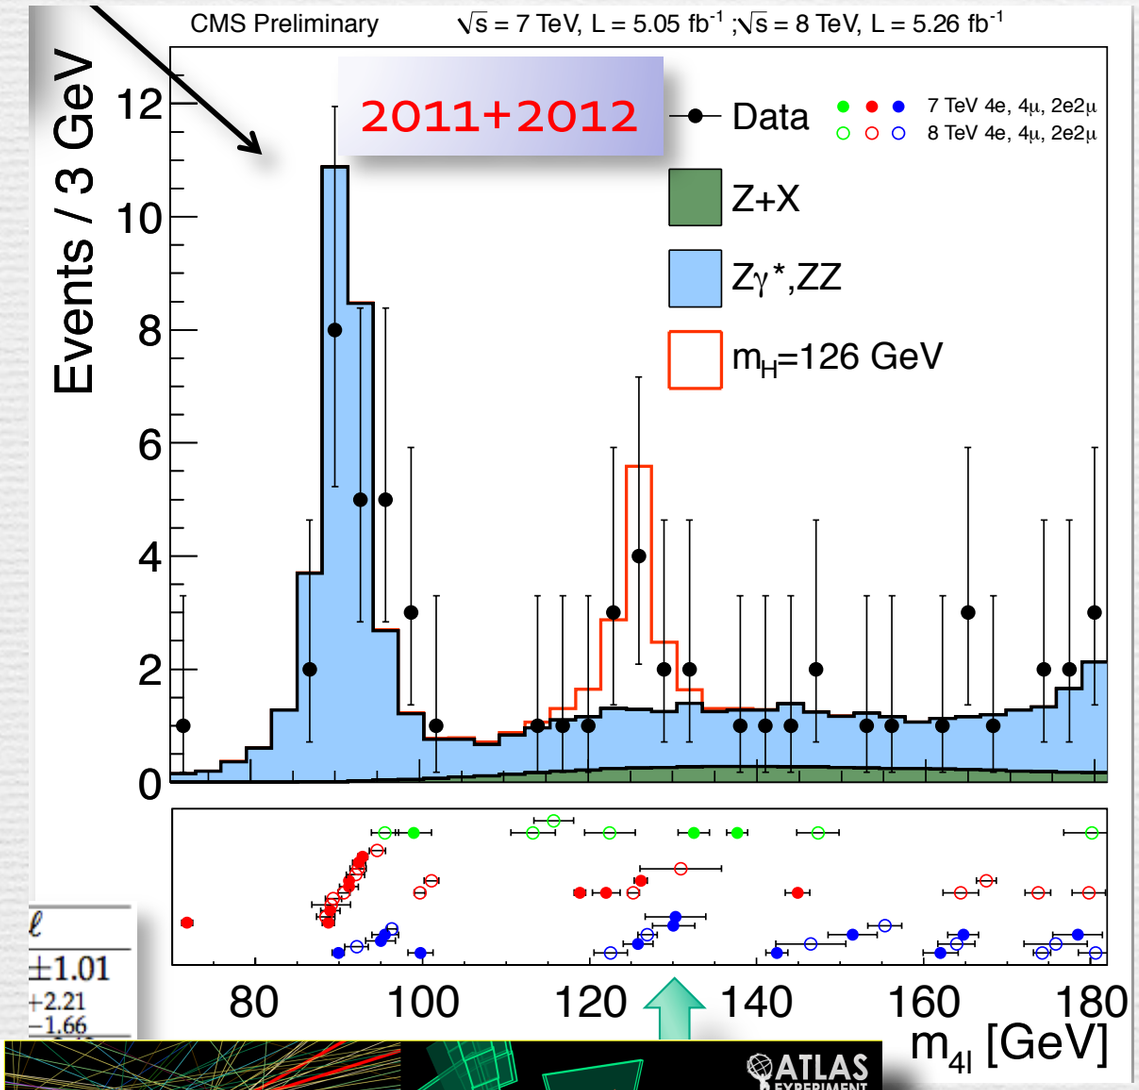
CMS Preliminary
√s = 7 TeV, L = 5.05 fb-1 ;√s = 8 TeV, L = 5.26 fb-1
2011+2012
Data
7 TeV 4e, 4μ, 2e2μ
8 TeV 4e, 4μ, 2e2μ
Z+X
Zγ*,ZZ
mH=126 GeV
Events / 3 GeV
m4l [GeV]

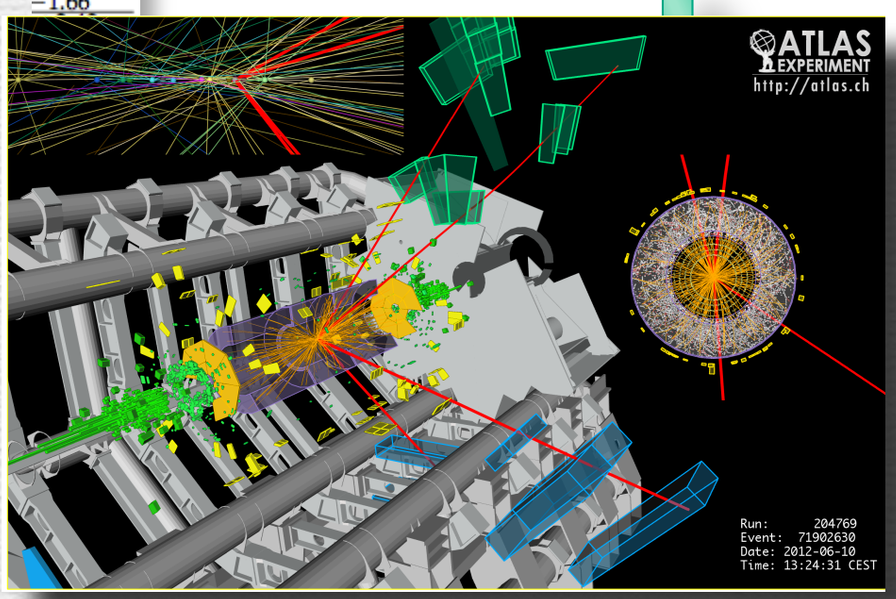
ATLAS
EXPERIMENT
http://atlas.ch
Run: 204769
Event: 71902630
Date: 2012-06-10
Time: 13:24:31 CEST

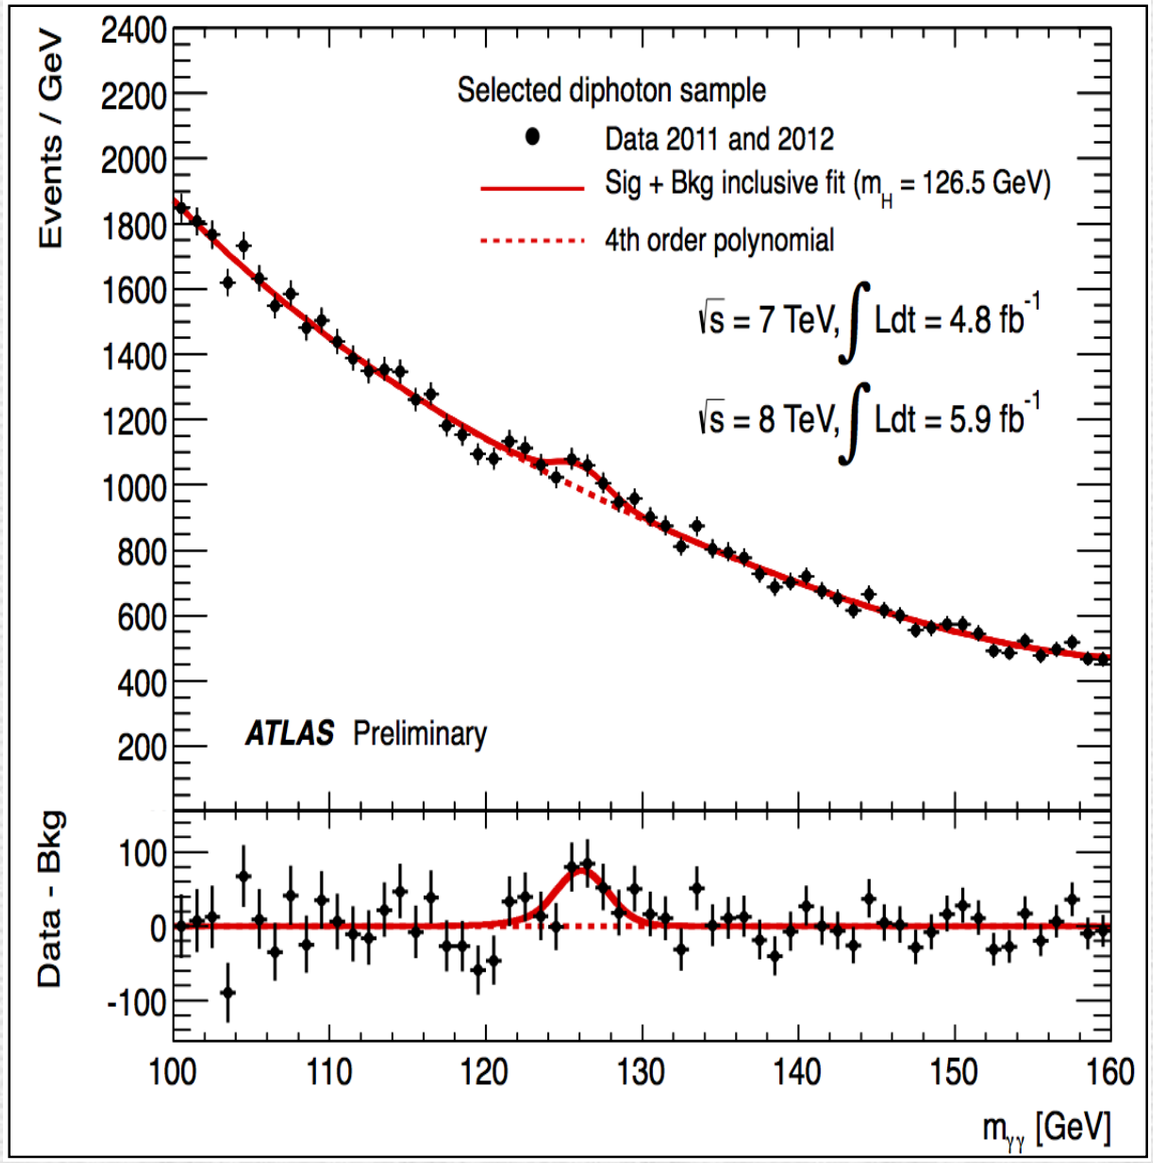
Selected diphoton sample
Data 2011 and 2012
Sig + Bkg inclusive fit (mH = 126.5 GeV)
4th order polynomial
√s = 7 TeV, ∫ Ldt = 4.8 fb-1
√s = 8 TeV, ∫ Ldt = 5.9 fb-1
ATLAS Preliminary
Events / GeV
Data - Bkg
mγγ [GeV]

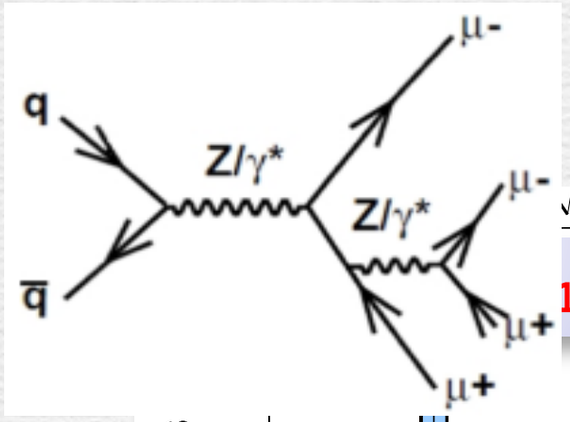
q
q̄
Z/γ*
Z/γ*
μ-
μ-
μ+
μ+

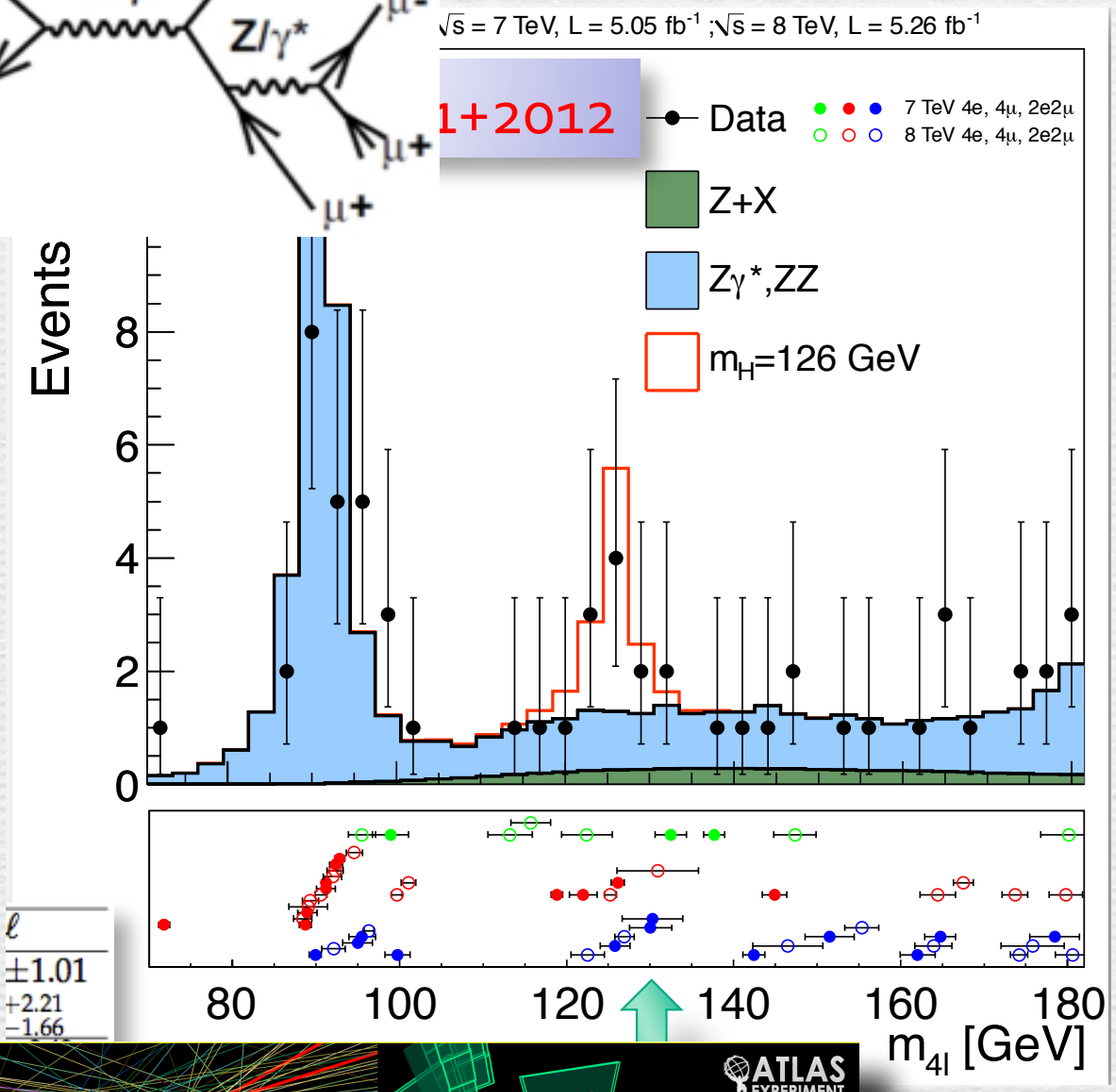
√s = 7 TeV, L = 5.05 fb-1 ;√s = 8 TeV, L = 5.26 fb-1
1+2012
Data
7 TeV 4e, 4μ, 2e2μ
8 TeV 4e, 4μ, 2e2μ
Z+X
Zγ*,ZZ
mH=126 GeV
Events
0
2
4
6
8
80
100
120
140
160
180
m4l [GeV]

