Processing Petabytes per Second with the ATLAS Experiment at the Large Hadron Collider in CERN

GPU Technology Conference 2010

P.J. Clark, J. Henderson, C. Jones, M. Rovatsou, A. Washbrook
(P.J.Clark@ed.ac.uk)

University of Edinburgh

22nd September 2010

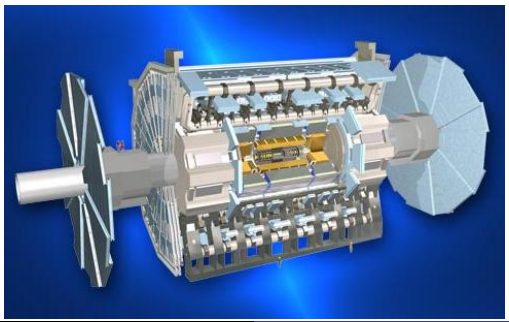

Outline

1. Introduction to the LHC and ATLAS
   - The Large Hadron Collider (LHC)
   - The ATLAS detector
   - The Higgs Boson
2. Areas of study for GPGPU adoption
   - Particle tracking in a magnetic field
   - The ATLAS trigger and data acquisition
   - The worldwide LHC computing grid
3. Using GPUs in the high level trigger
   - The Level 2 Trigger
   - The Z finder algorithm
   - The Kalman Filter

Outline

1. Introduction to the LHC and ATLAS
   - The Large Hadron Collider (LHC)
   - The ATLAS detector
   - The Higgs Boson
2. Areas of study for GPGPU adoption
   - Particle tracking in a magnetic field
   - The ATLAS trigger and data acquisition
   - The worldwide LHC computing grid
3. Using GPUs in the high level trigger
   - The Level 2 Trigger
   - The Z finder algorithm
   - The Kalman Filter

The Large Hadron Collider

The Large Hadron Collider

The Large Hadron Collider

- World's largest collider (27 km circumference)

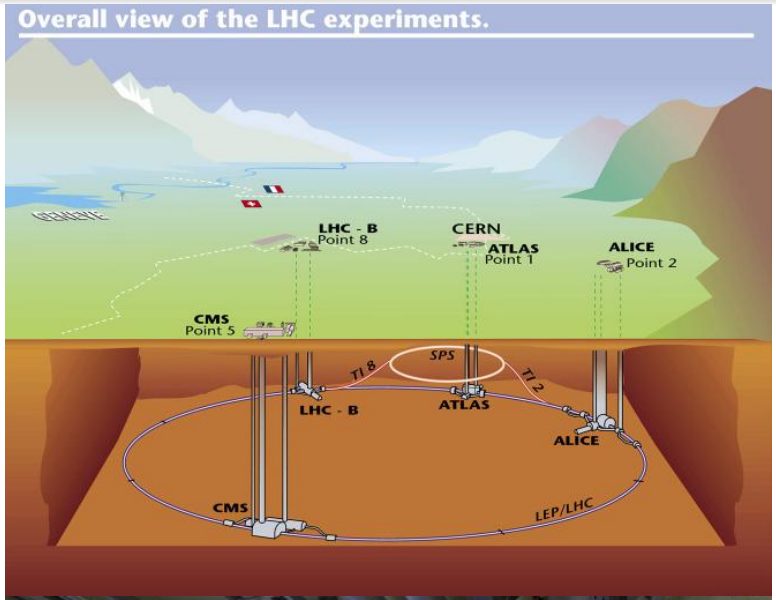

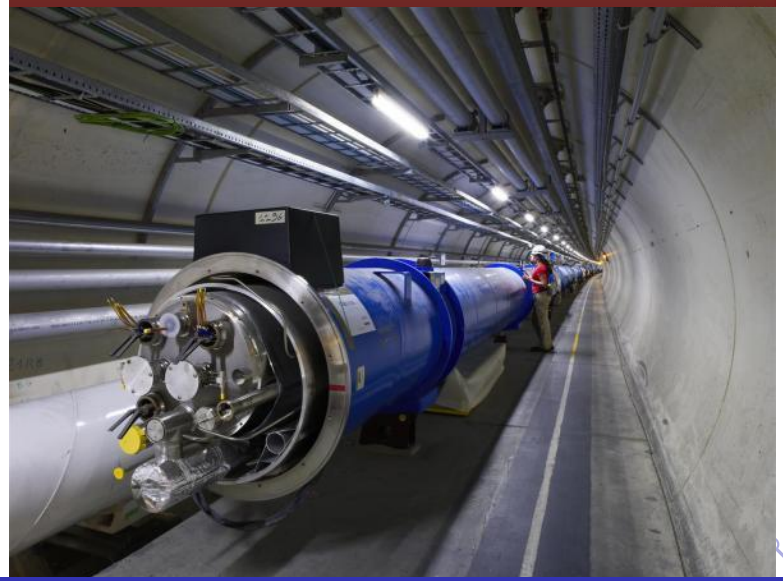

The Large Hadron Collider

- World's largest collider
  (27 km circumference)
- Highest energy: protons with 7 TeV
  (99.9999991% of speed of light)

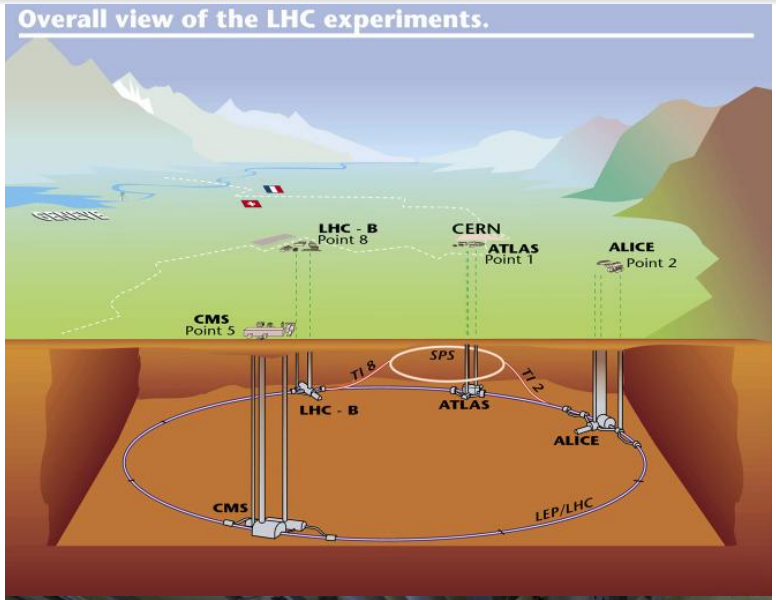

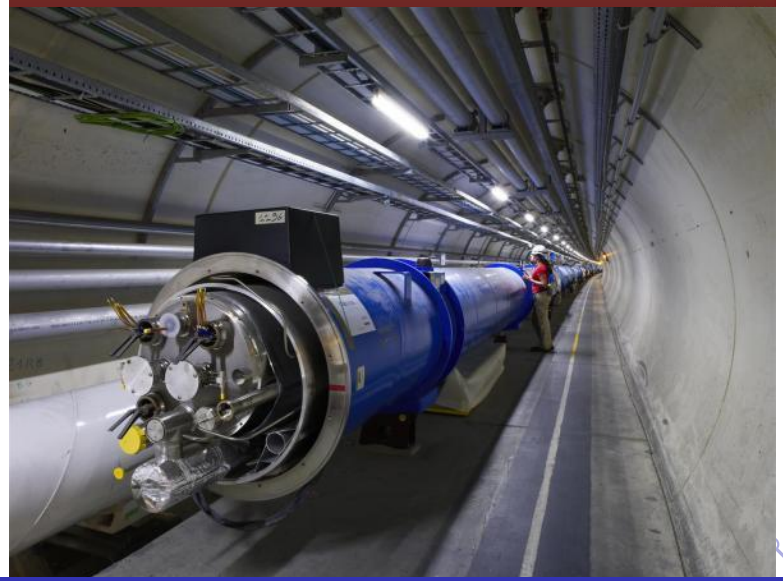

The Large Hadron Collider

- World's largest collider
  (27 km circumference)
- Highest energy: protons with 7 TeV
  (99.9999991% of speed of light)
- 8.3 T magnets cooled to 1.9 K
  (cooler than the universe)

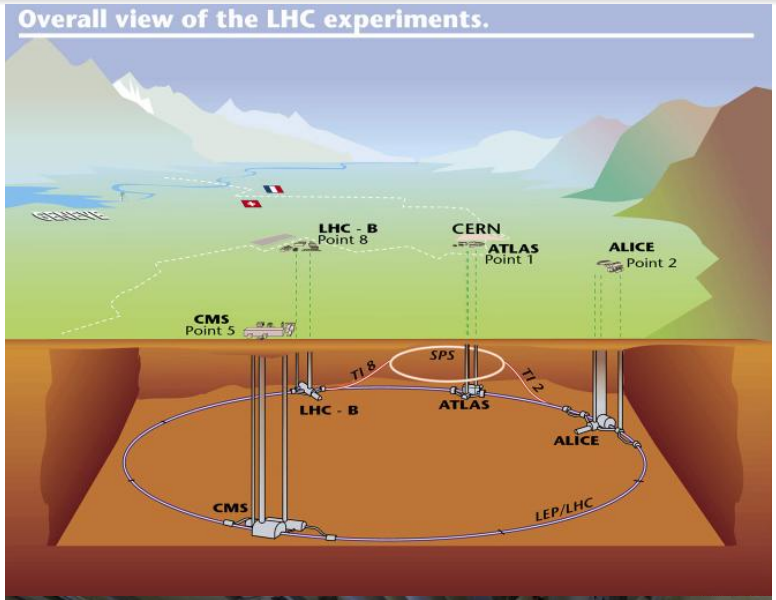

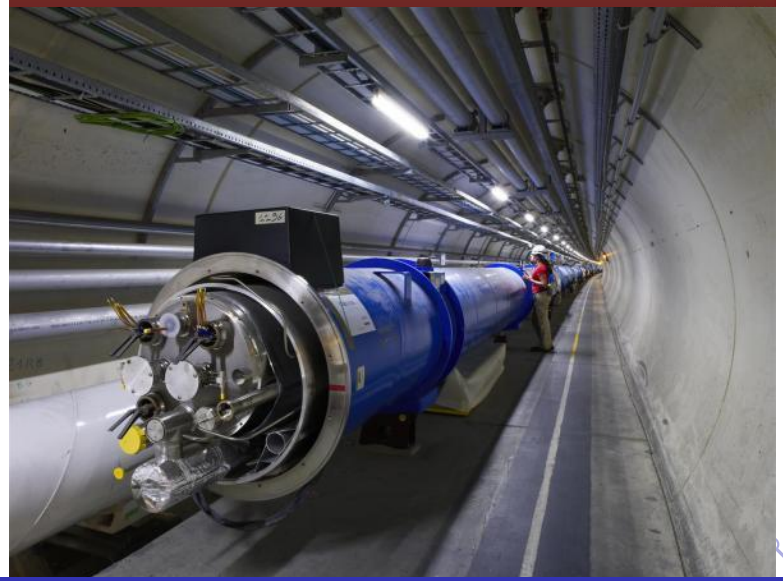

The Large Hadron Collider

- World's largest collider
  (27 km circumference)
- Highest energy: protons with 7 TeV
  (99.9999991% of speed of light)
- 8.3 T magnets cooled to 1.9 K
  (cooler than the universe)
- Total beam energy 724 MJ
  (Nimitz aircraft carrier at 14 km/h)