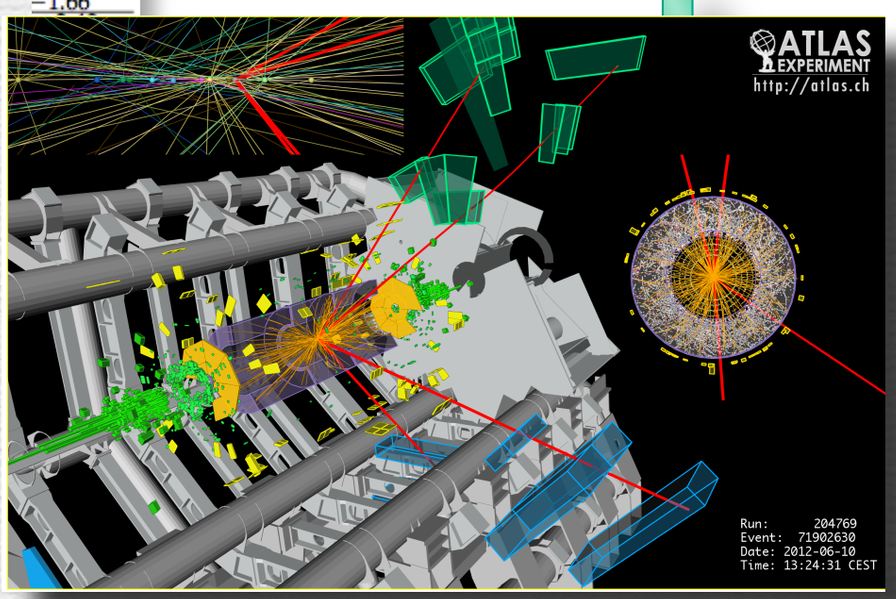
ATLAS
EXPERIMENT
http://atlas.ch
Run: 204769
Event: 71902630
Date: 2012-06-10
Time: 13:24:31 CEST

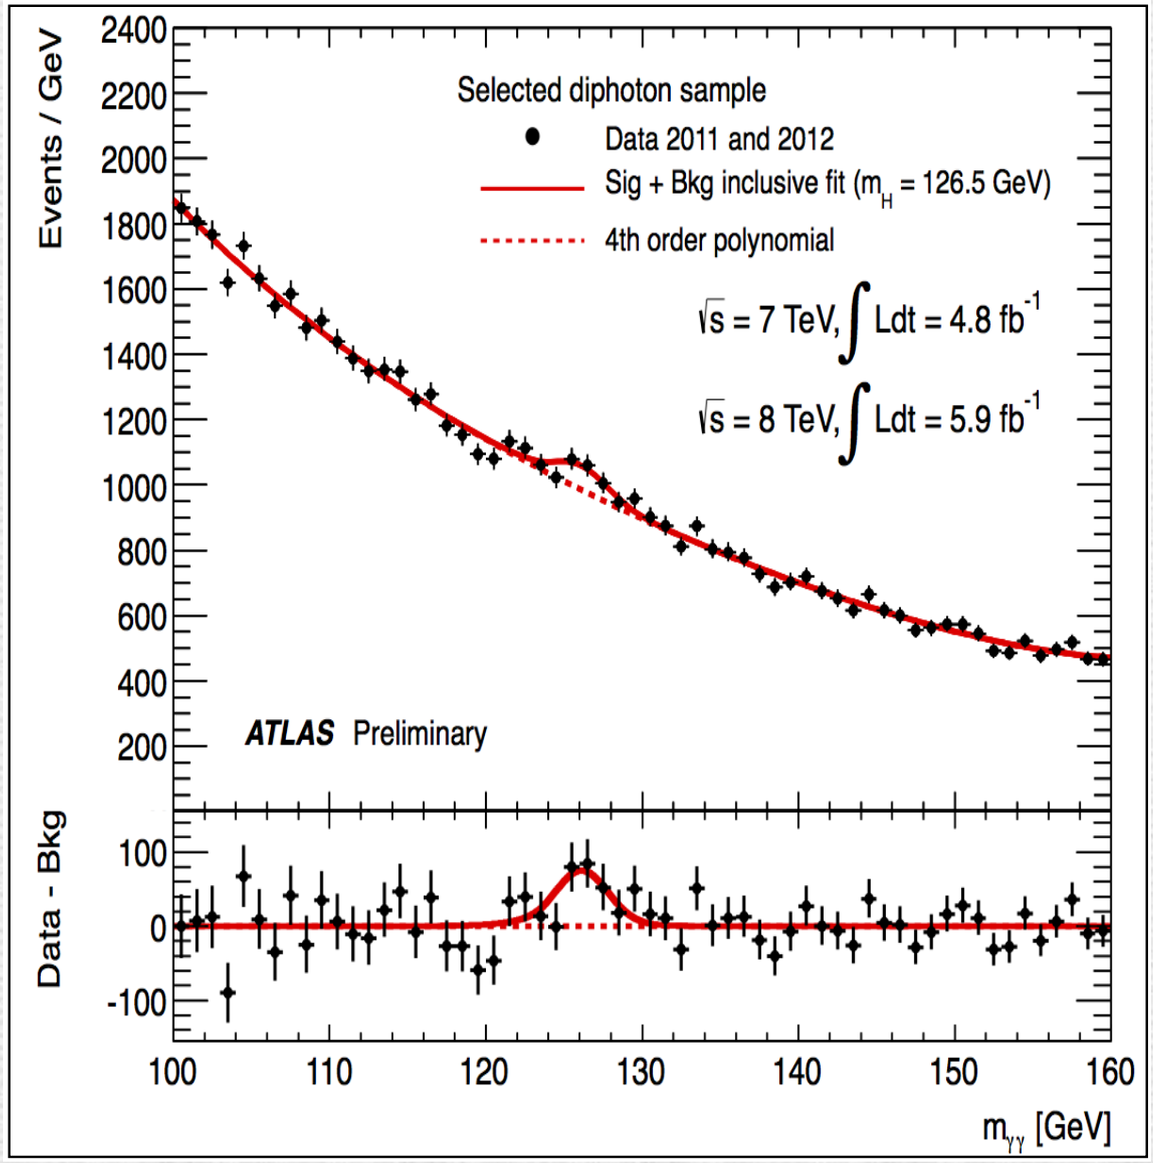
Selected diphoton sample
Data 2011 and 2012
Sig + Bkg inclusive fit (mH = 126.5 GeV)
4th order polynomial
√s = 7 TeV, ∫ Ldt = 4.8 fb-1
√s = 8 TeV, ∫ Ldt = 5.9 fb-1
ATLAS Preliminary
Events / GeV
Data - Bkg
mγγ [GeV]

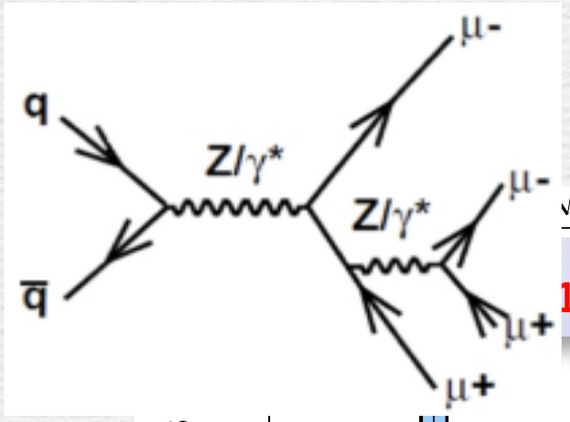

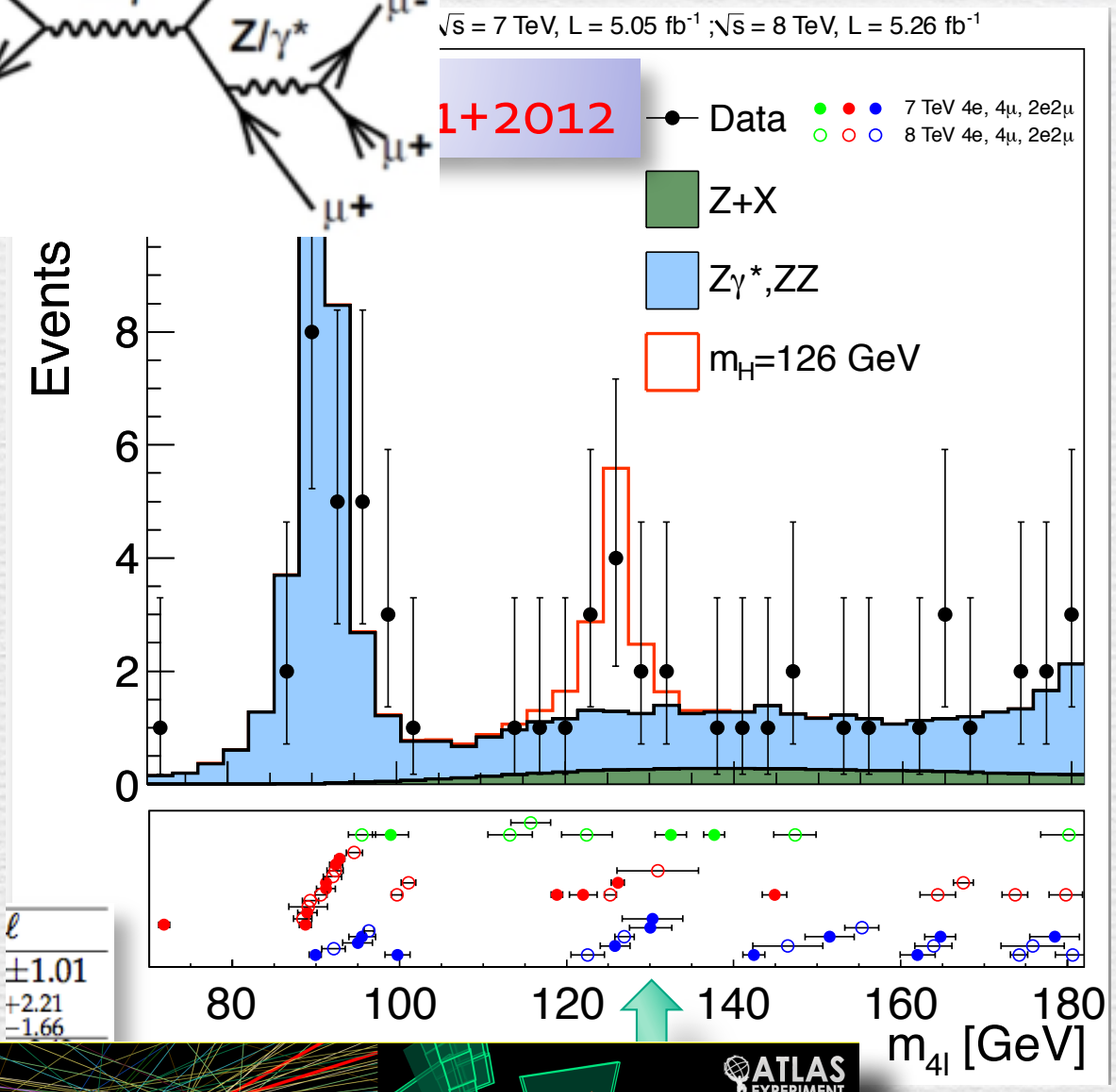

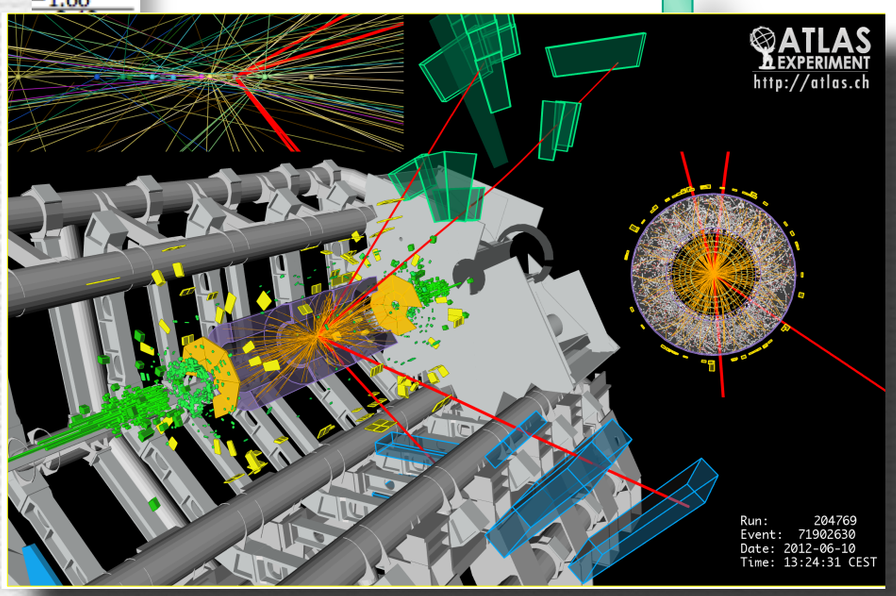

Need to get theory predictions.

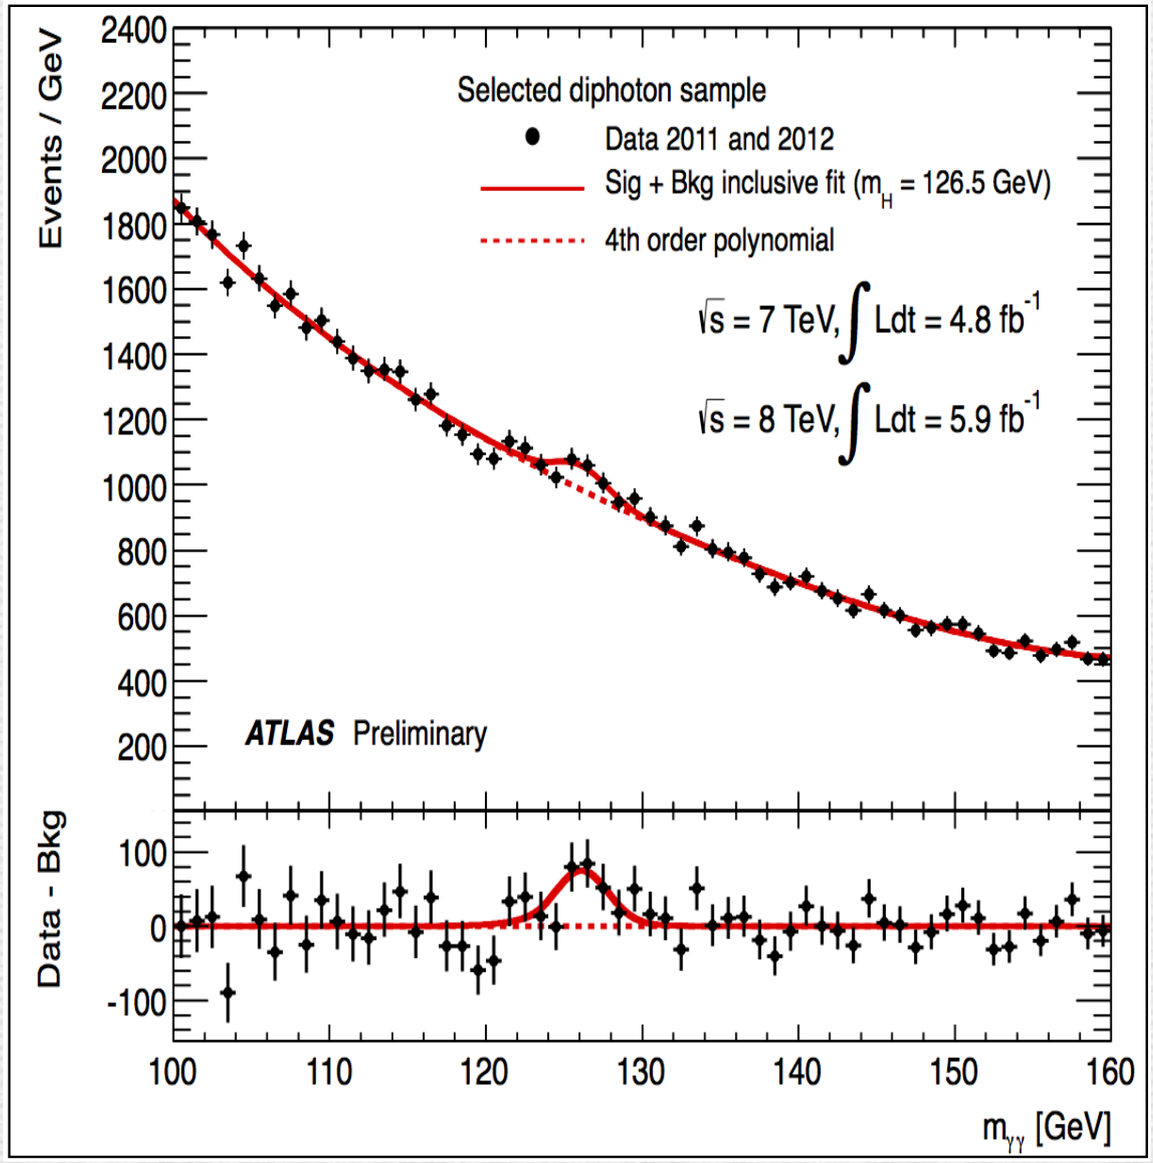

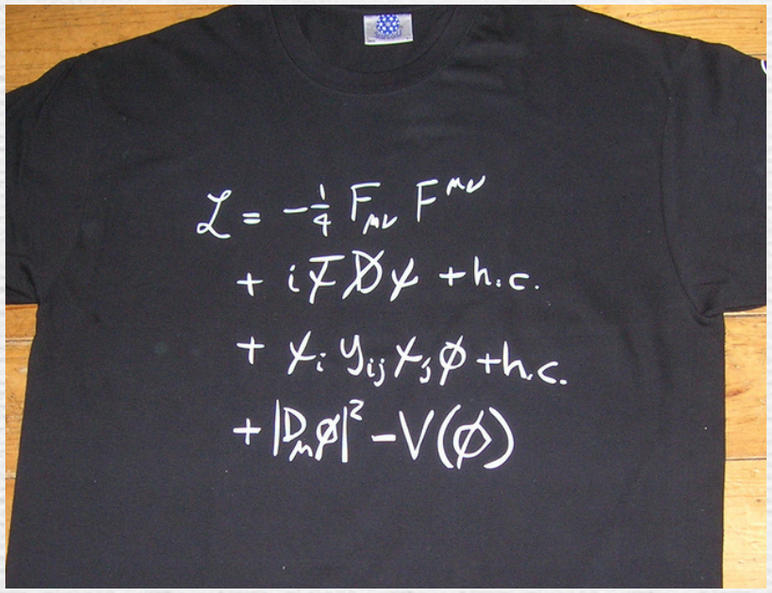
$\mathcal{L} = -\frac{1}{4} F_{\mu\nu} F^{\mu\nu}$
$+ i\bar{\mathcal{F}} \not{D} \mathcal{F} + h.c.$
$+ \mathcal{F}_i y_{ij} \mathcal{F}_j \phi + h.c.$
$+ |D_\mu \phi|^2 - V(\phi)$

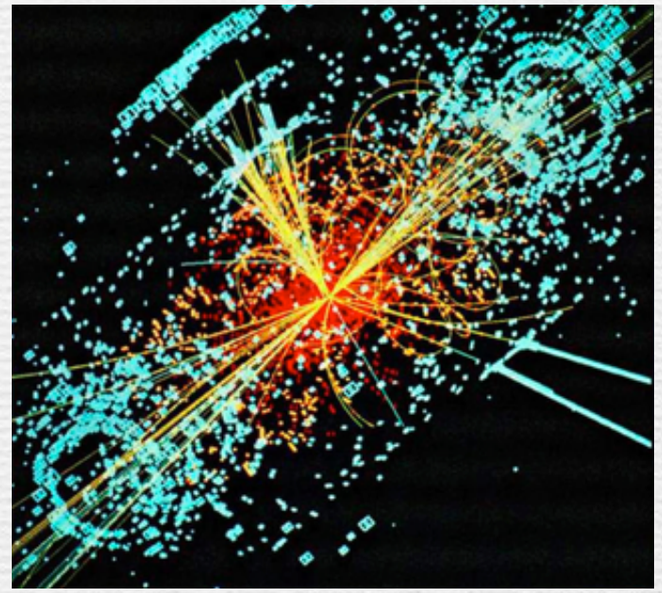

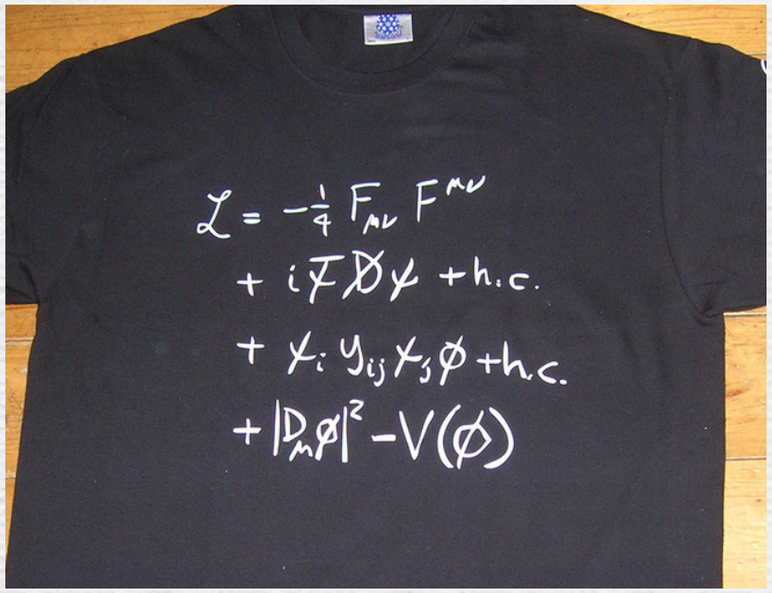
$\mathcal{L} = -\frac{1}{4} F_{\mu\nu} F^{\mu\nu}$
$+ i\bar{\mathcal{F}} \not{D} \mathcal{F} + h.c.$
$+ \mathcal{F}_i y_{ij} \mathcal{F}_j \phi + h.c.$
$+ |D_\mu \phi|^2 - V(\phi)$