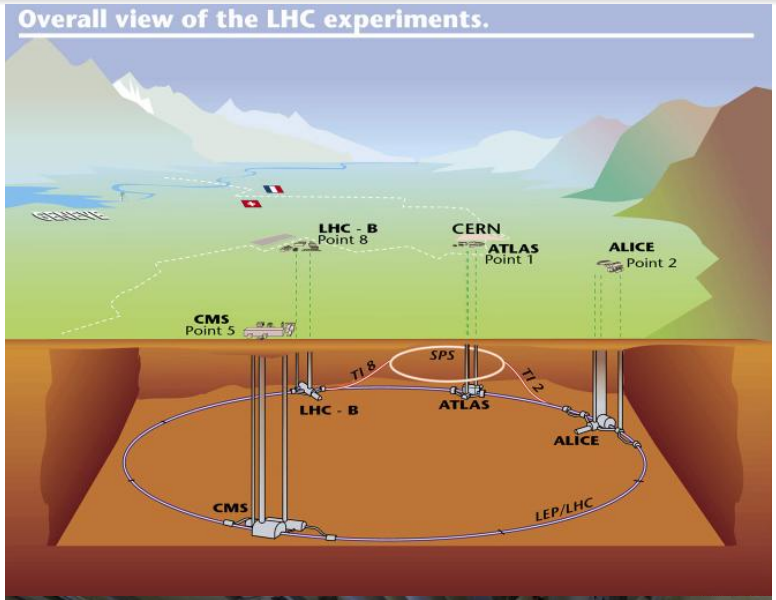

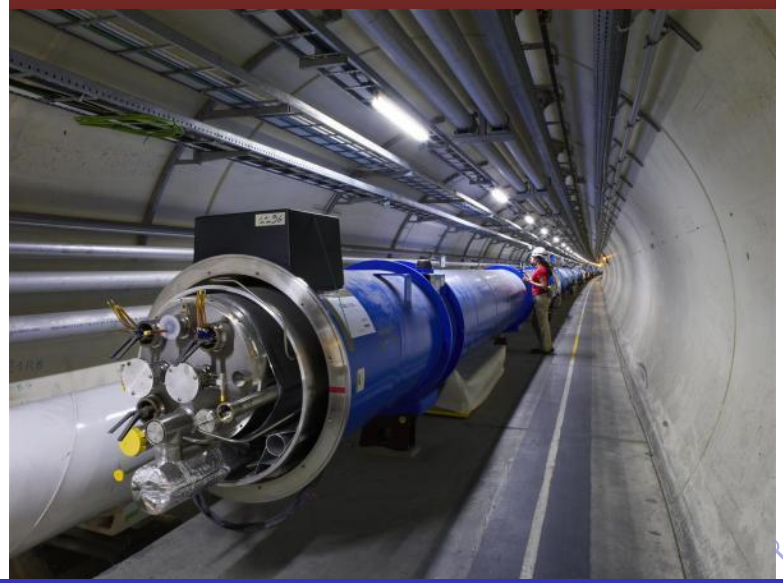

The Large Hadron Collider

- World's largest collider
  (27 km circumference)
- Highest energy: protons with 7 TeV
  (99.9999991% of speed of light)
- 8.3 T magnets cooled to 1.9 K
  (cooler than the universe)
- Total beam energy 724 MJ
  (Nimitz aircraft carrier at 14 km/h)
- Beam size 0.2 mm
  (smaller than a needle)

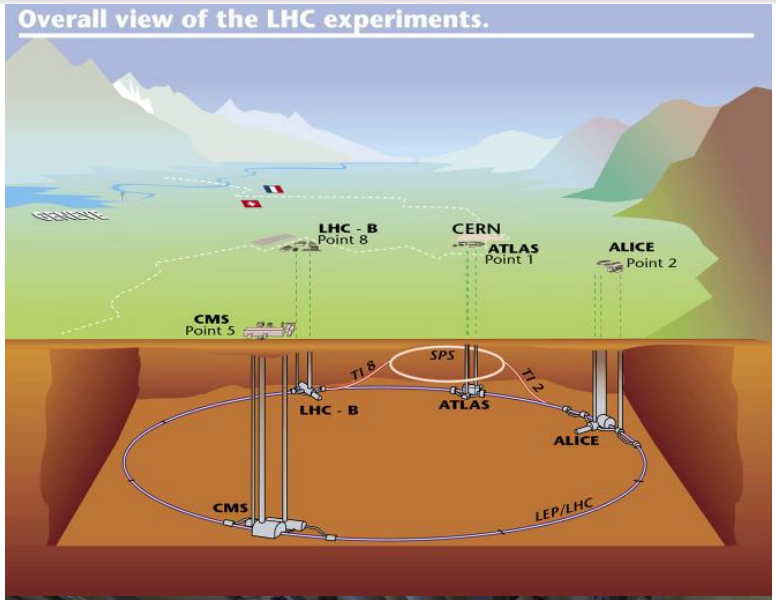

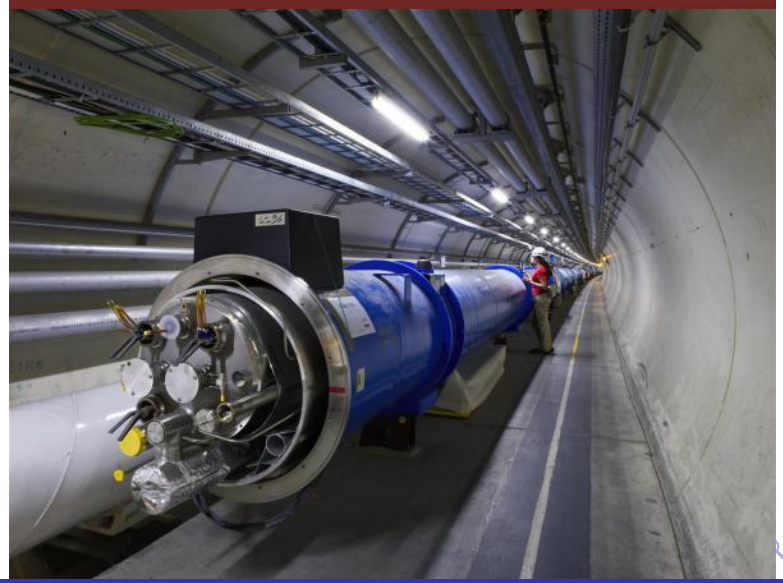

The Large Hadron Collider

- World's largest collider
  (27 km circumference)
- Highest energy: protons with 7 TeV
  (99.9999991% of speed of light)
- 8.3 T magnets cooled to 1.9 K
  (cooler than the universe)
- Total beam energy 724 MJ
  (Nimitz aircraft carrier at 14 km/h)
- Beam size 0.2 mm
  (smaller than a needle)
- Interaction scale studied $< 10^{-16}$ m
  (0.0000000000001 mm)

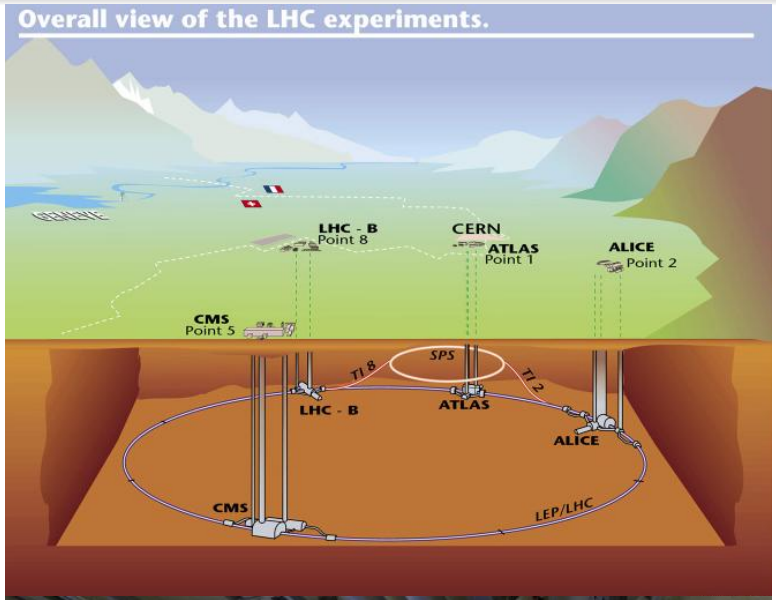

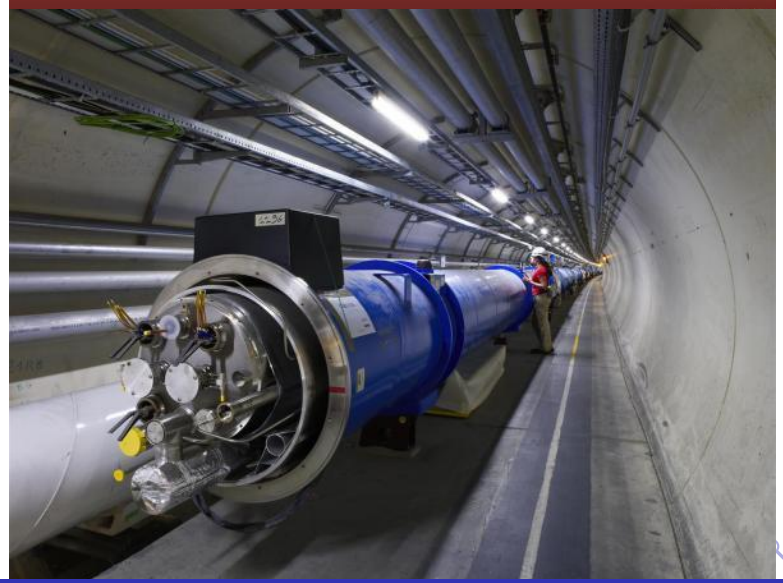

The Large Hadron Collider

- World's largest collider
  (27 km circumference)
- Highest energy: protons with 7 TeV
  (99.9999991% of speed of light)
- 8.3 T magnets cooled to 1.9 K
  (cooler than the universe)
- Total beam energy 724 MJ
  (Nimitz aircraft carrier at 14 km/h)
- Beam size 0.2 mm
  (smaller than a needle)
- Interaction scale studied $< 10^{-16}$ m
  (0.0000000000001 mm)