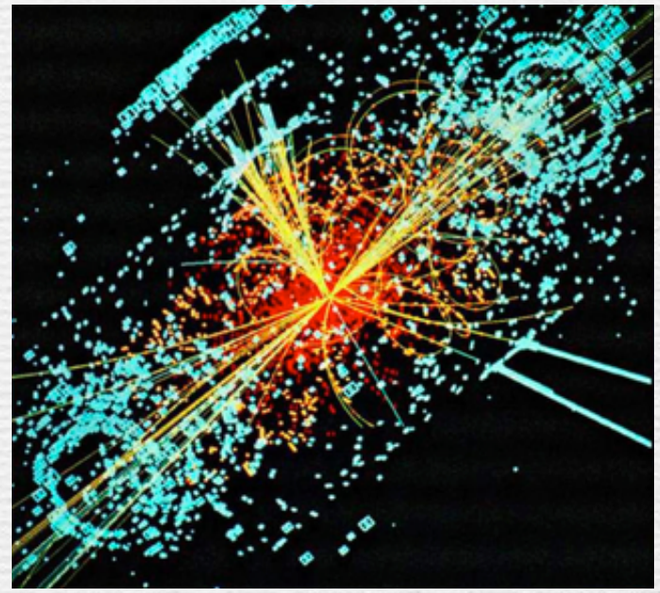

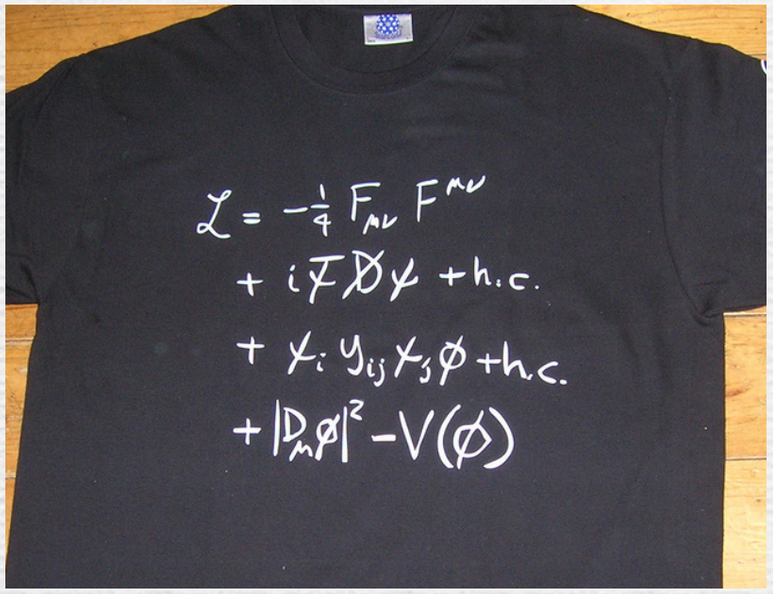

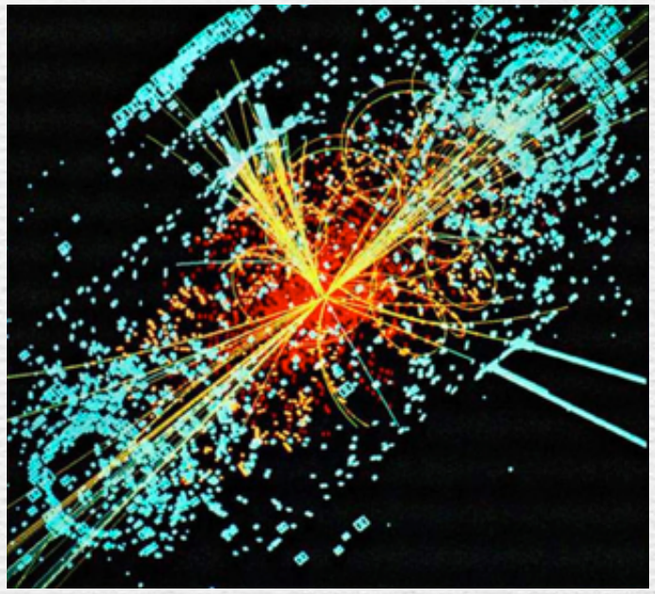

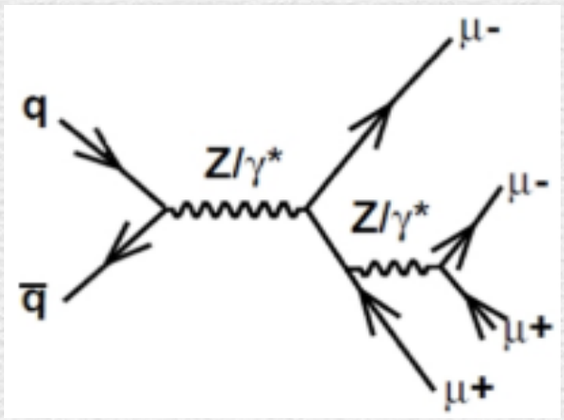

Monte Carlo event generators

Nature

Theory model

Detector

Event generator

Trigger

Simulated events

Reconstructed events

Analysis

Event generator

Matrix element

Parton shower

Hadronization

Decays

Event generator

Matrix element

Monte-Carlo integration

Parton shower

Markov chain

Hadronization

book-keeping

Decays