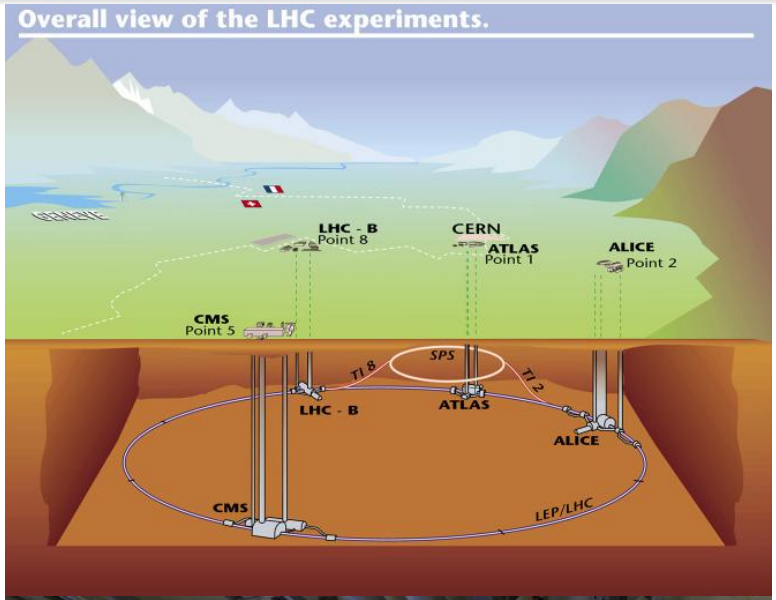

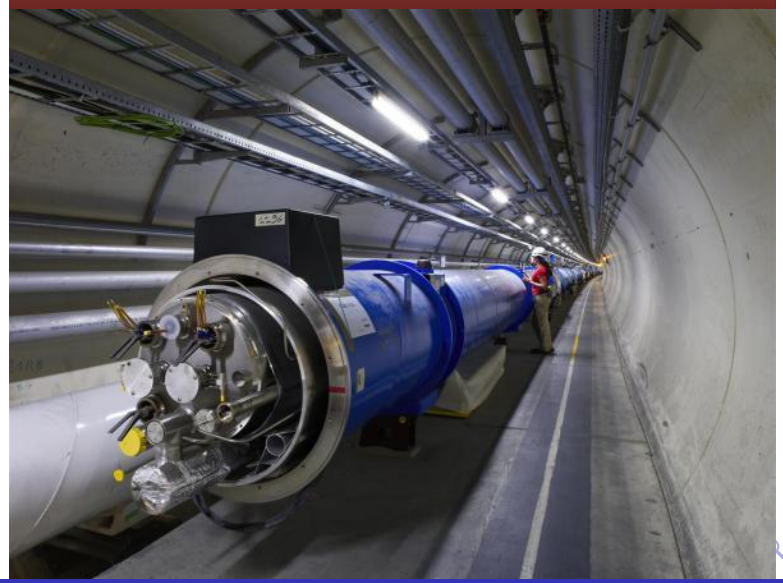

Attracted Large Media Interest

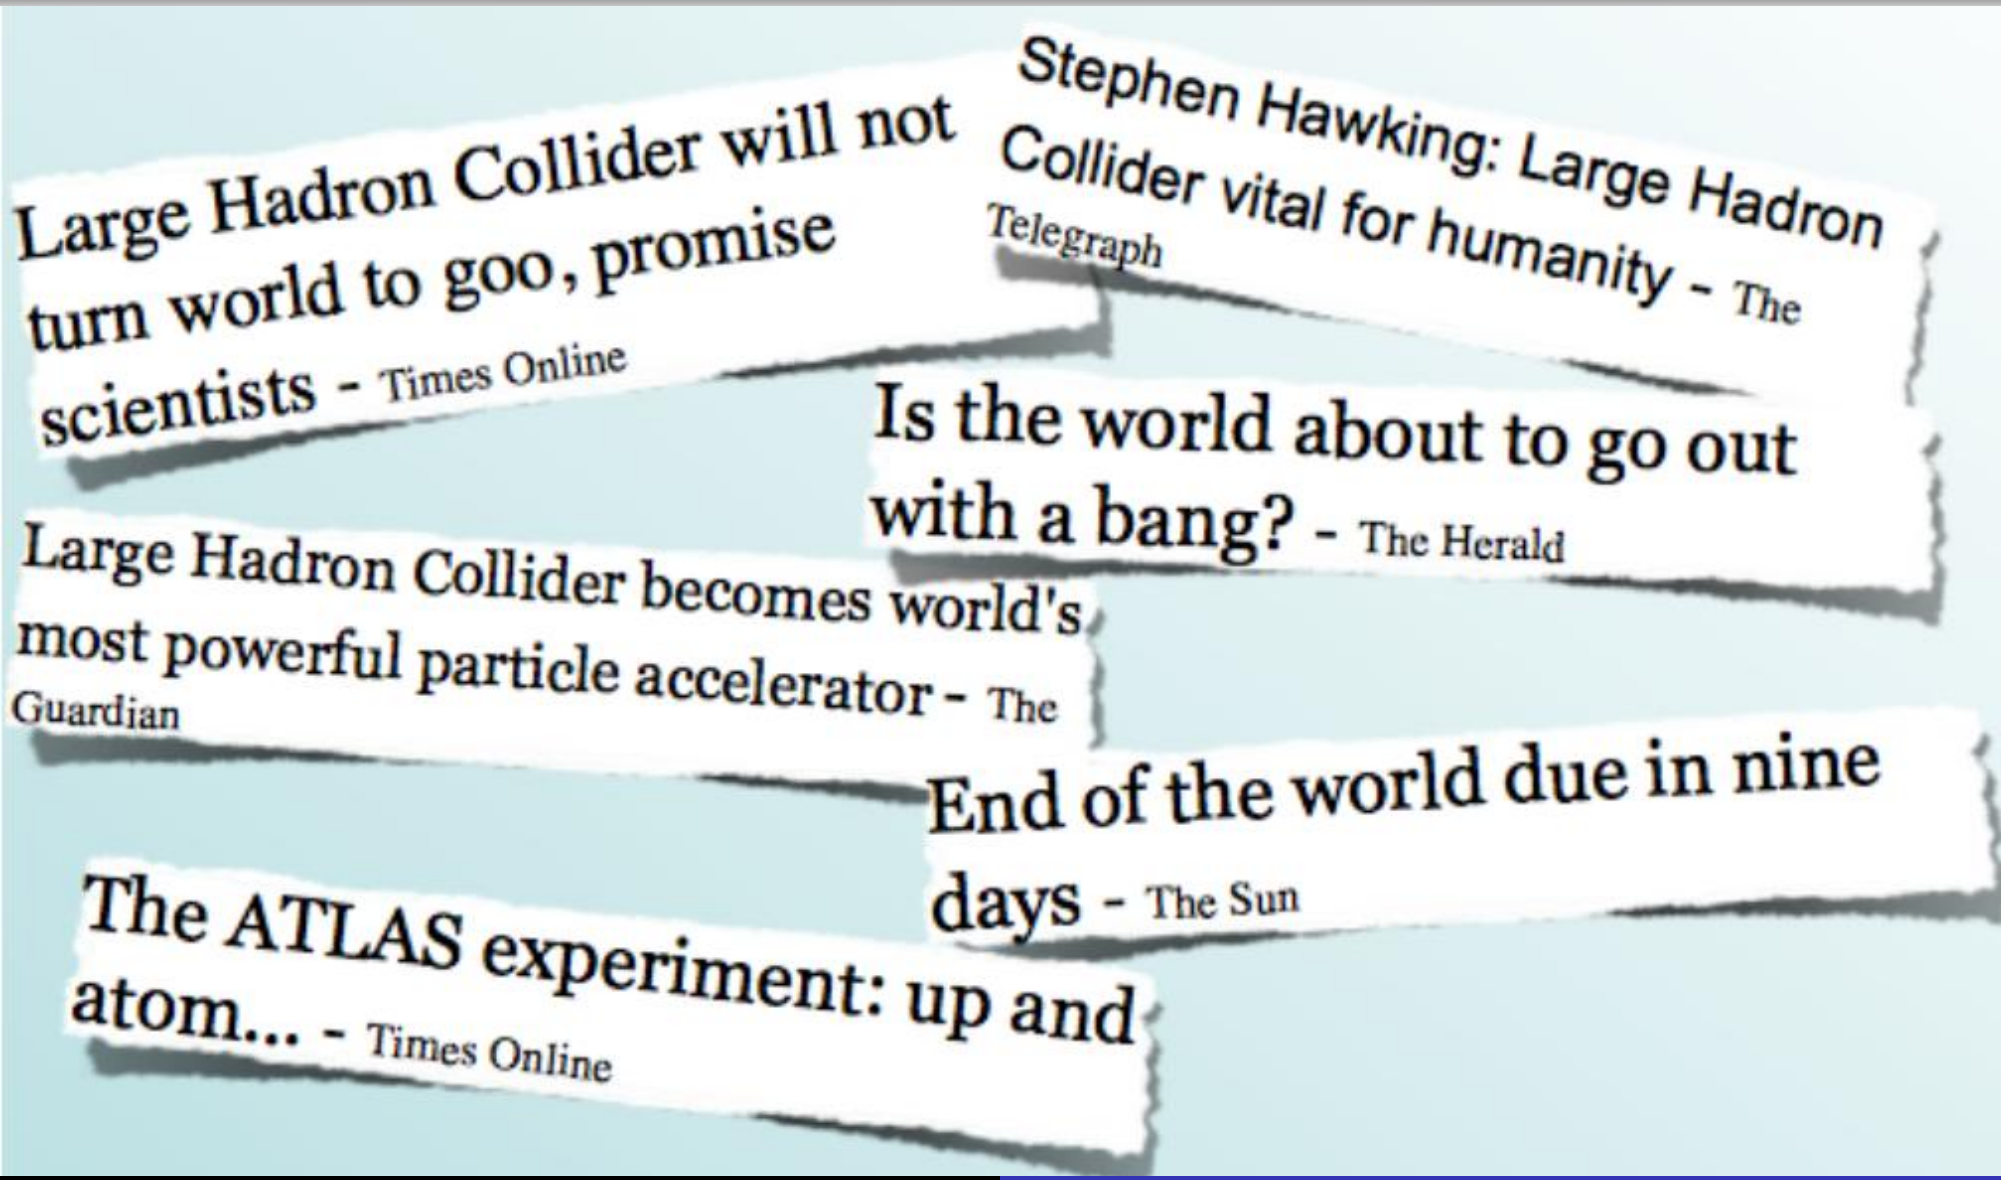

Global CERN project (∼10,000 scientists & engineers)

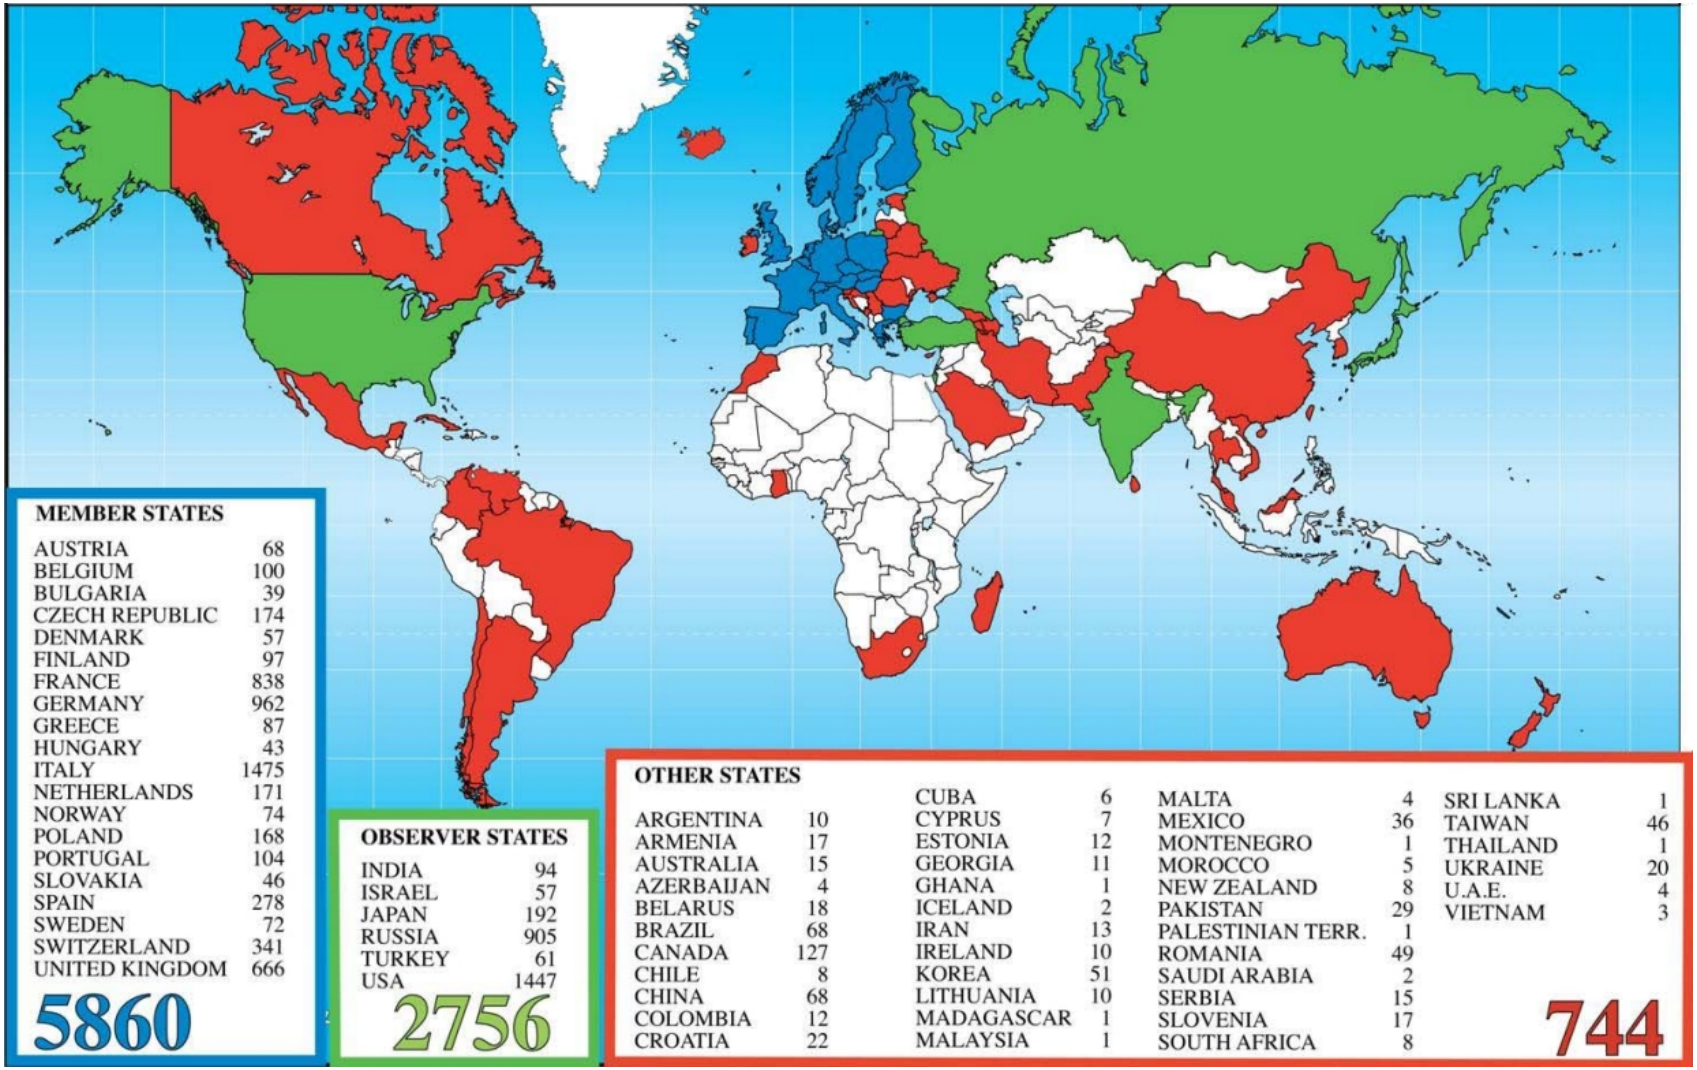

Outline

1. Introduction to the LHC and ATLAS
   - The Large Hadron Collider (LHC)
   - The ATLAS detector
   - The Higgs Boson
2. Areas of study for GPGPU adoption
   - Particle tracking in a magnetic field
   - The ATLAS trigger and data acquisition
   - The worldwide LHC computing grid
3. Using GPUs in the high level trigger
   - The Level 2 Trigger
   - The Z finder algorithm
   - The Kalman Filter

Two of the LHC collision points

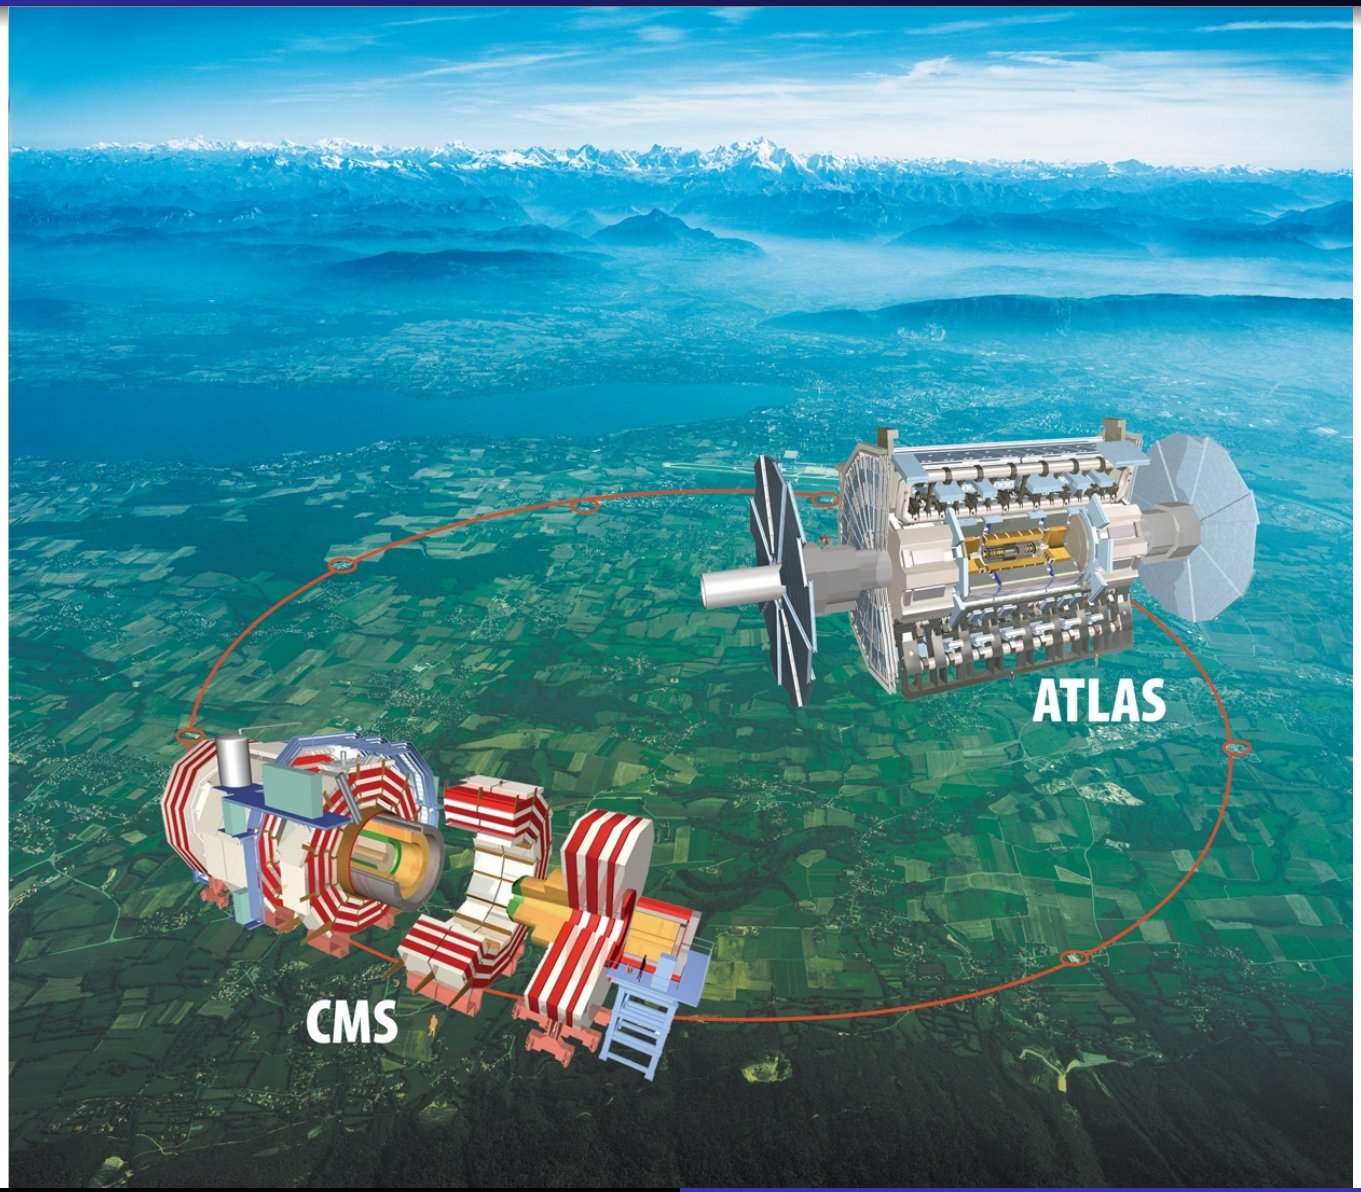

The ATLAS detector

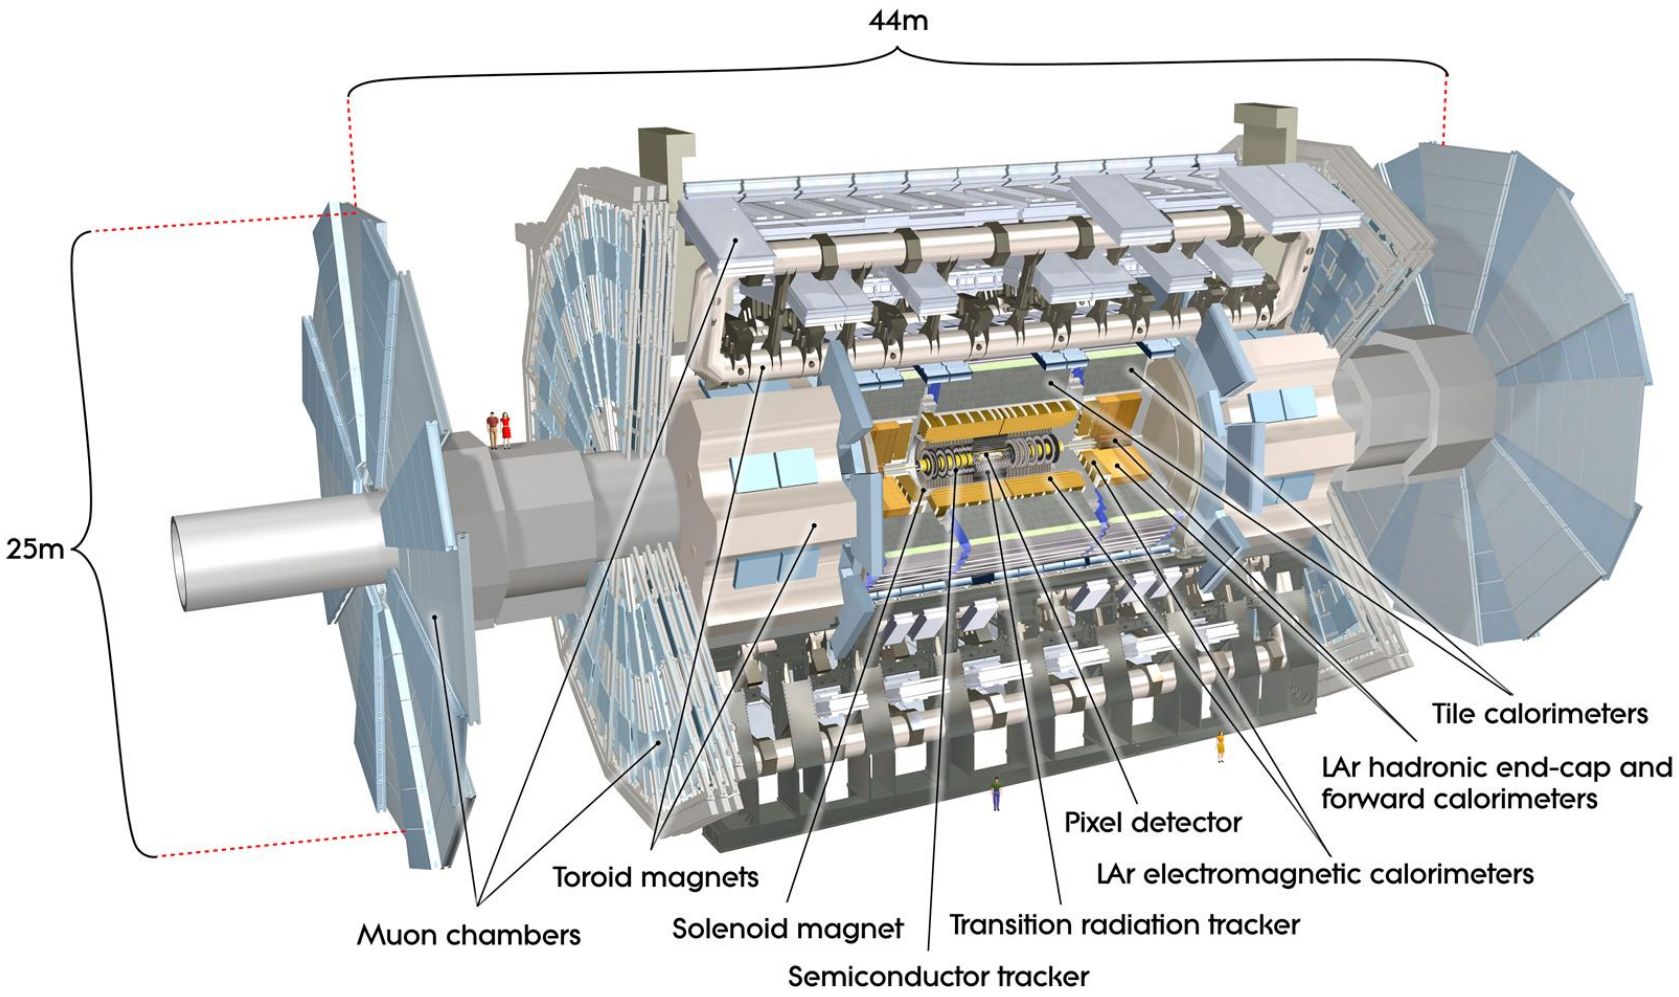

The ATLAS detector

The ATLAS detector

Physicist couple (10^{30} protons)