Monte-Carlo integration

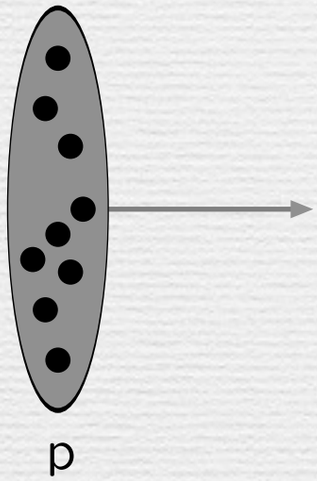

p

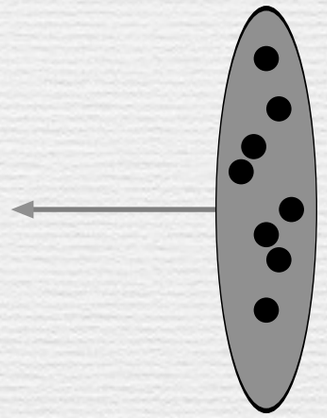

$p/\overline{p}$

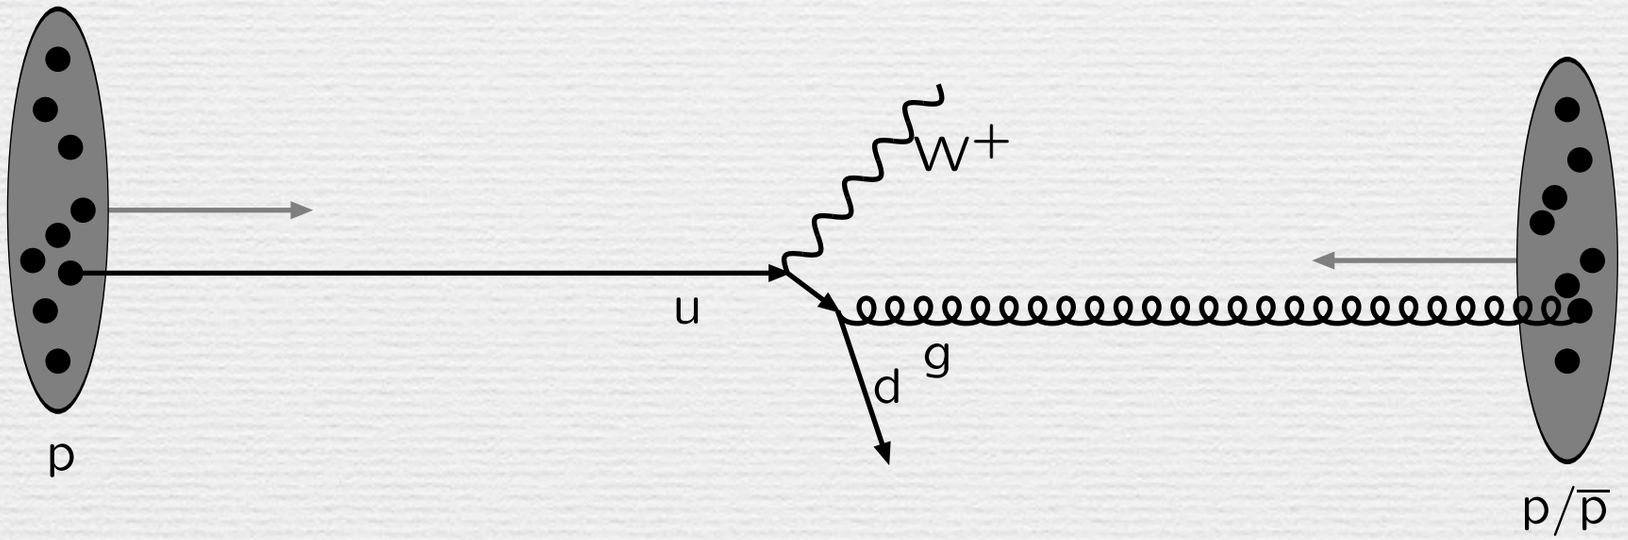
W+
u
g
d
p
p/p̄

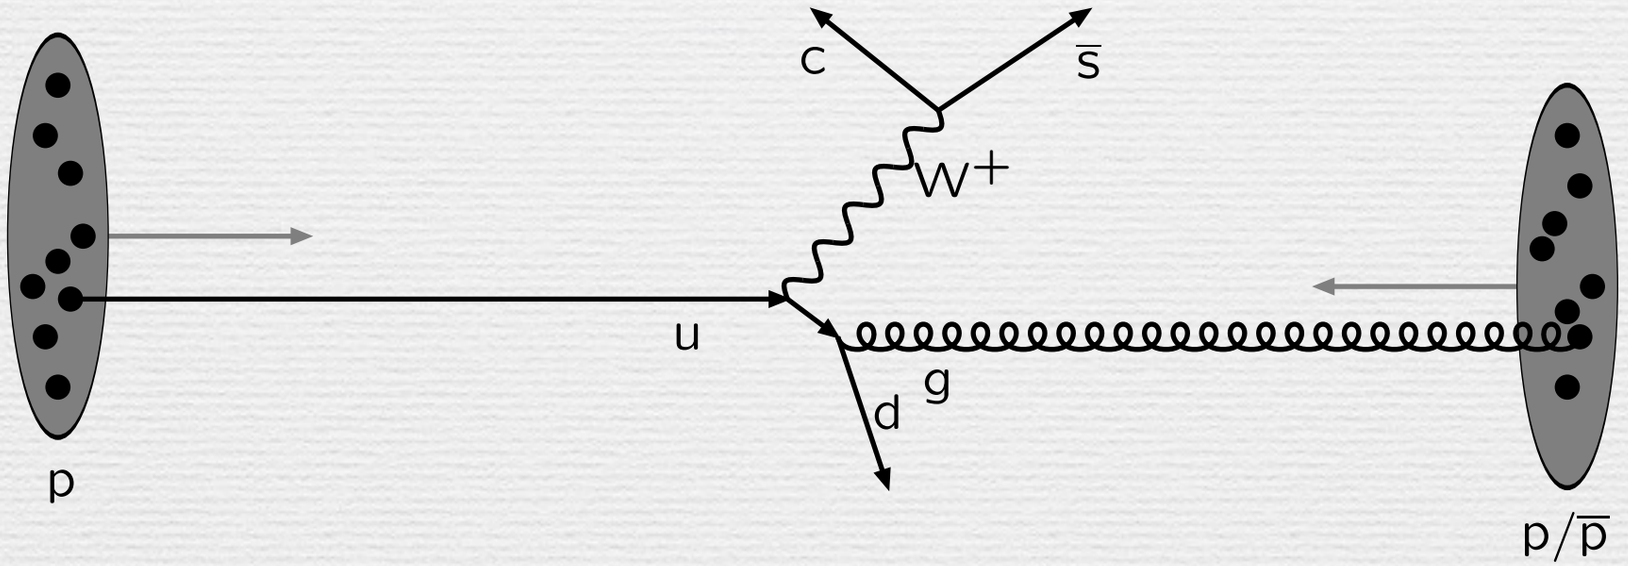
c
$\bar{s}$
W+
u
g
d
p
p/$\bar{p}$

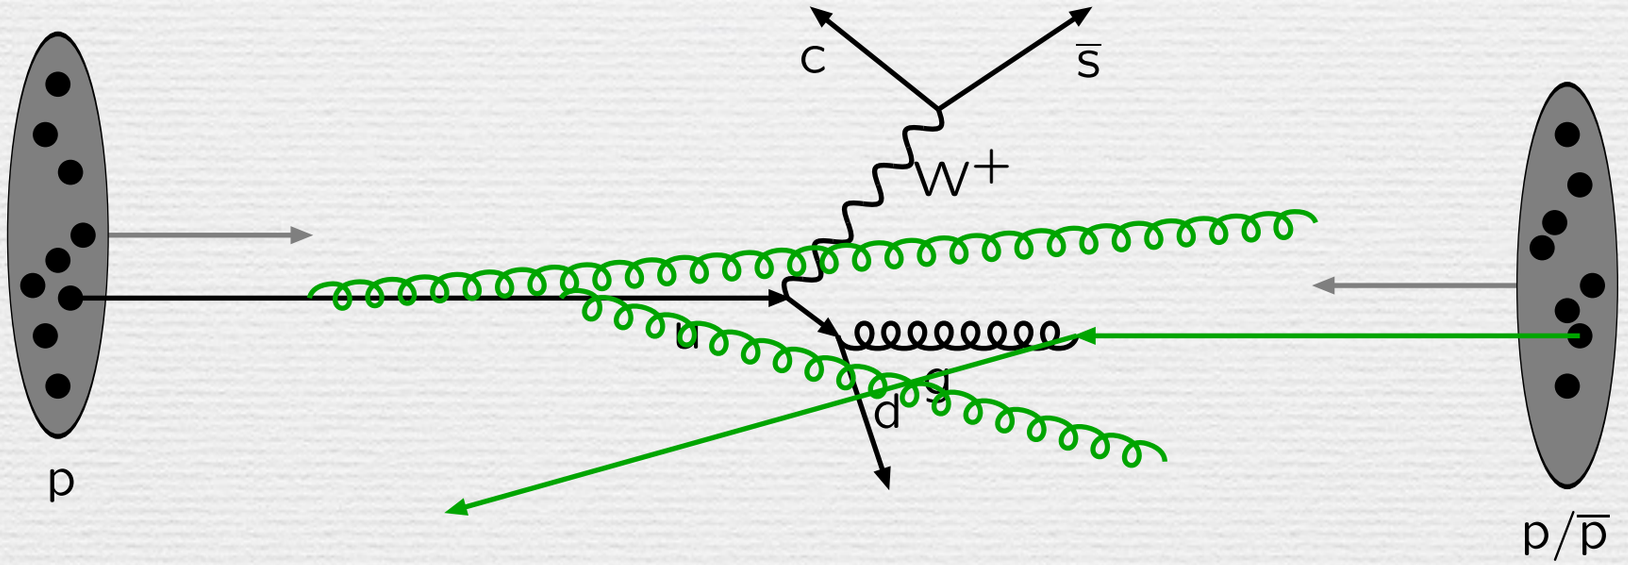
c
s̄
W+
u
d̄
d
p
p/p̄

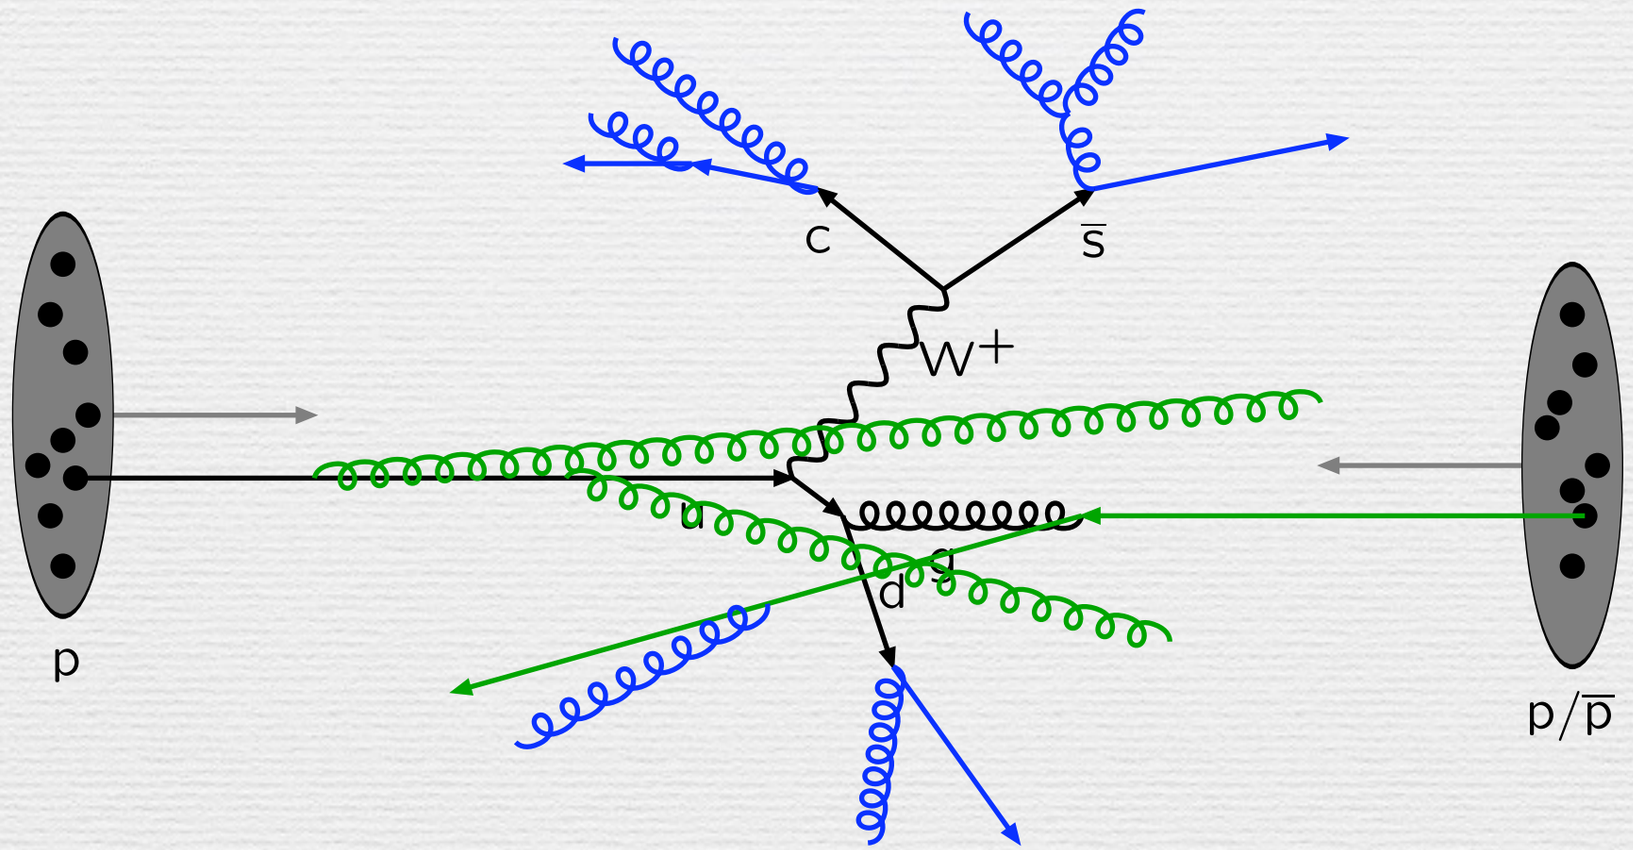
c
s̄
W+
u
d̄
d
p
p/p̄

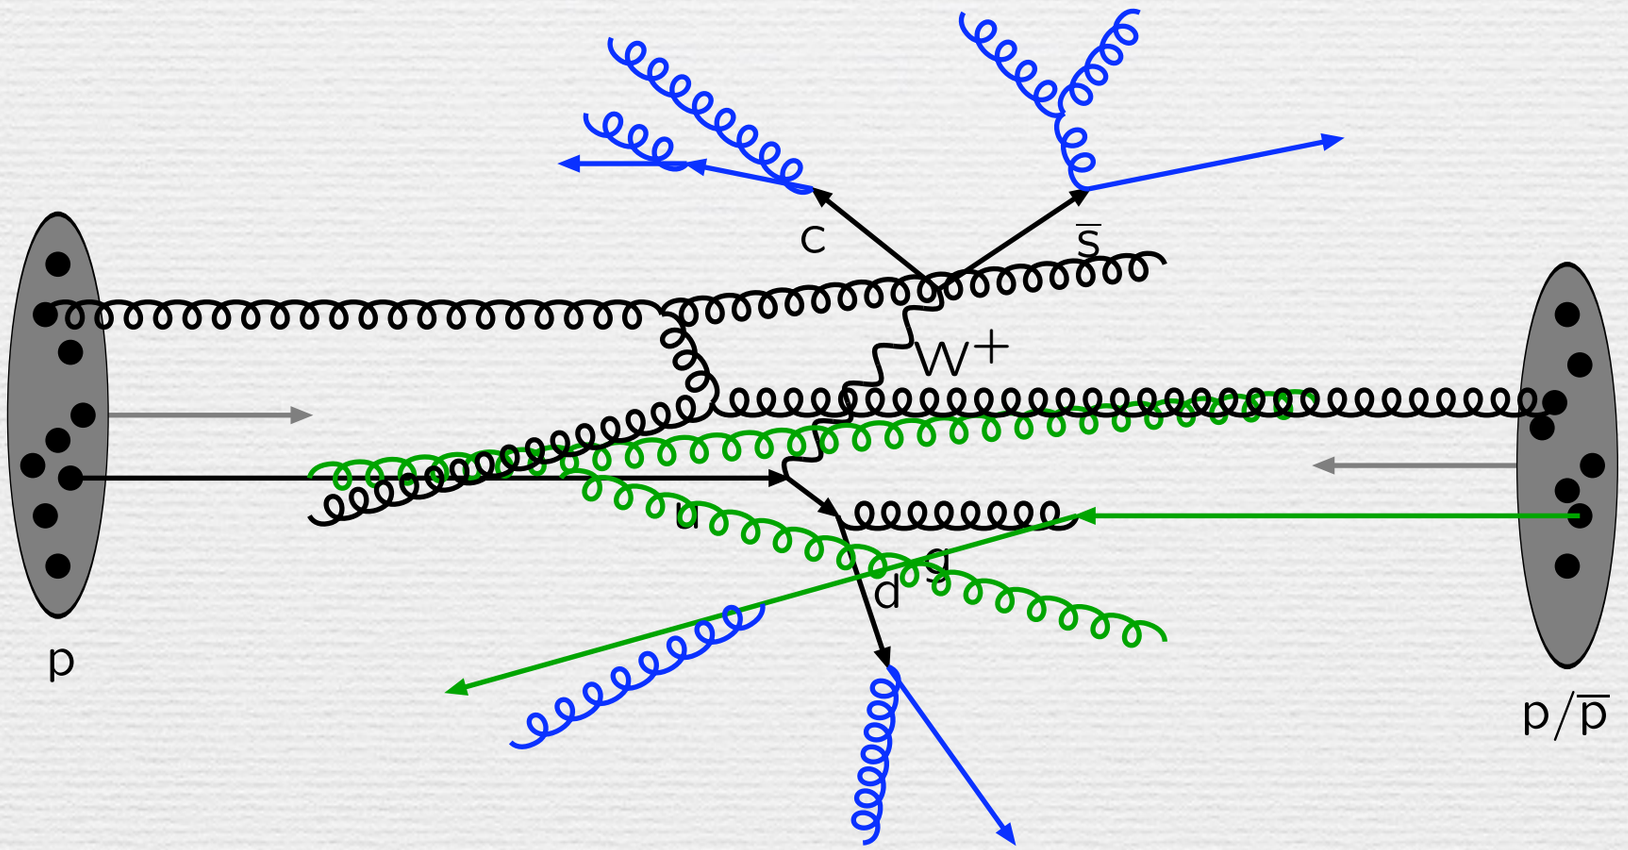
c
s̄
W+
u
d
s
p
p/p̄

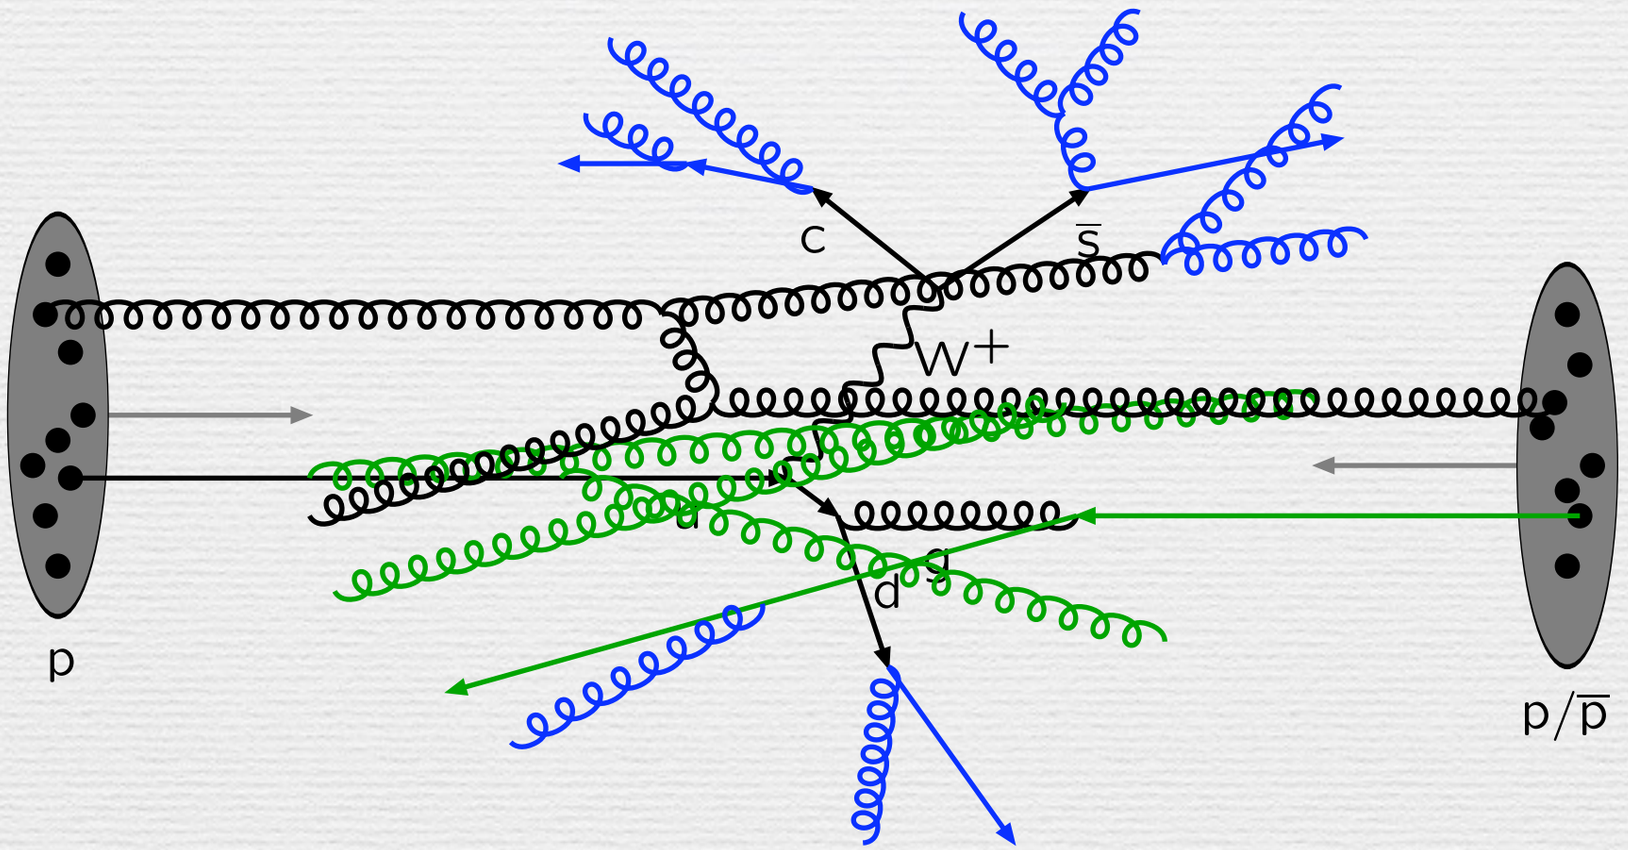
c
s̄
W+
d
s
d
p
p/p̄