The ATLAS detector

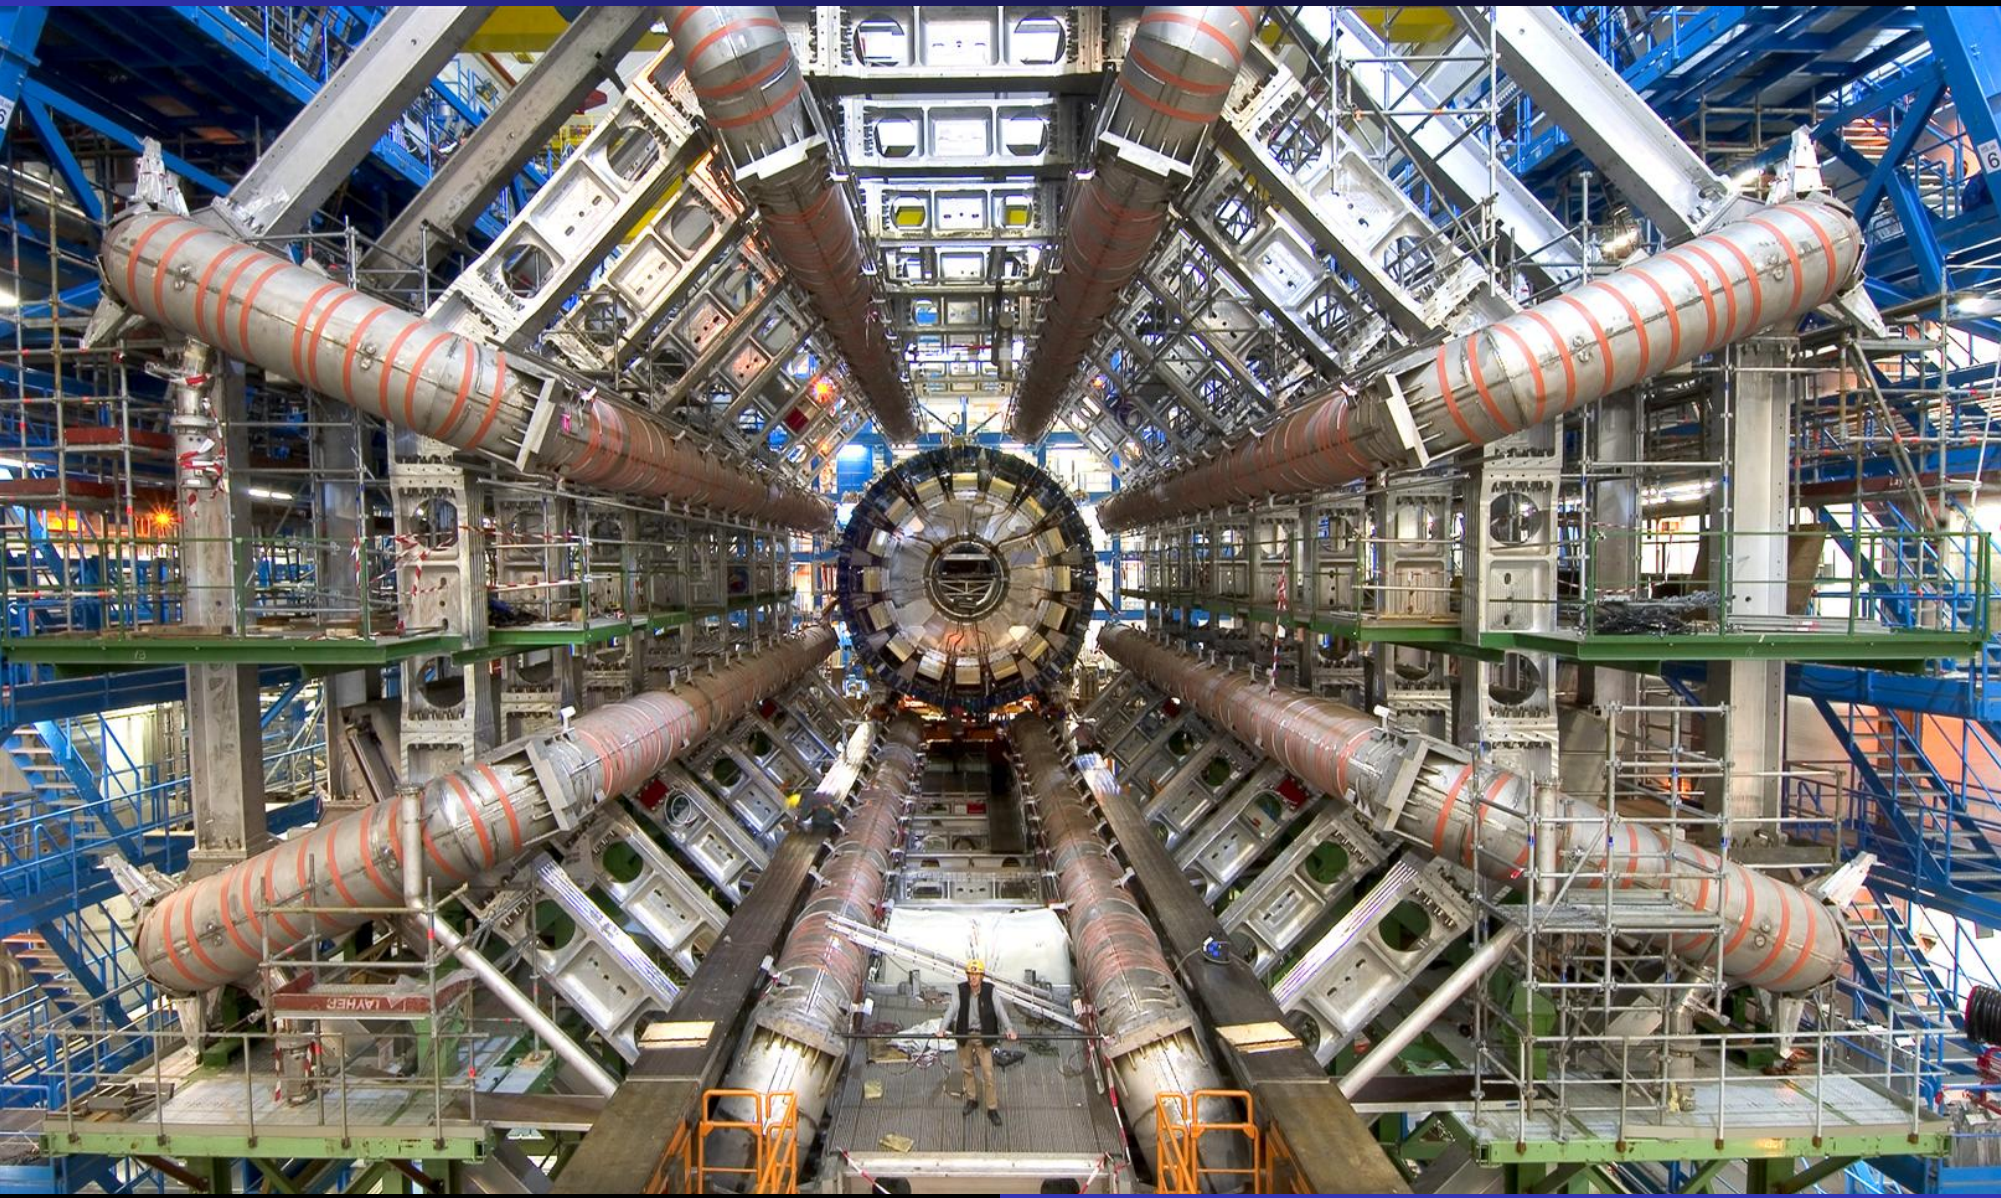

The ATLAS detector

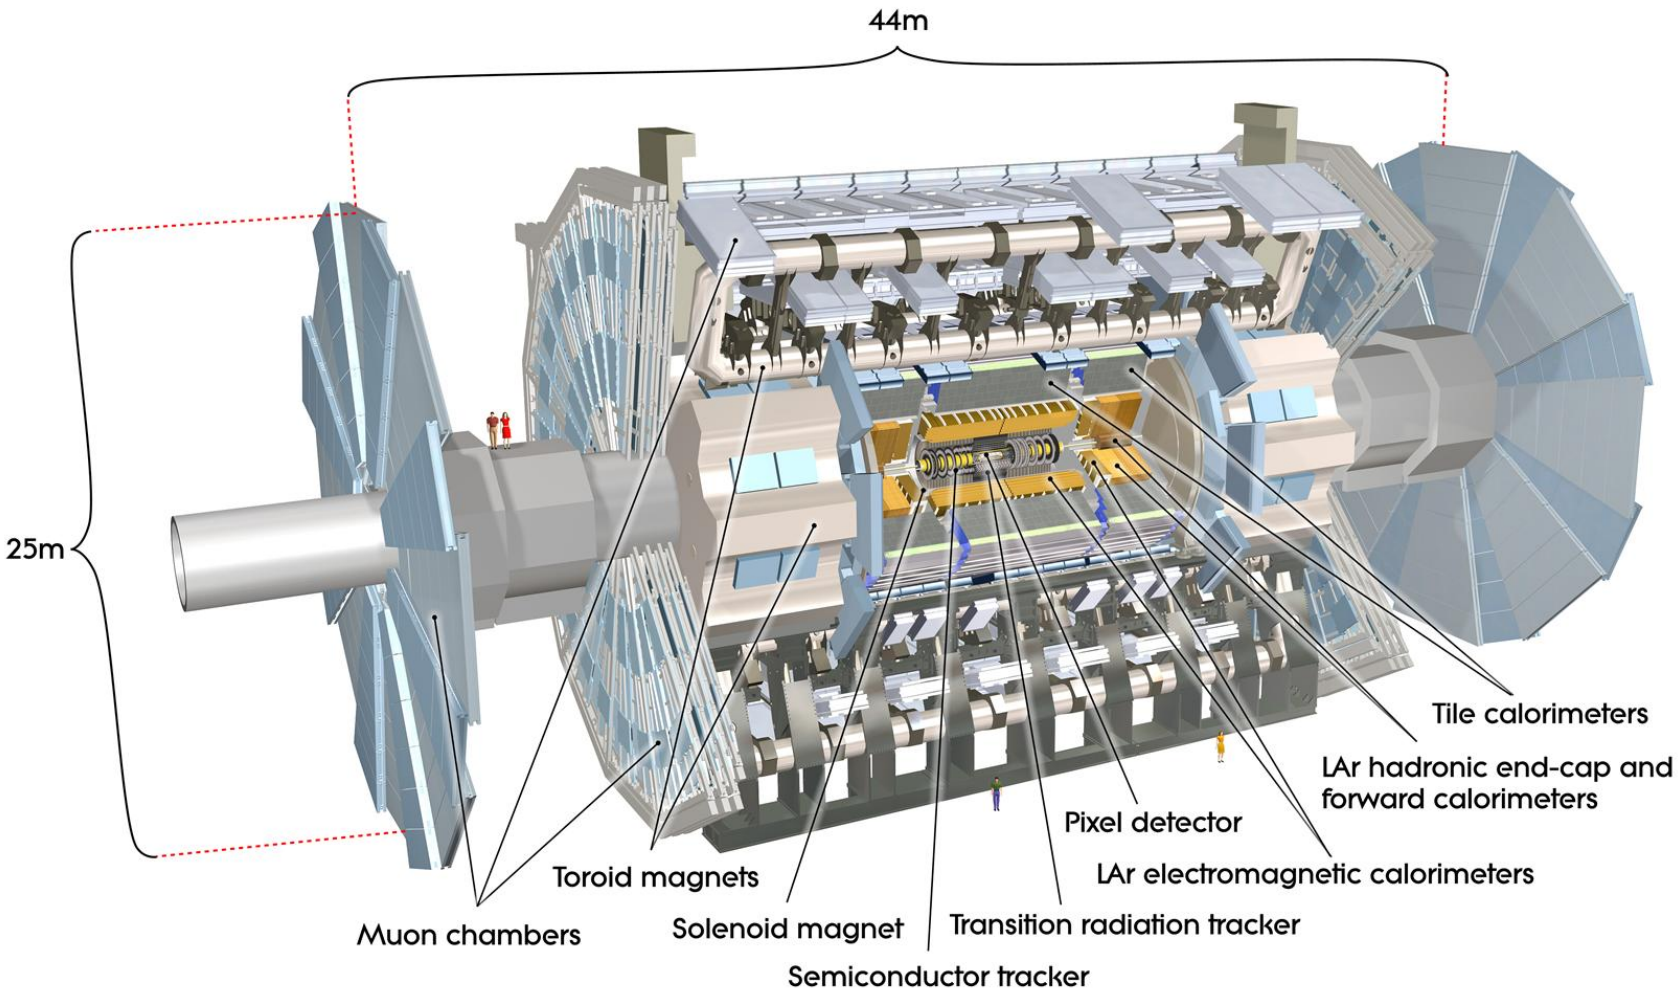

The ATLAS detector

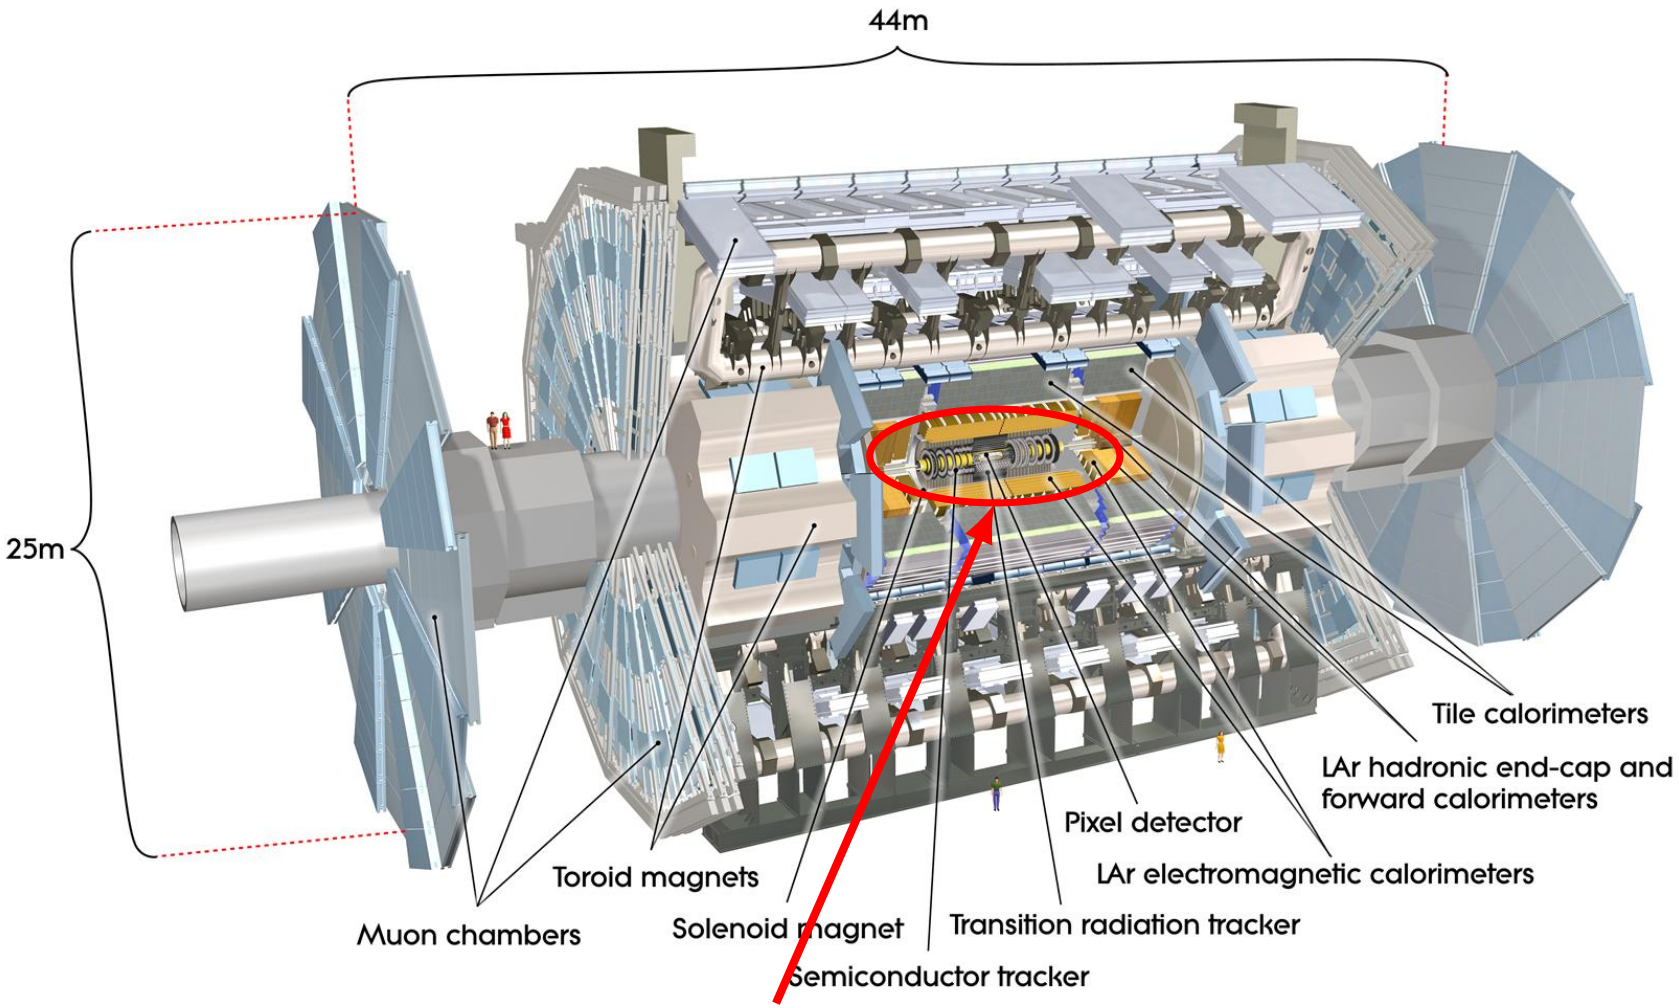

The ATLAS Inner Detector

Detector hits provide space points, used to track particles

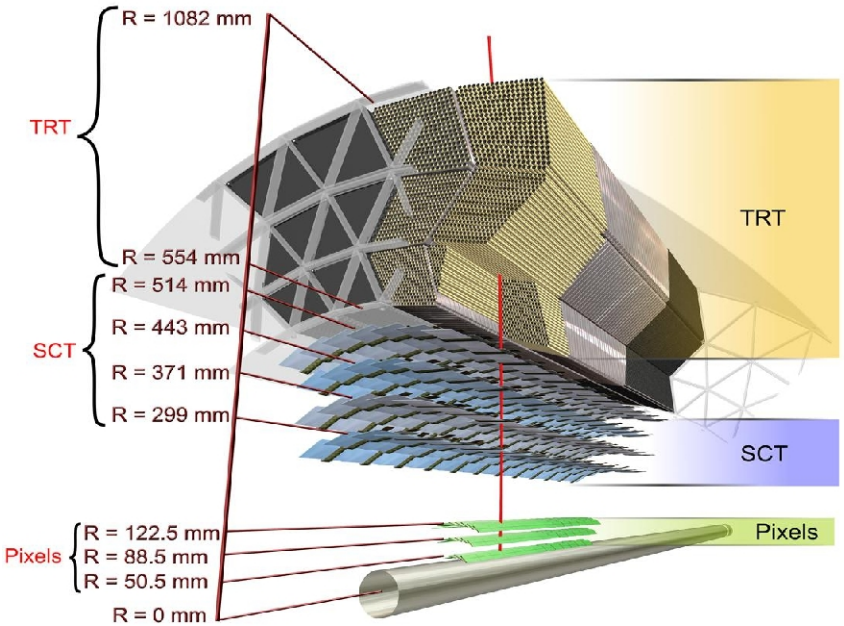

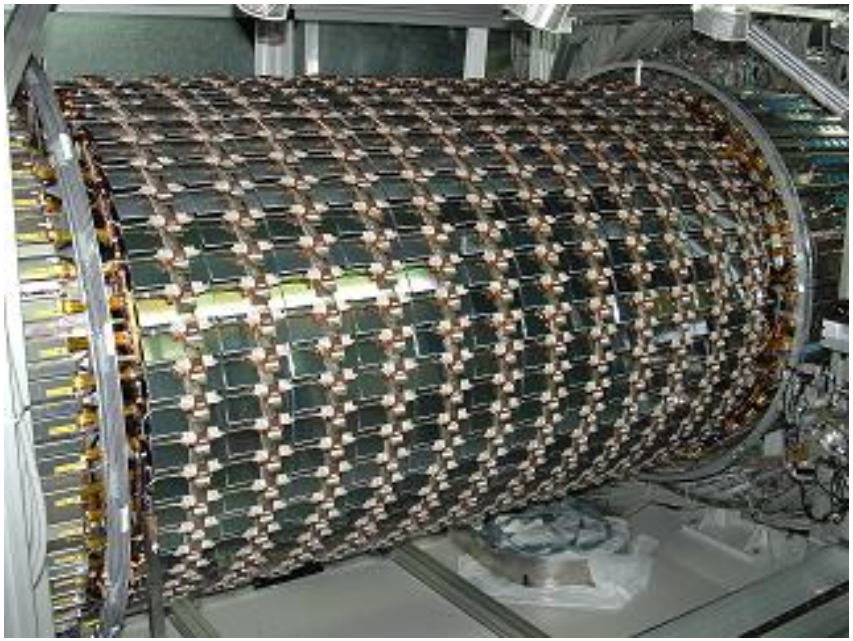

- Pixels pin-point the particle production vertex
- 80 million readout channels!
- SCT provides large area tracking: 61 m^2 silicon
- 6.3 million readout channels

A slice of the detector

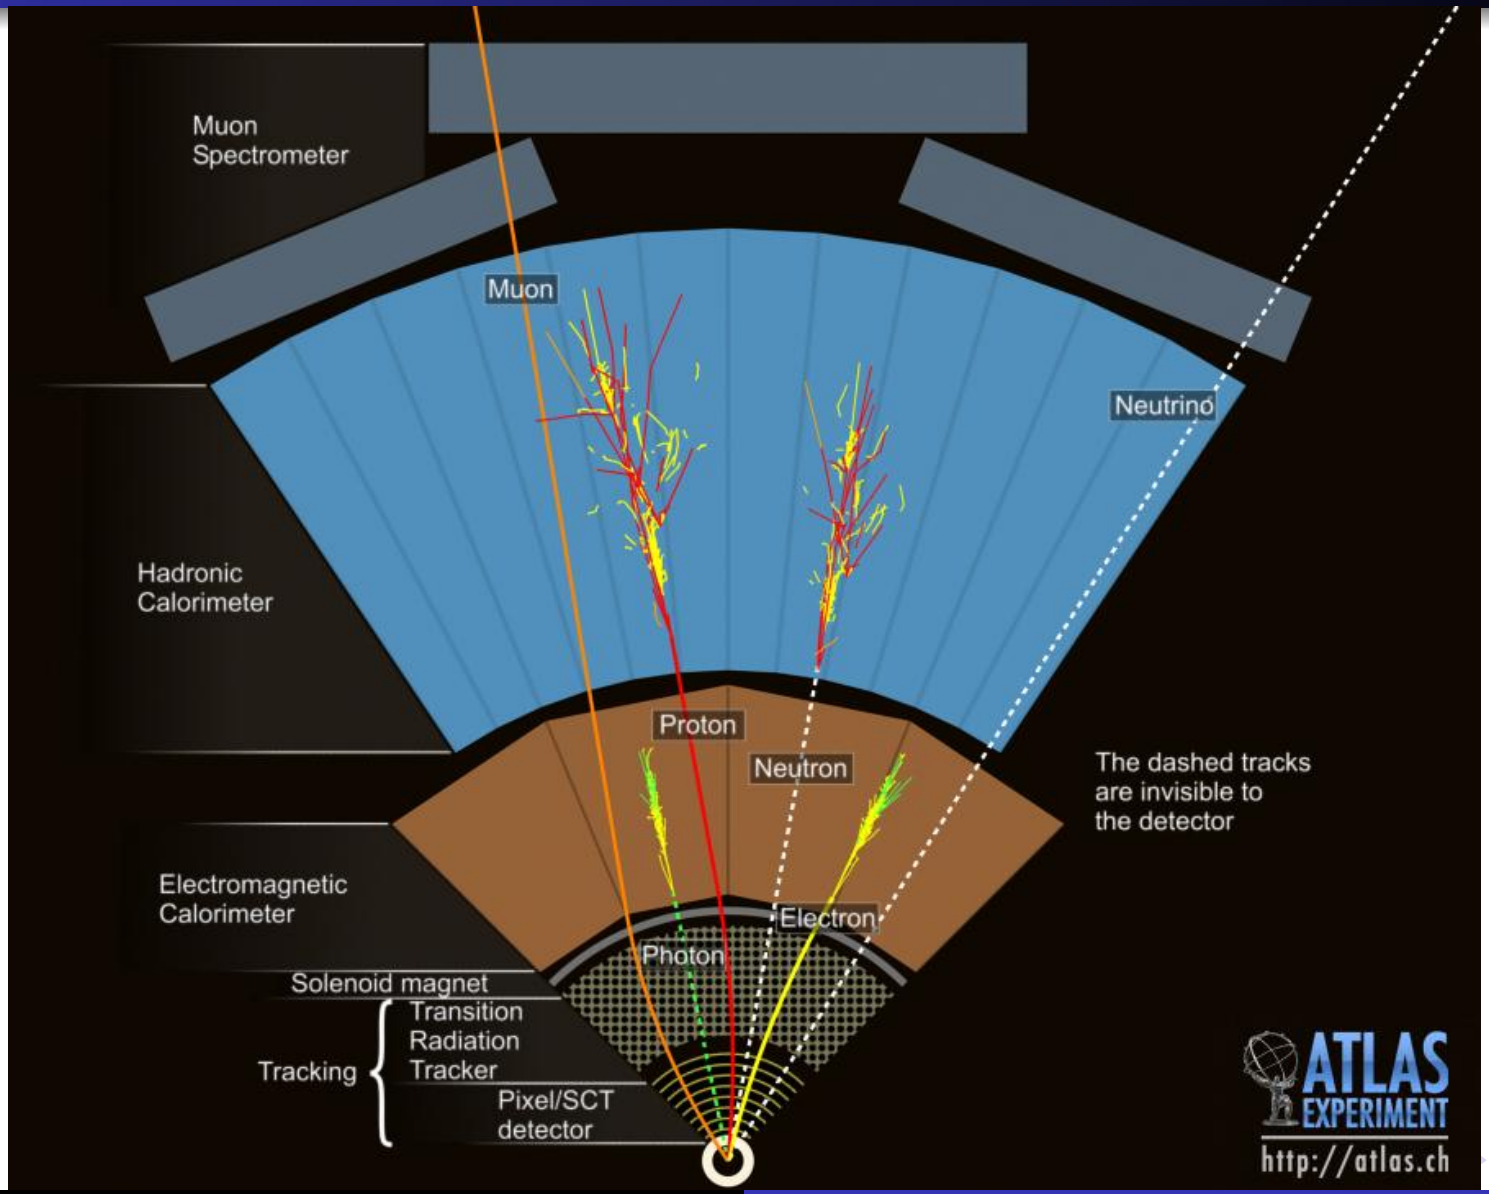

Outline

1. Introduction to the LHC and ATLAS
   - The Large Hadron Collider (LHC)
   - The ATLAS detector
   - The Higgs Boson
2. Areas of study for GPGPU adoption
   - Particle tracking in a magnetic field
   - The ATLAS trigger and data acquisition
   - The worldwide LHC computing grid
3. Using GPUs in the high level trigger
   - The Level 2 Trigger
   - The Z finder algorithm
   - The Kalman Filter