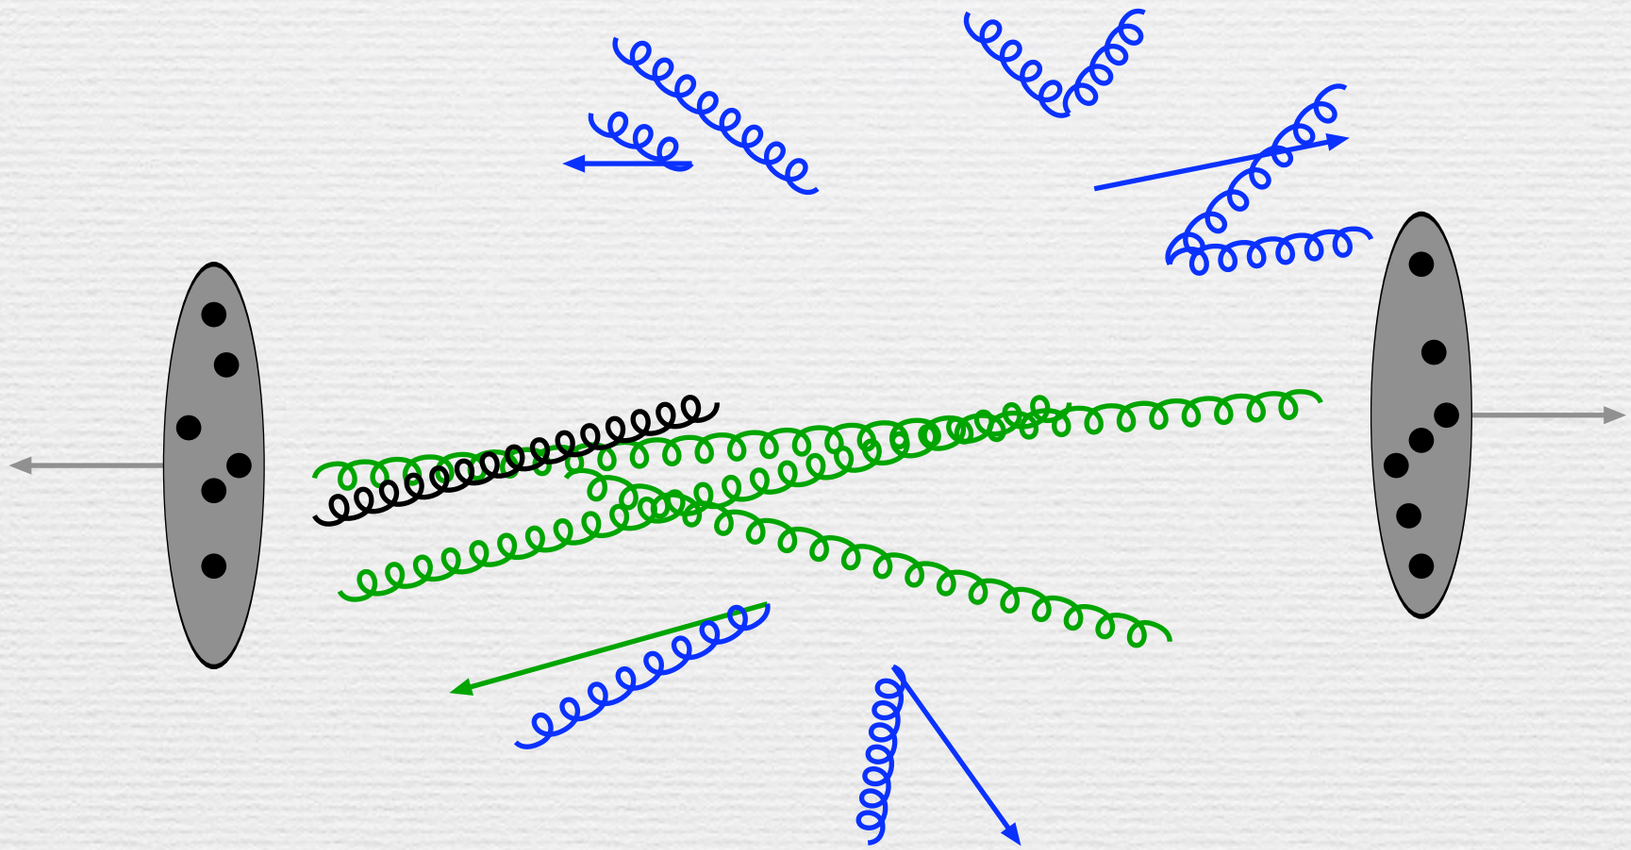

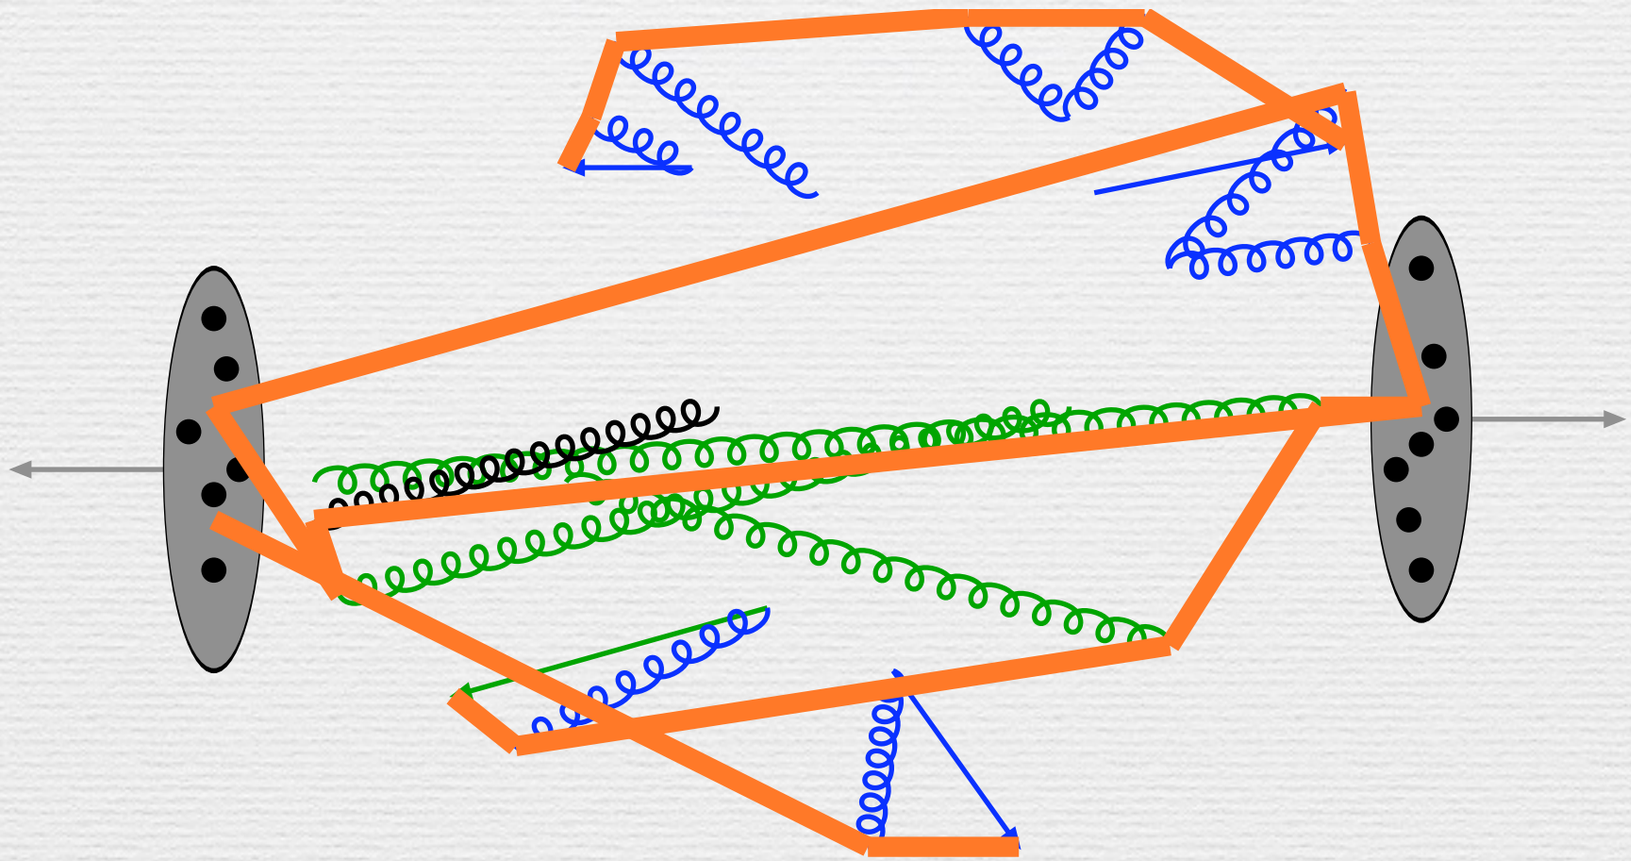

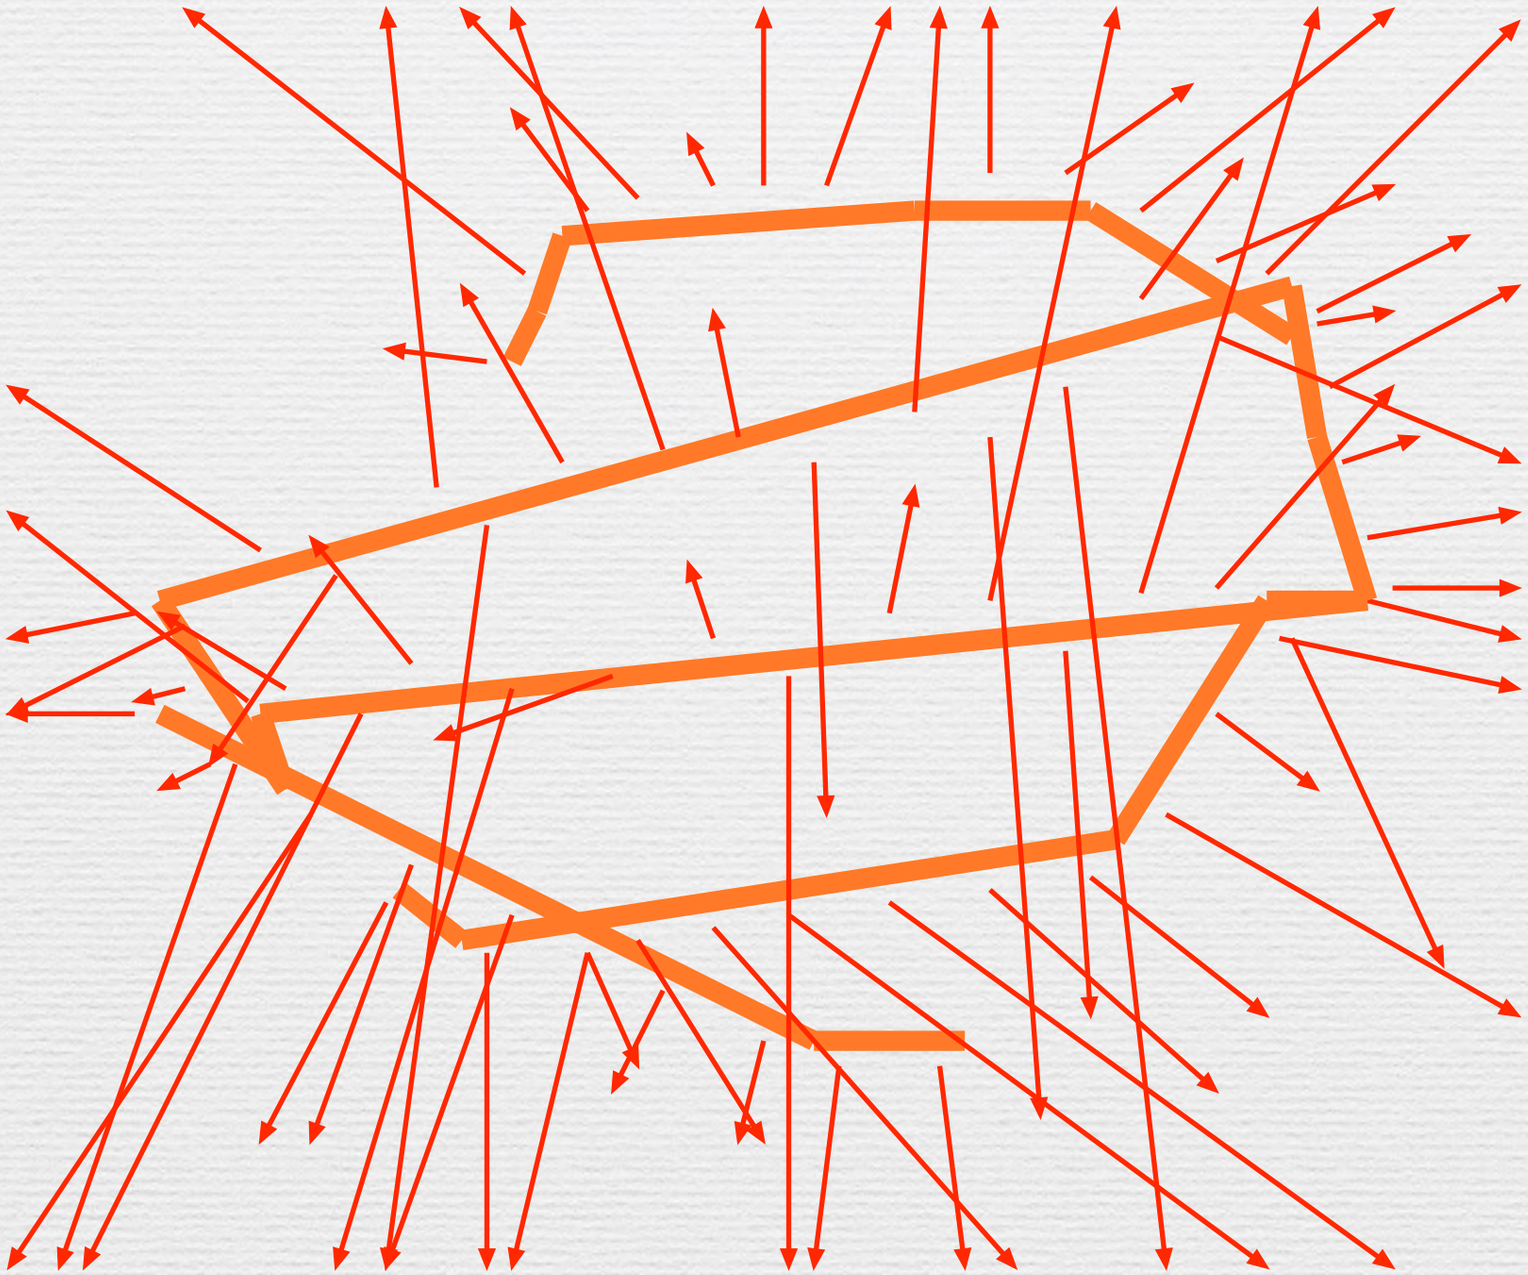

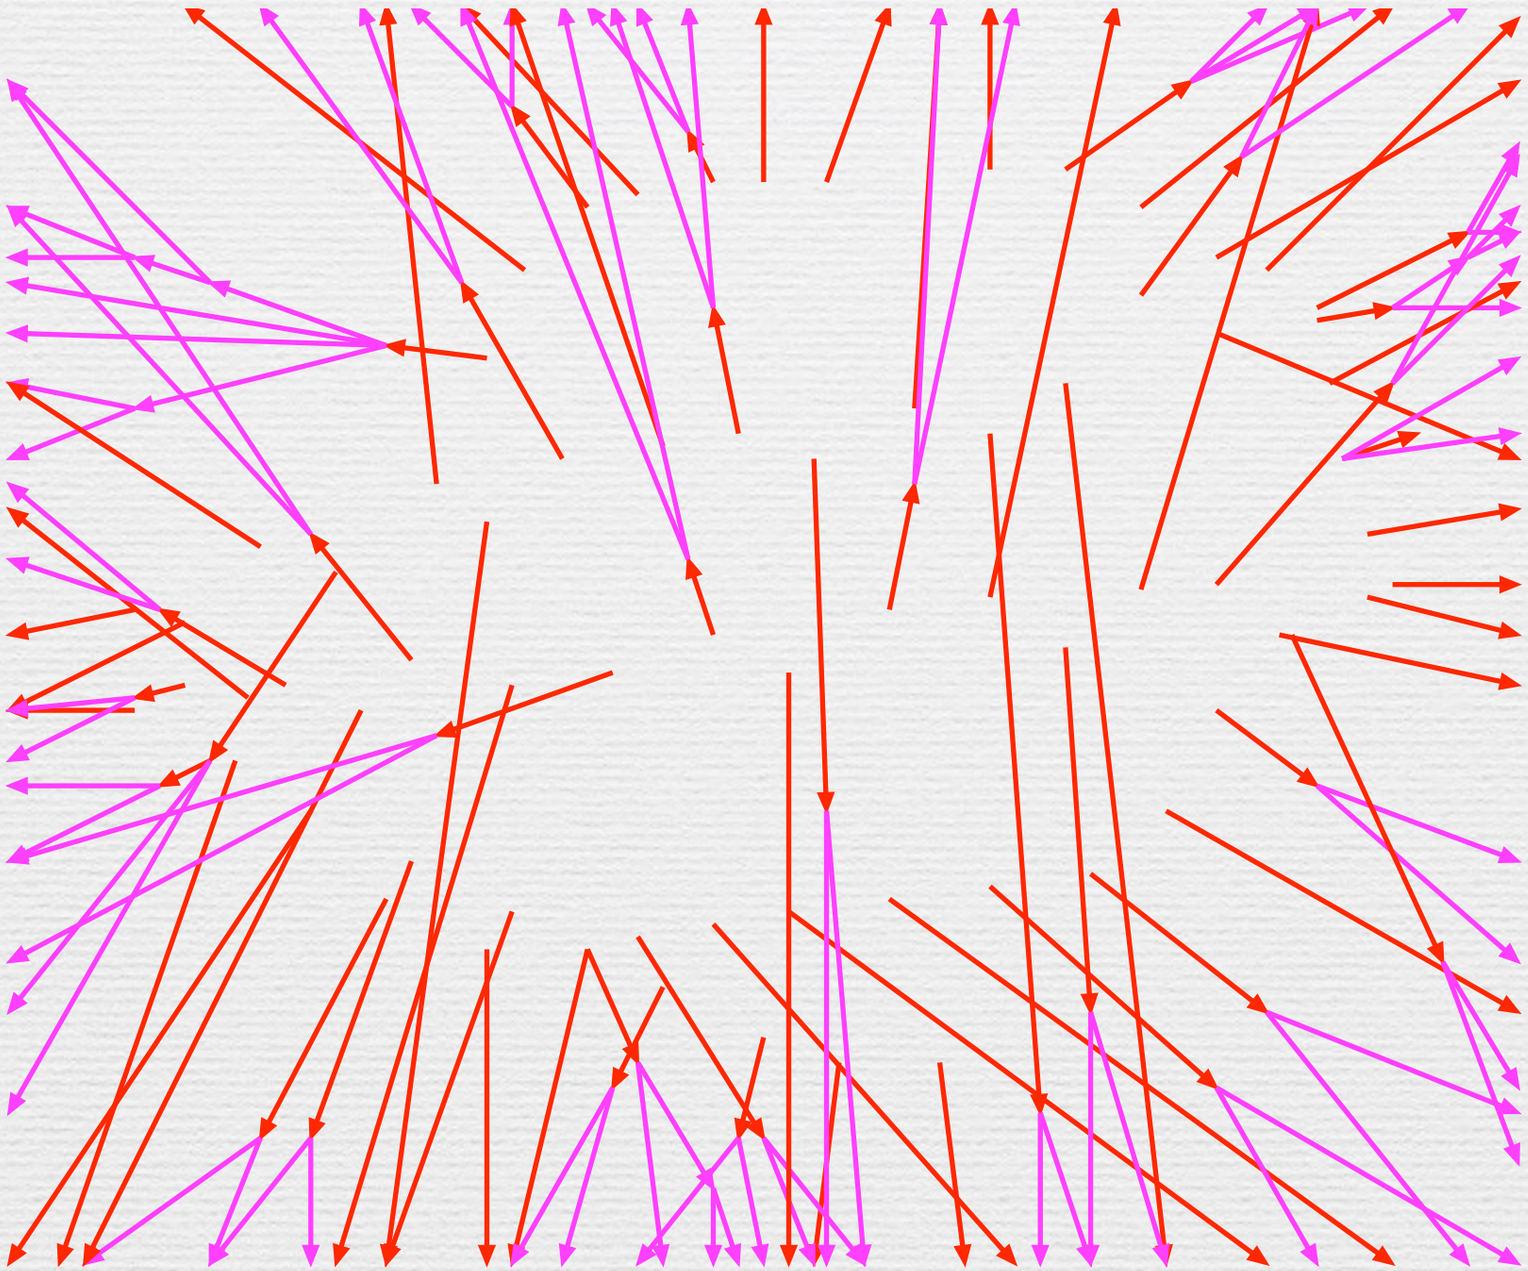

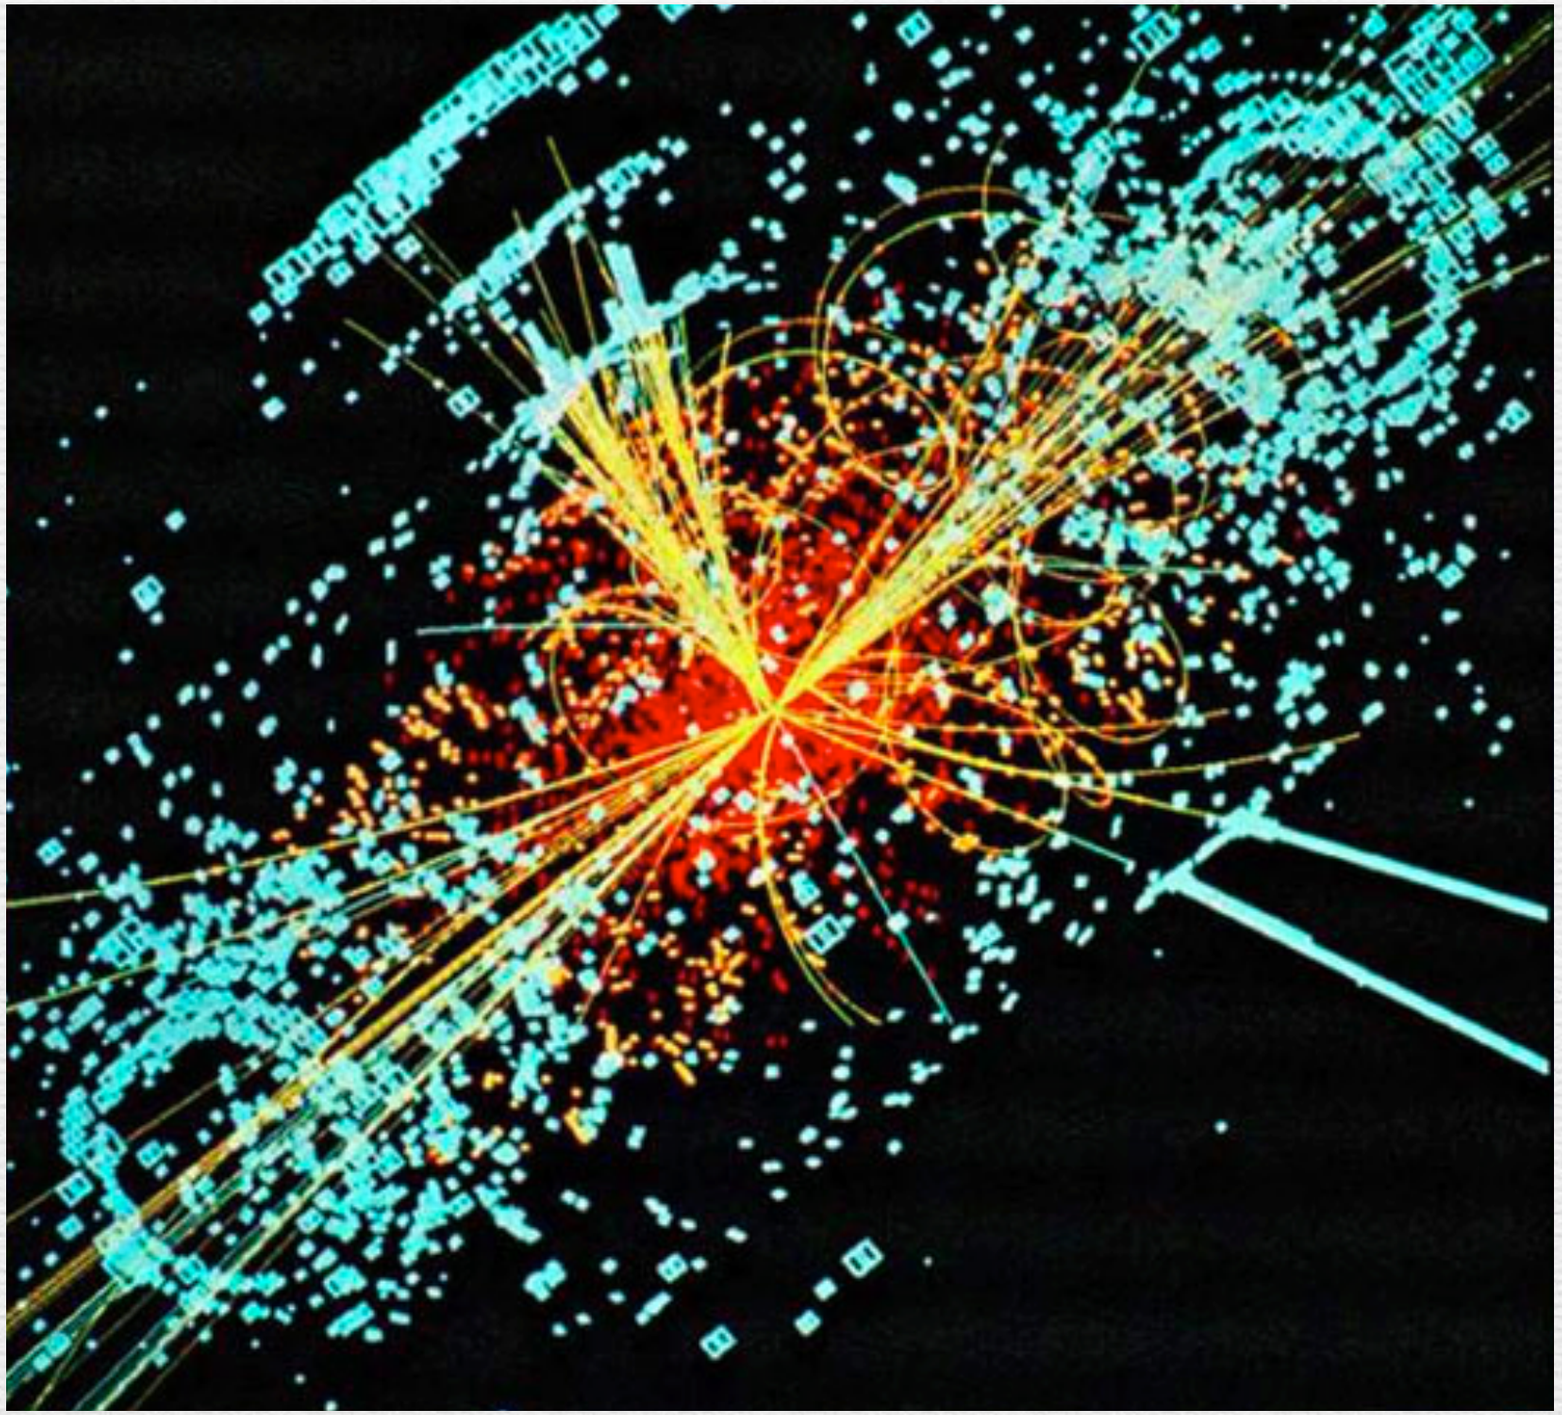

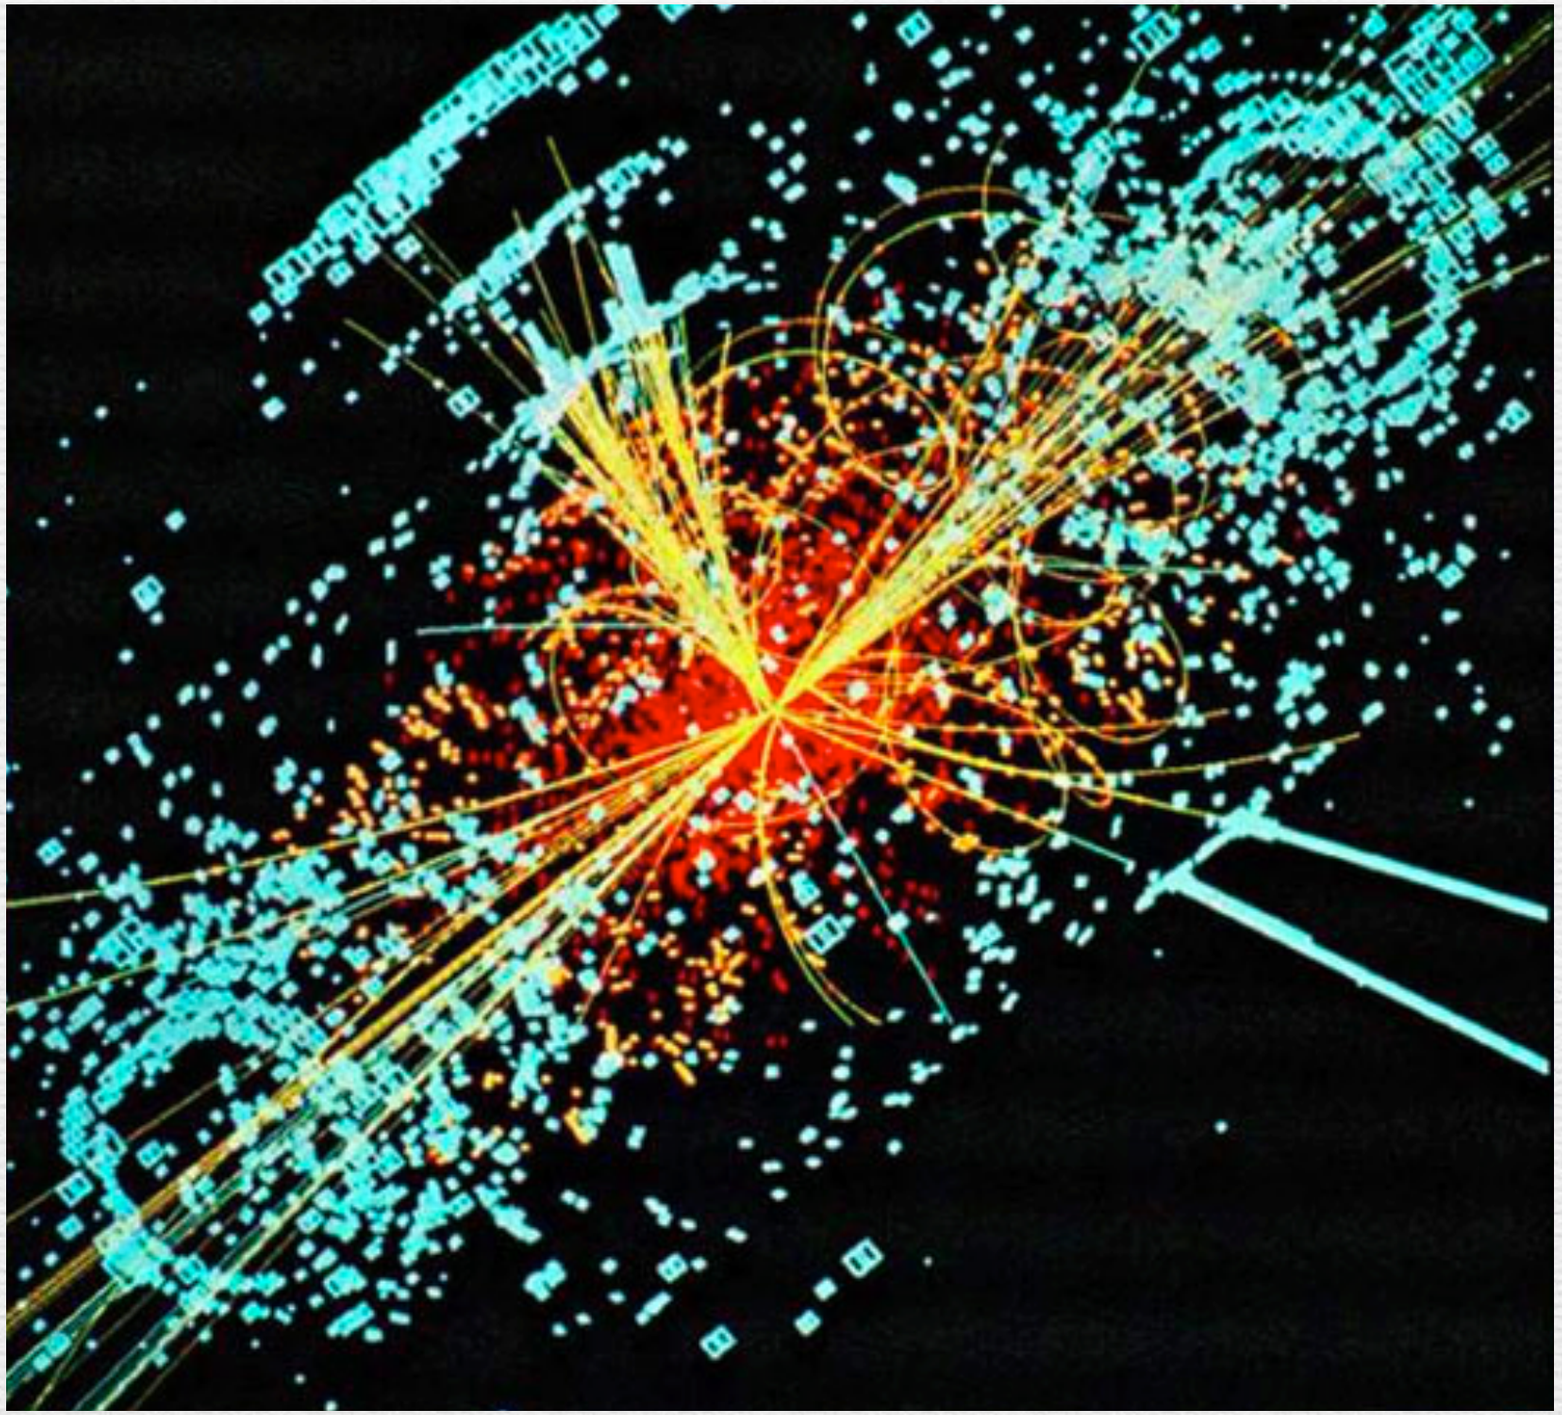

Simulated data sets of millions of events

Nature

Detector

Trigger

Reconstructed events

Theory model

Event generator

Simulated events

Analysis

Nature

Detector

Trigger

Reconstructed events

Theory model

Event generator

Simulated events

need millions!

Analysis

Each event independent

Batch farms are OK, but typical university clusters not large enough

Connect all participants transparently:

Worldwide LHC Computing Grid

WLCG tiered structure

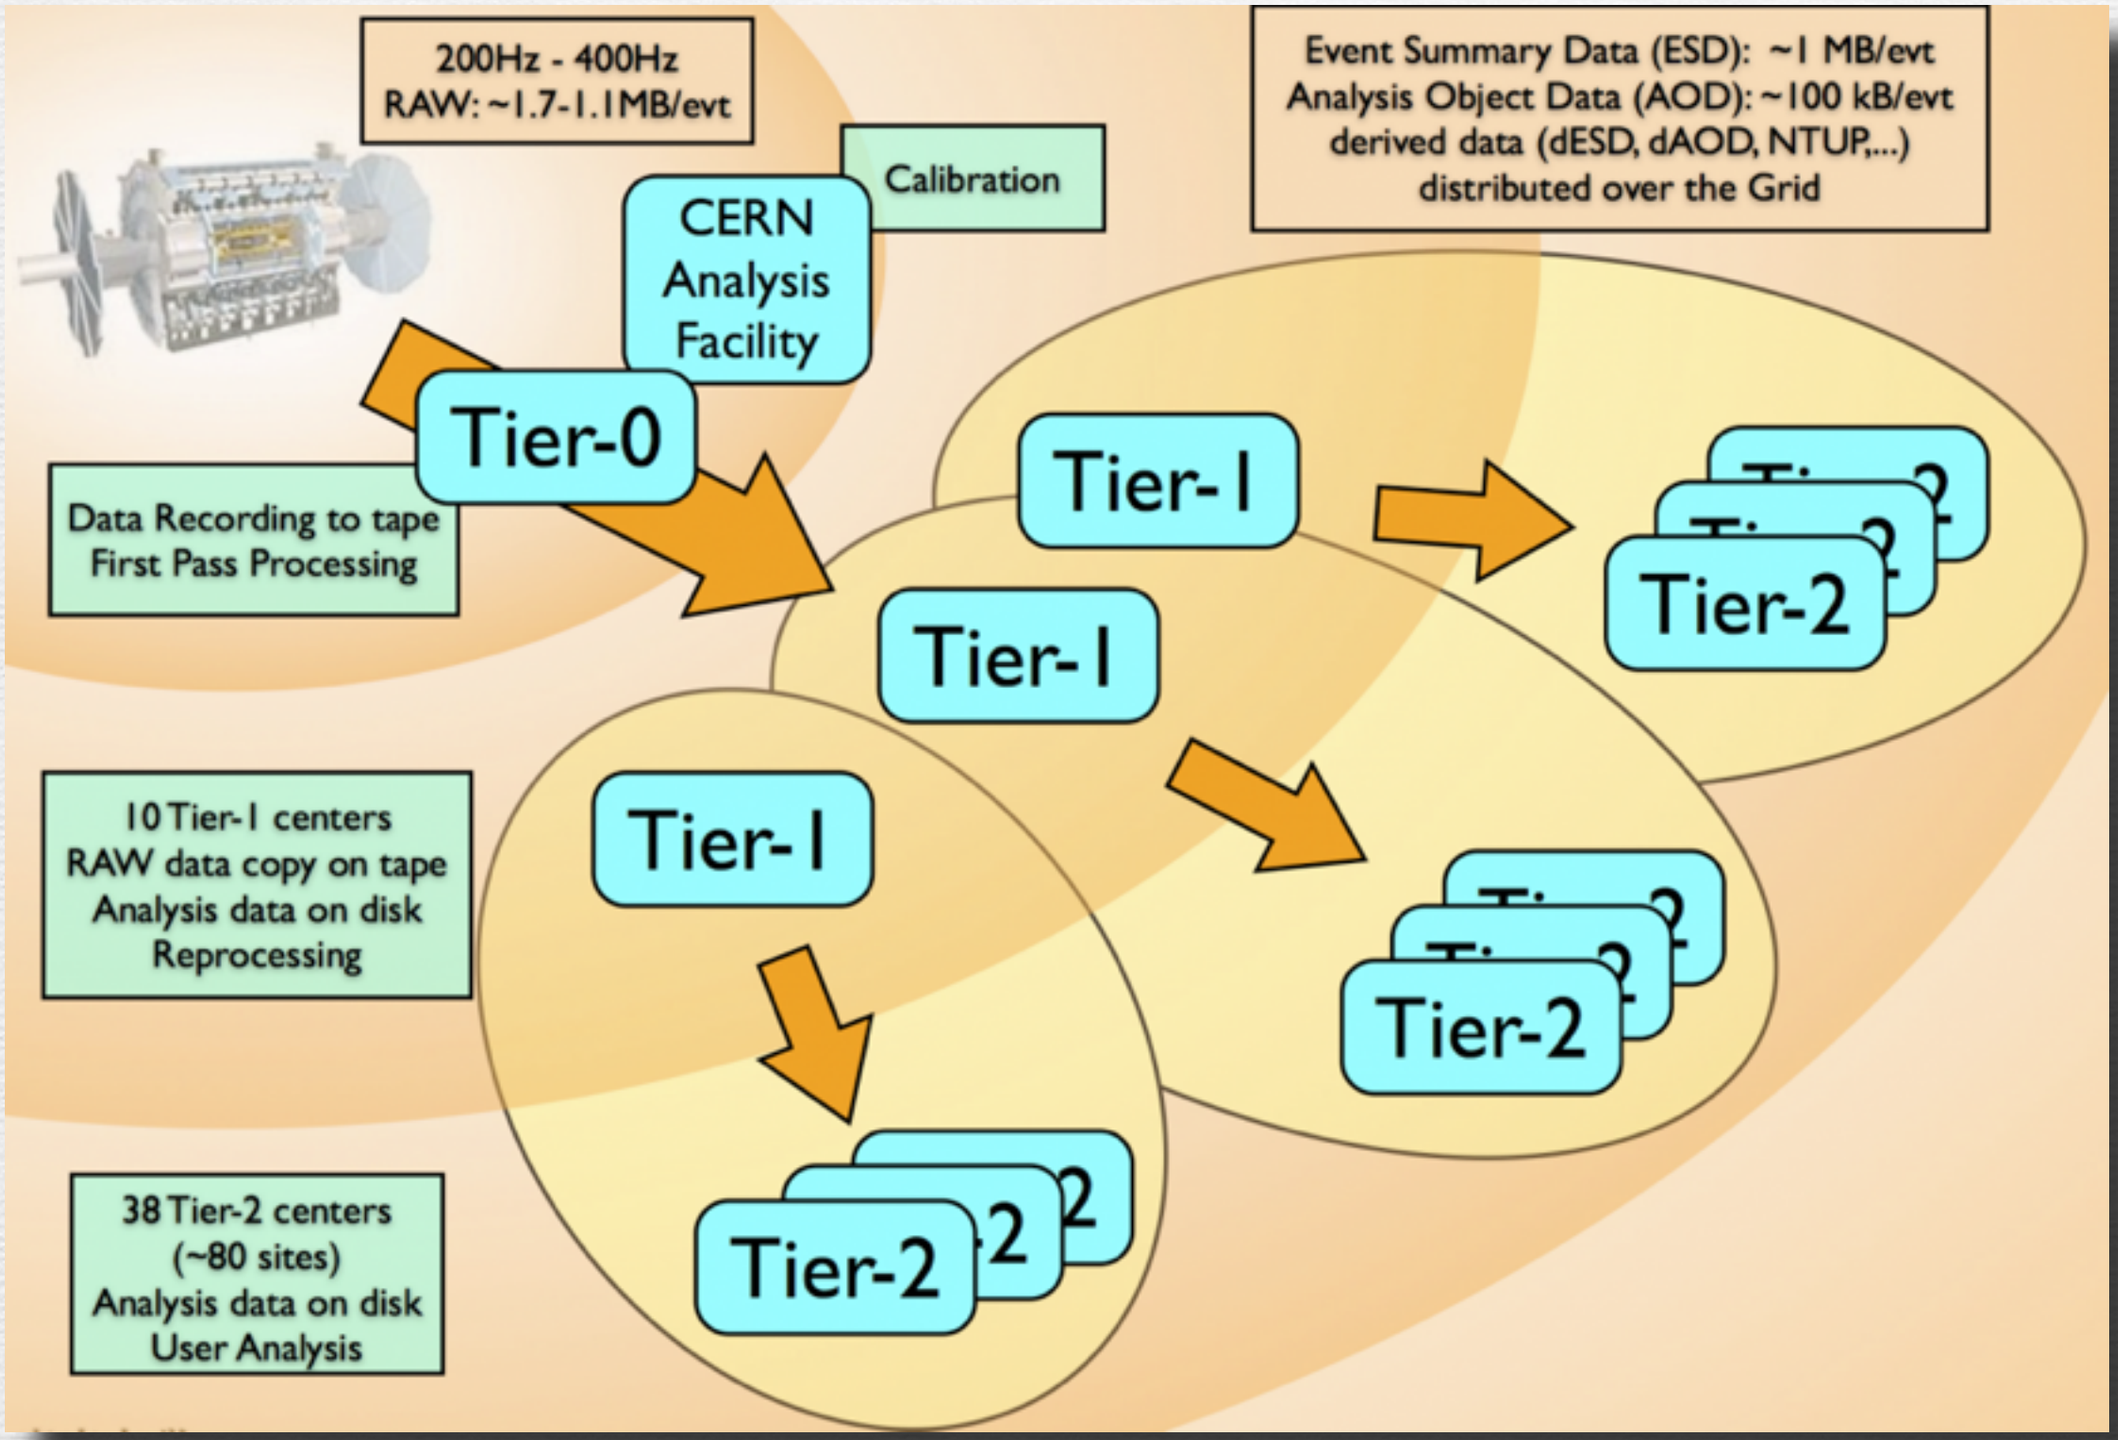

WLCG tiered structure

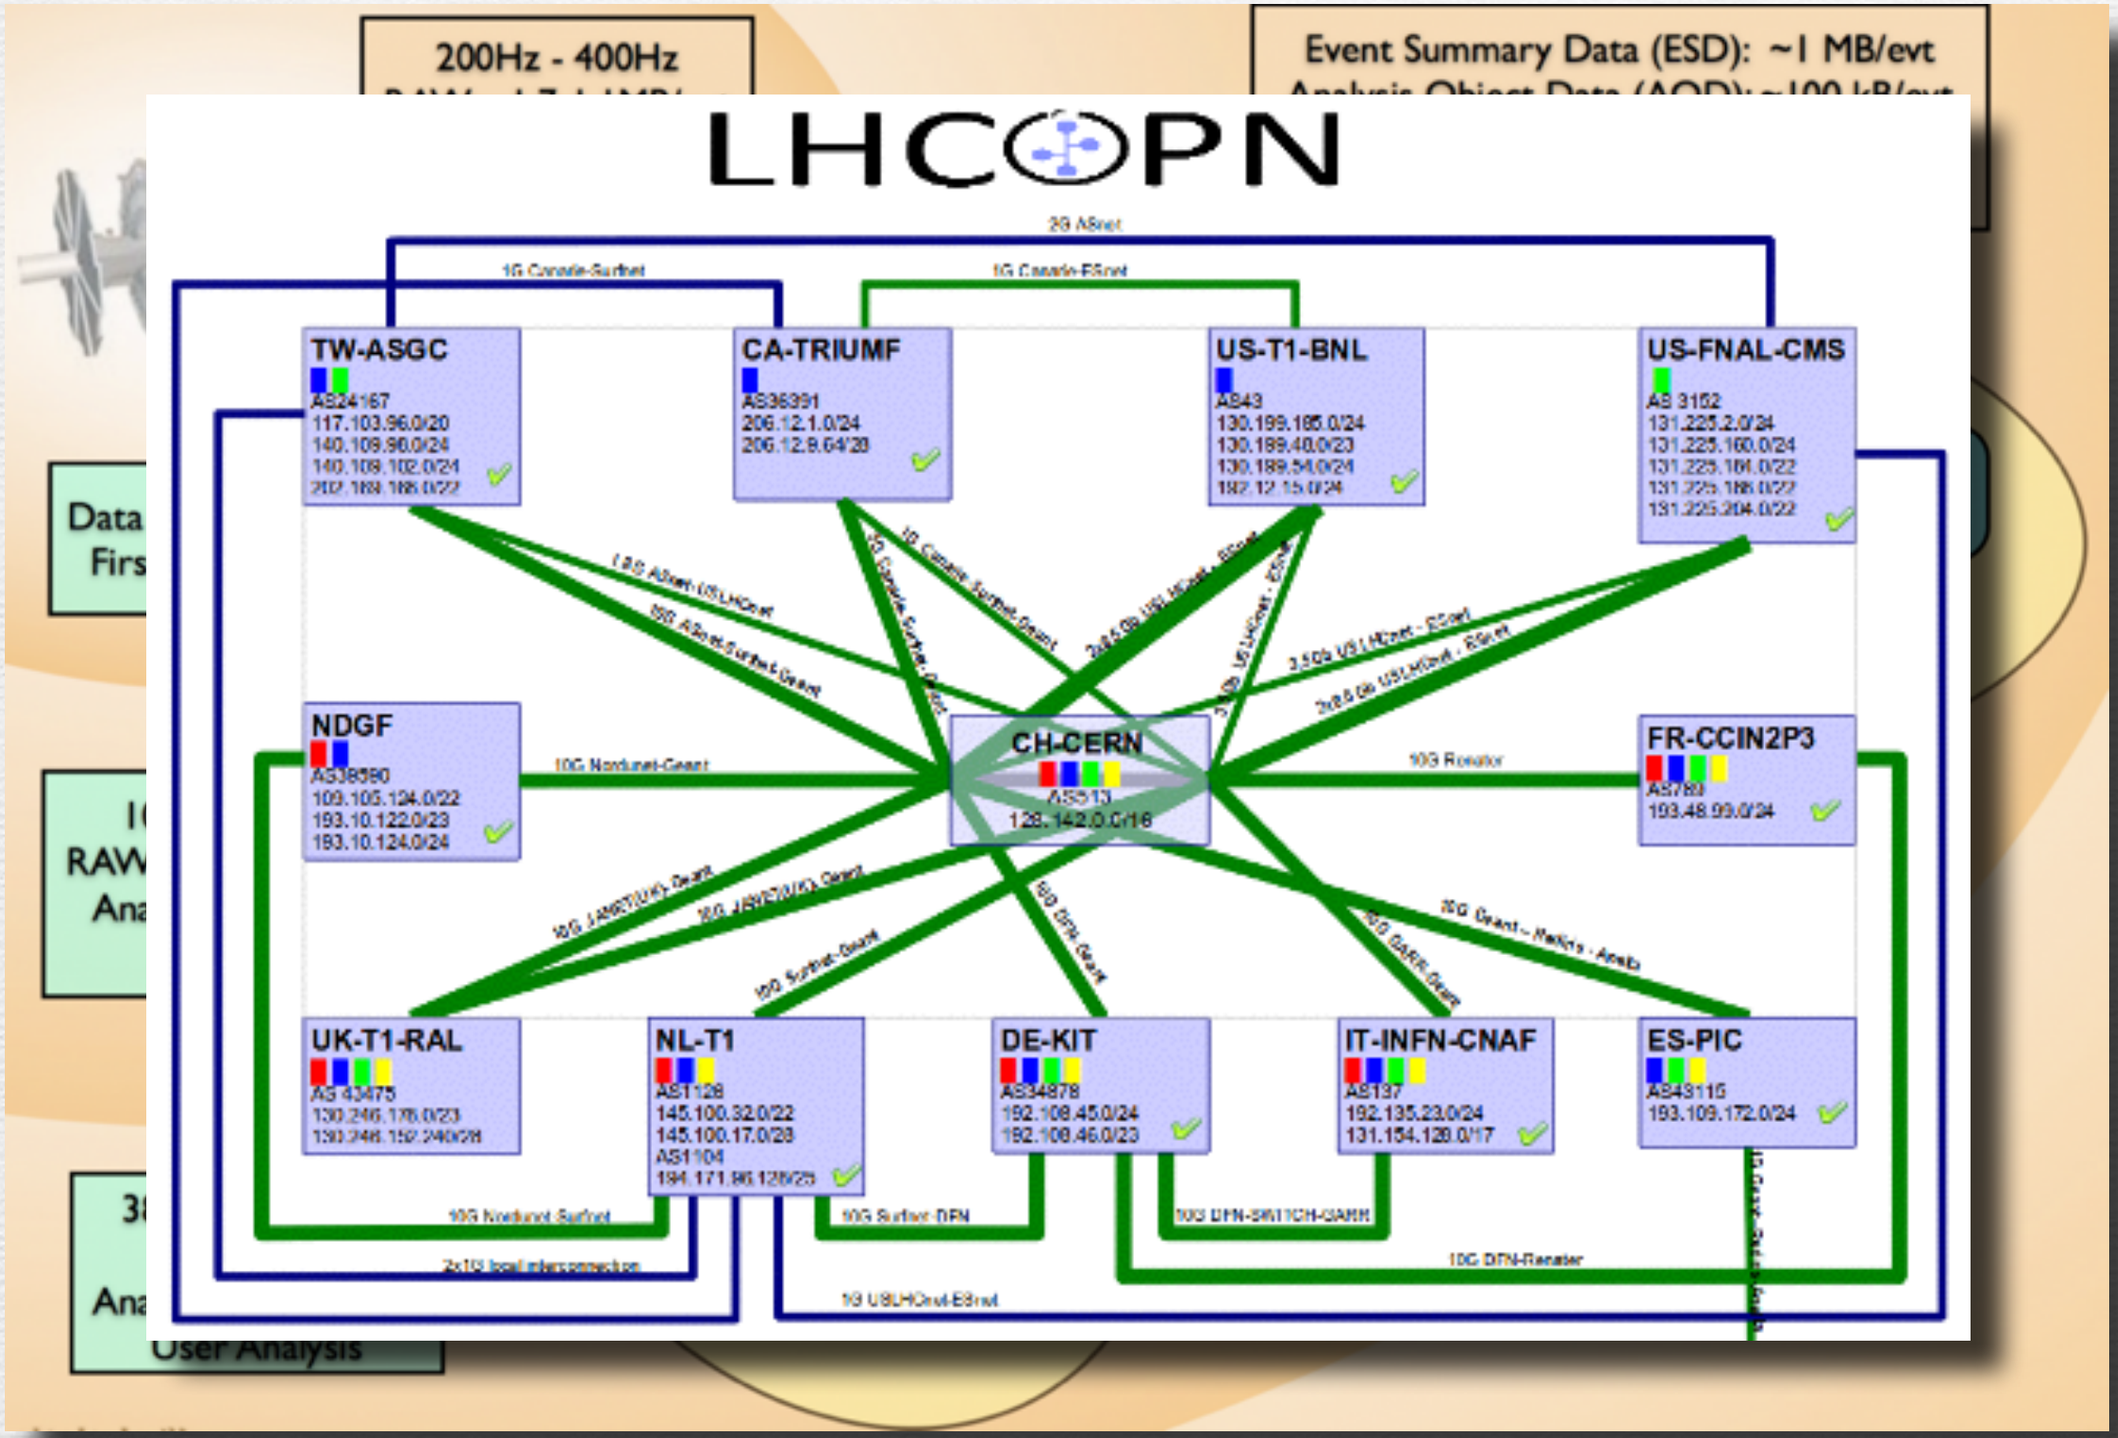

WLCG tiered structure

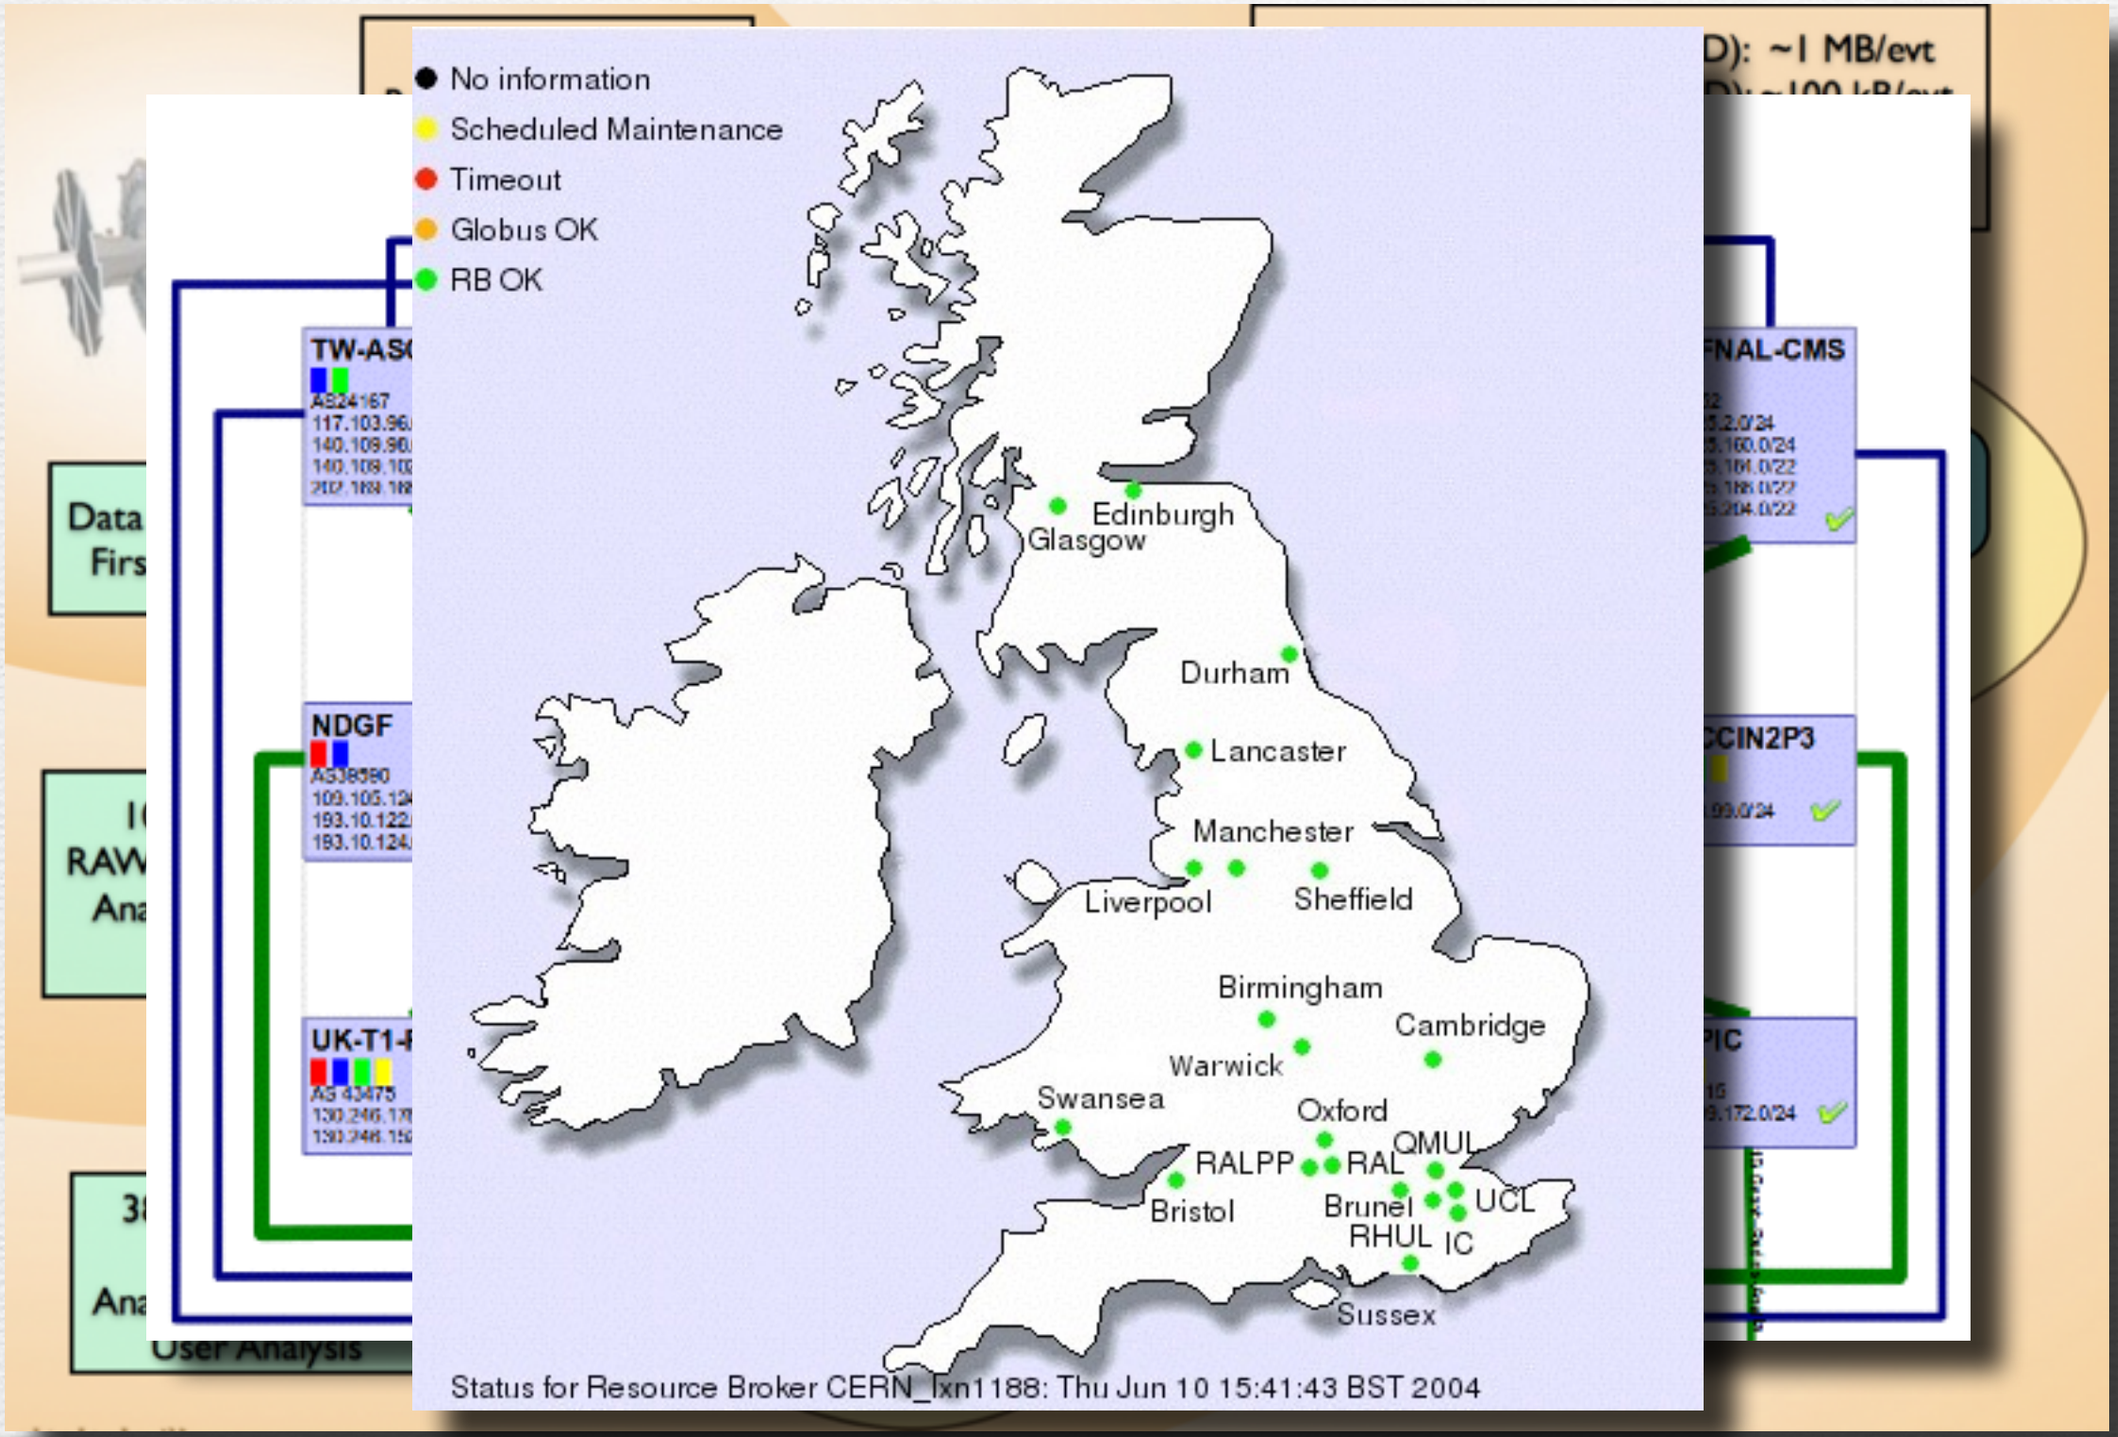

WLCG tiered structure

Parallel programs not considered for a long time

Why?

easier to keep CPUs busy by running independent copies

Nature

Detector

Trigger

Reconstructed events

Theory model

Event generator

Simulated events

need millions!

Analysis

Event generator

Matrix element

Monte-Carlo integration

Parton shower

Markov chain

Hadronization

book-keeping

Decays

Monte-Carlo integration

Event generator

Matrix element

Parton shower

Hadronization

Decays

Monte-Carlo integration

60 events / second

sequential, branchy algorithms

Event generator

Matrix element

Parton shower

Hadronization

Decays

Detector sim

Monte-Carlo integration

60 events / second

sequential, branchy algorithms

Event generator

Matrix element

Parton shower

Hadronization

Decays

Detector sim

Monte-Carlo integration

60 events / second

sequential, branchy algorithms

0.05 events / second
(aim: 1 event / sec)

Event generator

Matrix element

hugely variable!
if simple: ~1000 ev/s

Parton shower

60 events / second

sequential, branchy algorithms

Hadronization

Decays

Detector sim

0.05 events / second
(aim: 1 event / sec)

Event generator

Matrix element

Parton shower

Hadronization

Decays

Detector sim

hugely variable!
if simple: ~1000 ev/s

60 events / second

sequential, branchy algorithms

0.05 events / second