The Higgs Boson

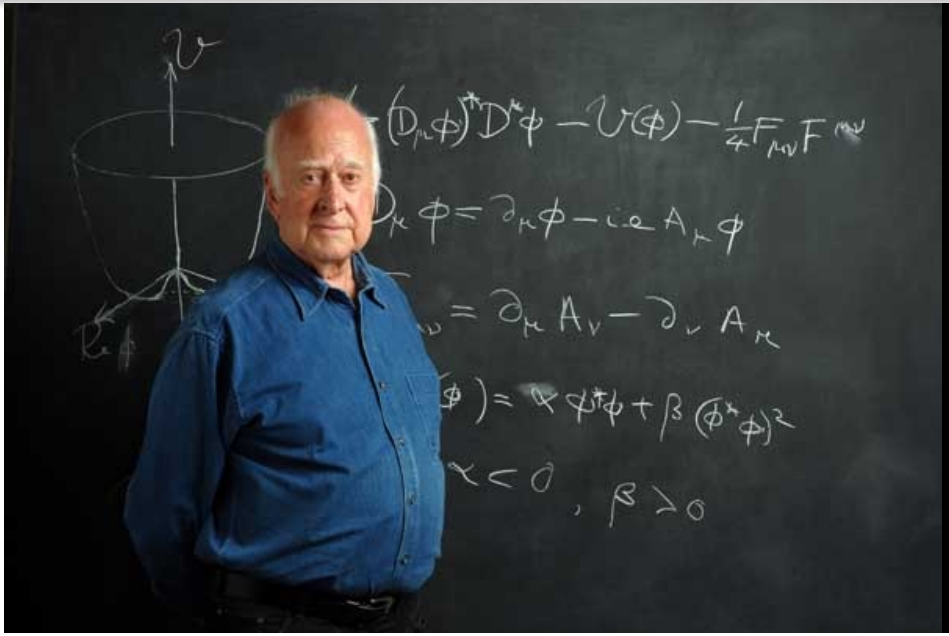

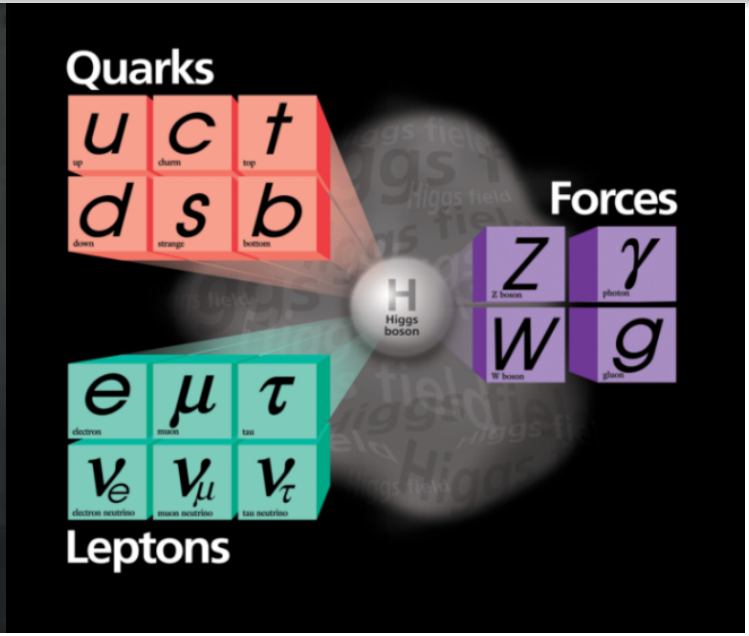

- Particle physics (Standard Model): extremely successful

The Higgs Boson

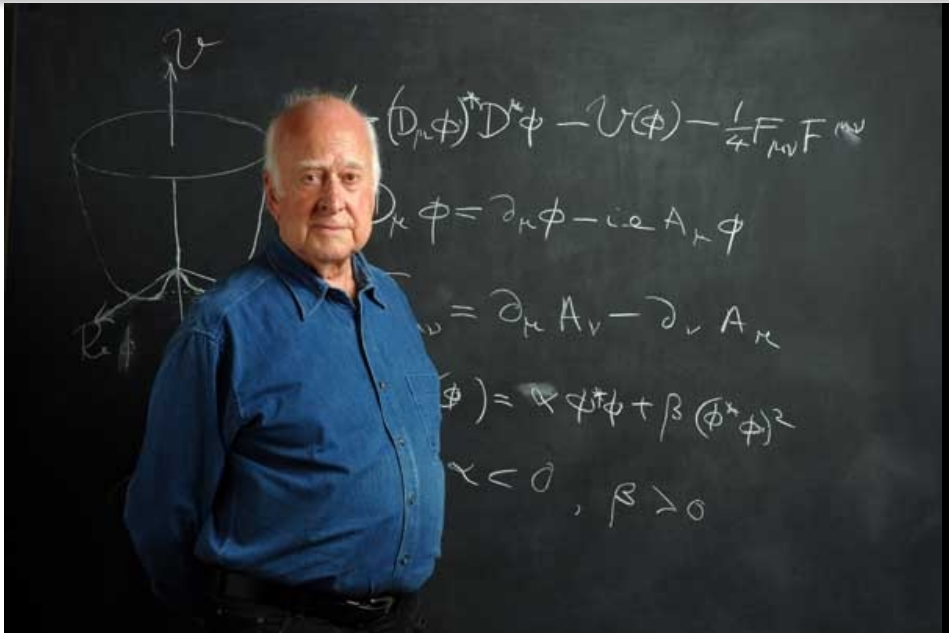

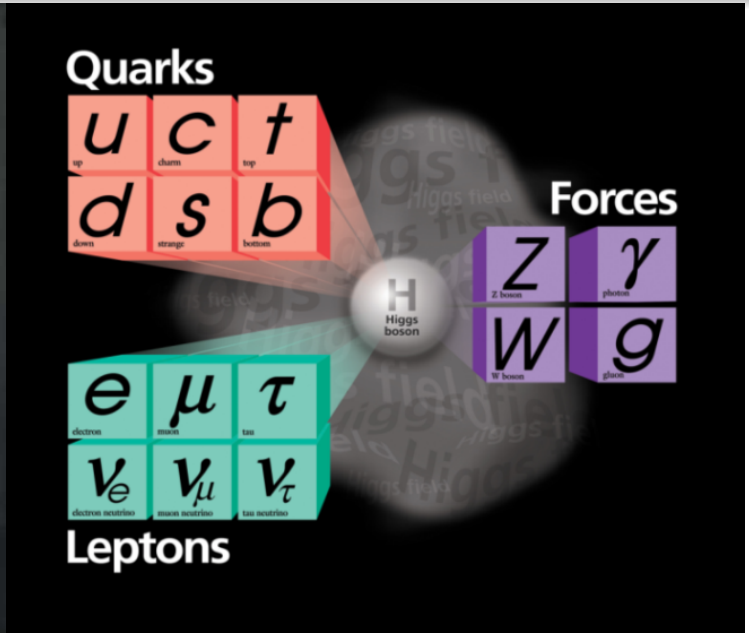

- Particle physics (Standard Model): extremely successful
  Highly predictive theory: has survived all experimental tests

The Higgs Boson

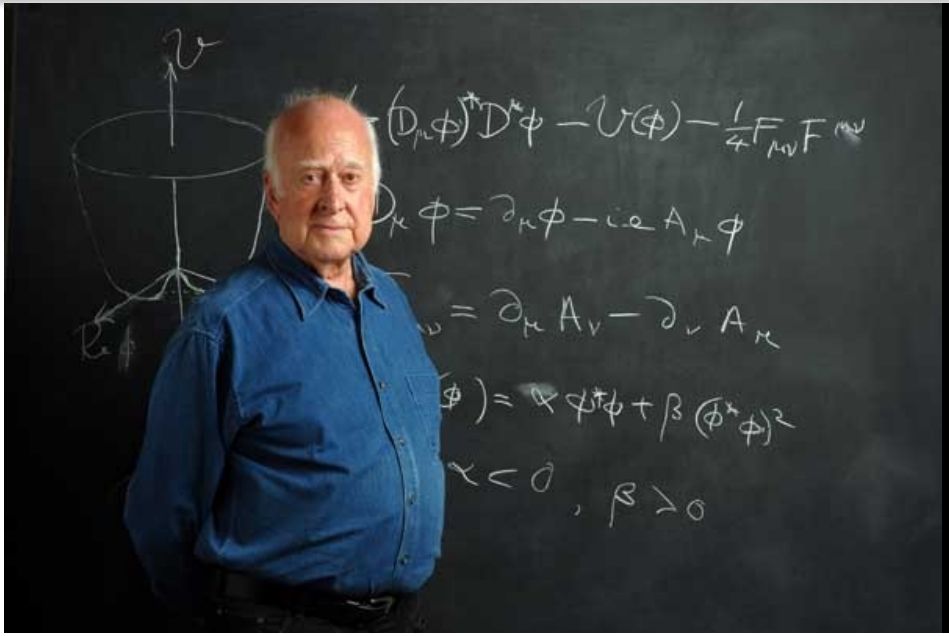

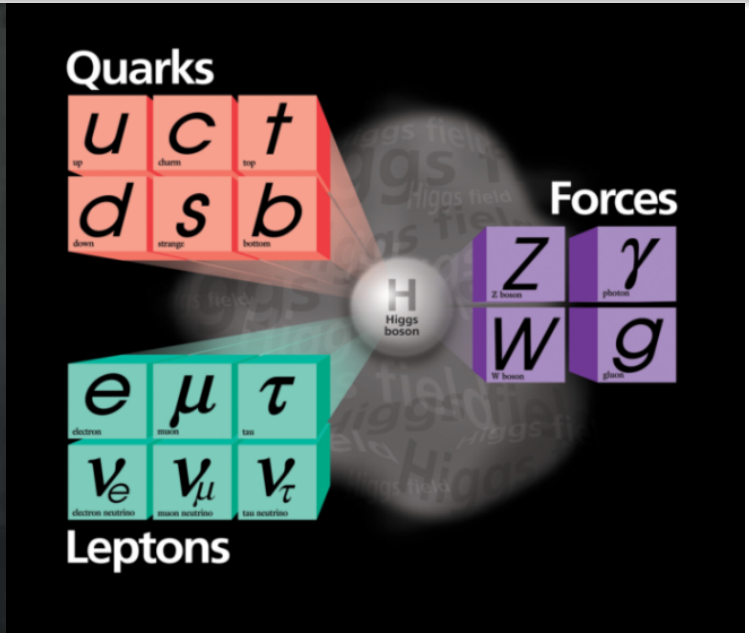

- Particle physics (Standard Model): extremely successful
  Highly predictive theory: has survived all experimental tests
- However, introducing elementary particle mass is difficult

The Higgs Boson

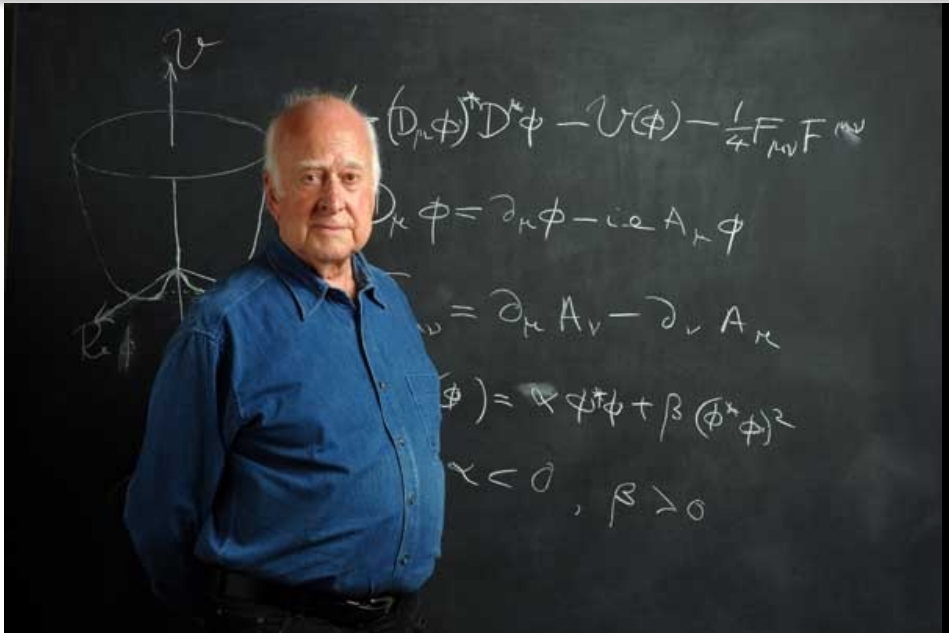

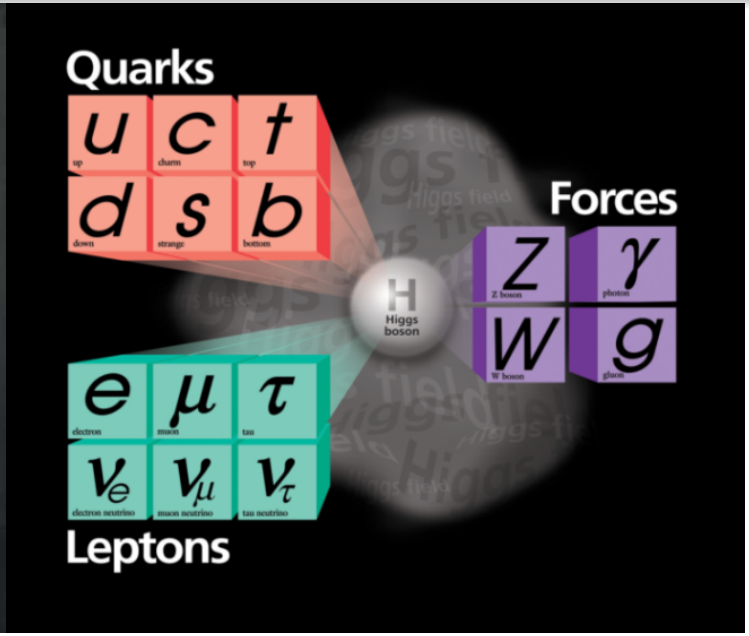

- Particle physics (Standard Model): extremely successful
  Highly predictive theory: has survived all experimental tests
- However, introducing elementary particle mass is difficult
- Peter Higgs invented a mechanism
  Gives mass to the leptons, quarks, W and Z force particles