(aim: 1 event / sec)

Event generator

Matrix element

Parton shower

Hadronization

Decays

Detector sim

hugely variable!
if simple: ~1000 ev/s

60 events / second

sequential, branchy algorithms

0.05 events / second
(aim: 1 event / sec)

Cost breakdown – Generation

Generation costs = once per event

Matrix element unweighting

- challenging matrix element calculations, e.g. table
  Höche, Gleisberg (2008)
- really expensive: unweighting
  efficiency for complicated processes as low as 0.001%
  $\rightarrow 10^5$ ME calculations per event
- additionally merging with parton shower including ME clustering

$\rightarrow$ $\mathcal{O}(1\text{ day})/1000$ events in complicated cases

Process	[ms/pt.]
$gg \to 2g$	0.073
$gg \to 3g$	0.339
$gg \to 4g$	1.67
$gg \to 5g$	8.98
$gg \to 6g$	49.6
$gg \to 7g$	298.
$gg \to 8g$	1990.
$gg \to 9g$	13100.
$gg \to 10g$	96000.

Remaining generation chain

- remaining cost of event independent of ME+PS: < 0.5 s/evt
- includes hadronisation, decays, QED FSR, multiple parton interactions

Example: W + 0,1,2j@NLO + 3,4,5j@LO

- from ATLAS central production: $\sim$ 12h/2000 evts $\approx$ 20 s/evt

$\rightarrow$ needs work if detector simulation gets to $\mathcal{O}(1\text{s})$/evt

Consequences for MC authors:

- Time spent on parallel implementations
useful, unless we rethink theory algorithms completely with parallel implementations in mind.

- Time spent on profiling and increasing serial execution speed is more promising. Several low hanging fruit still wait to be discovered.

Recently:
20% speedup from 3-line code change.
Relevant for situations where event generators are run stand-alone with fast MEs and fast downstream processing.