The Higgs Boson

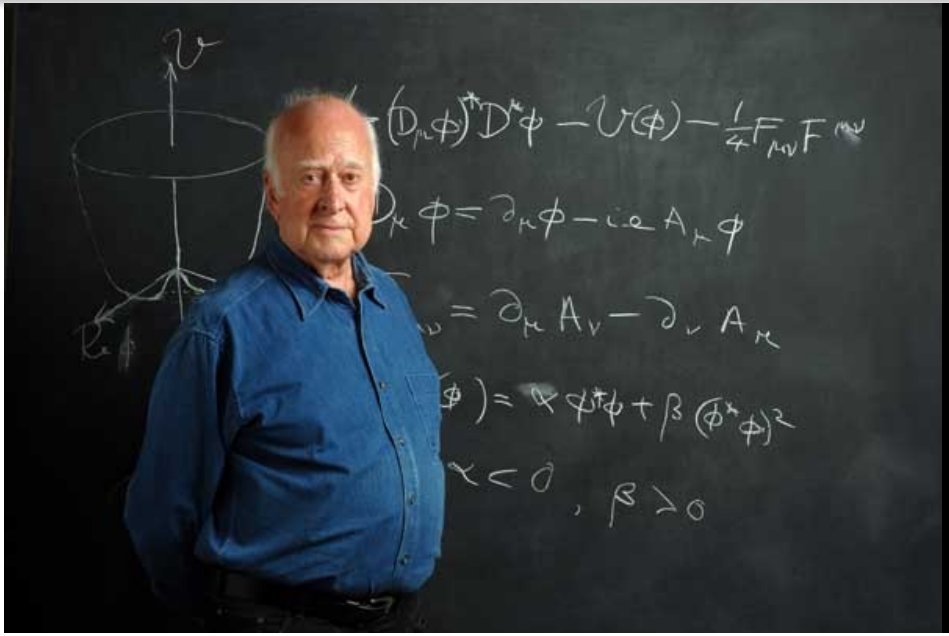

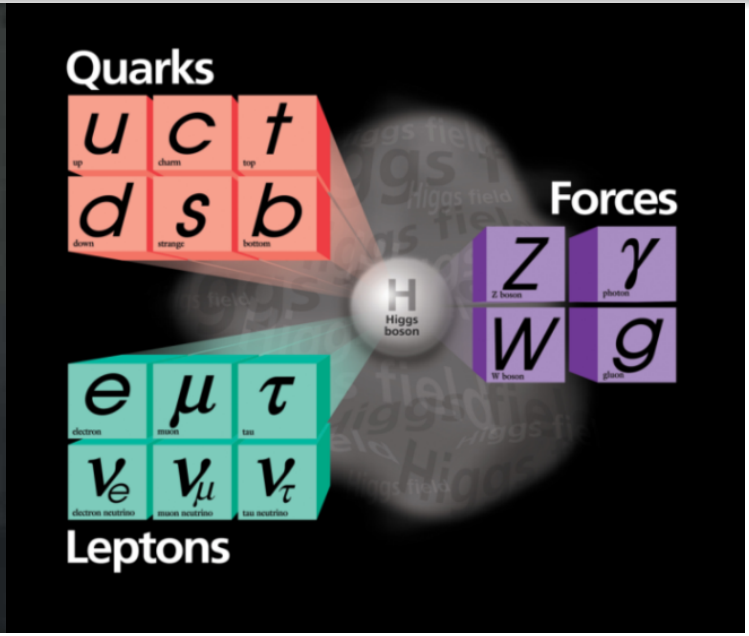

- Particle physics (Standard Model): extremely successful
  Highly predictive theory: has survived all experimental tests
- However, introducing elementary particle mass is difficult
- Peter Higgs invented a mechanism
  Gives mass to the leptons, quarks, W and Z force particles
- Requires a new type of particle to exist: the Higgs boson

Hunting for the Higgs boson

Most fundamental questions in nature

- Why do particles (and thus matter) have mass?

Hunting for the Higgs boson

Most fundamental questions in nature

- Why do particles (and thus matter) have mass?
- Why such different masses?

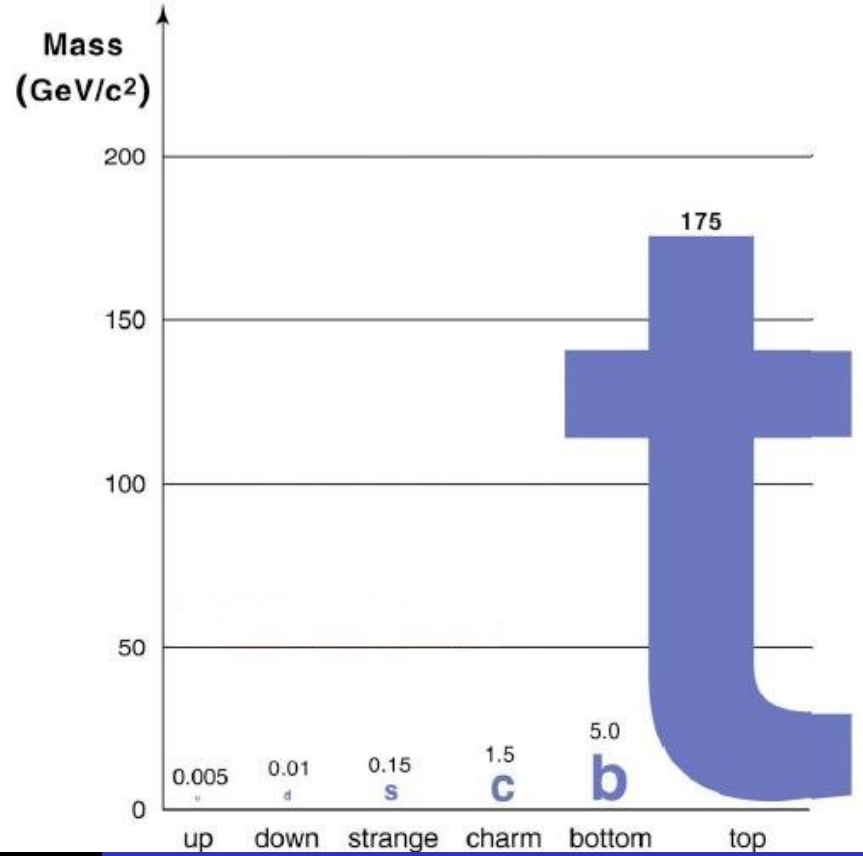

Hunting for the Higgs boson

Most fundamental questions in nature

- Why do particles (and thus matter) have mass?
- Why such different masses?

The search for the Higgs

- Physicists have searched for decades, but it has not yet been found.

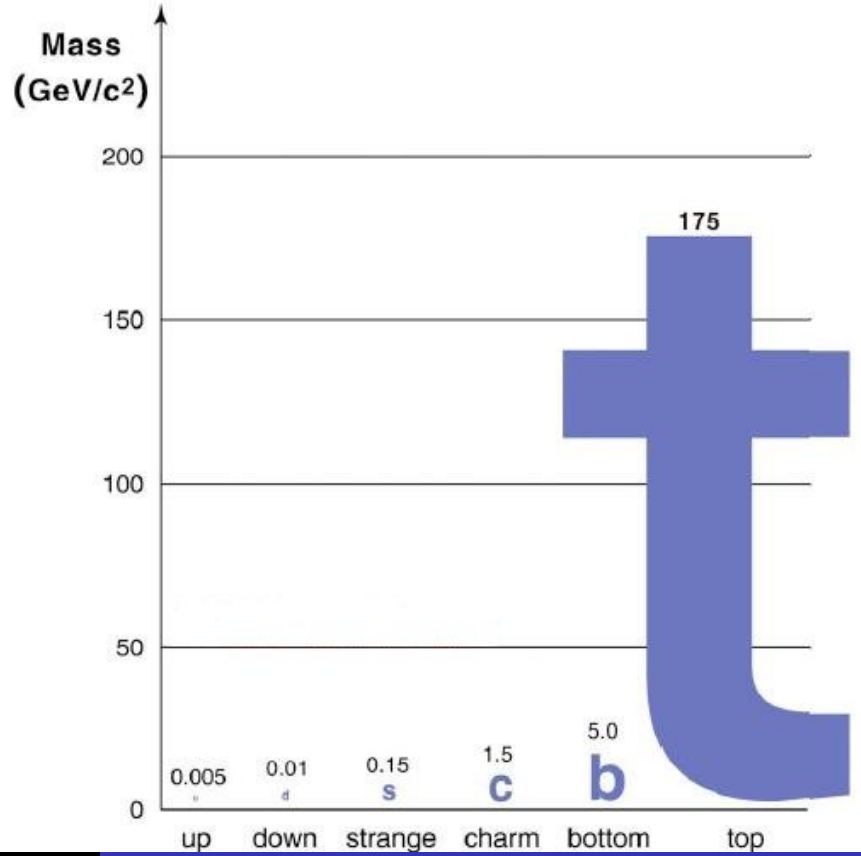

Hunting for the Higgs boson

Most fundamental questions in nature

- Why do particles (and thus matter) have mass?
- Why such different masses?

The search for the Higgs

- Physicists have searched for decades, but it has not yet been found.
- The LHC will have sufficient energy to produce it, if it exists.

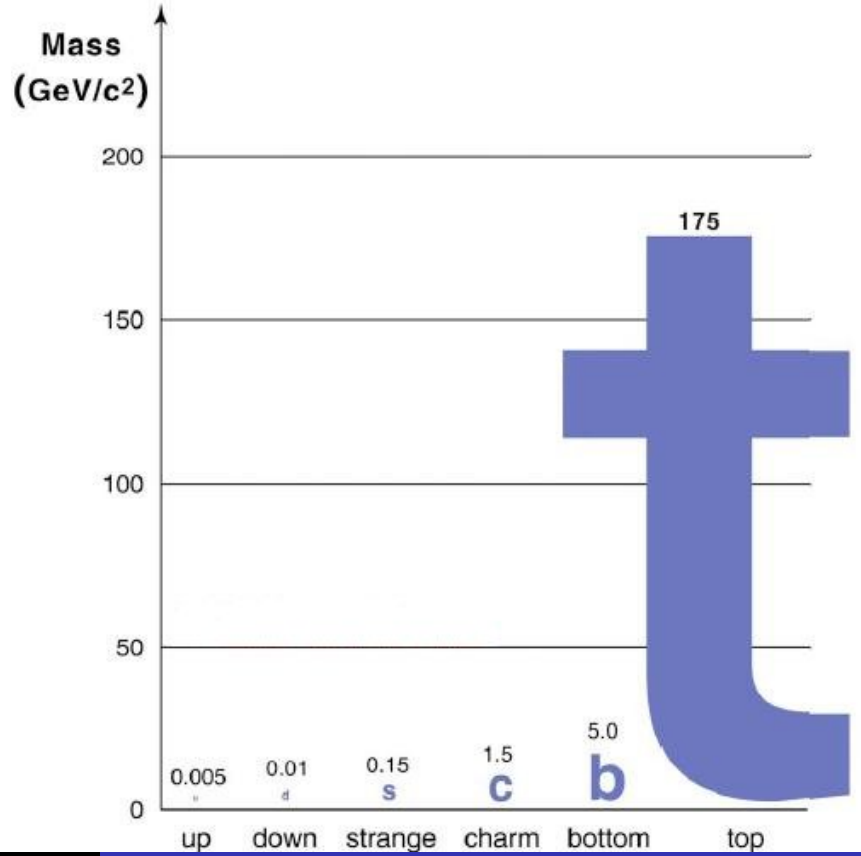

Hunting for the Higgs boson

Most fundamental questions in nature

- Why do particles (and thus matter) have mass?
- Why such different masses?

The search for the Higgs

- Physicists have searched for decades, but it has not yet been found.
- The LHC will have sufficient energy to produce it, if it exists.
- Conclusively

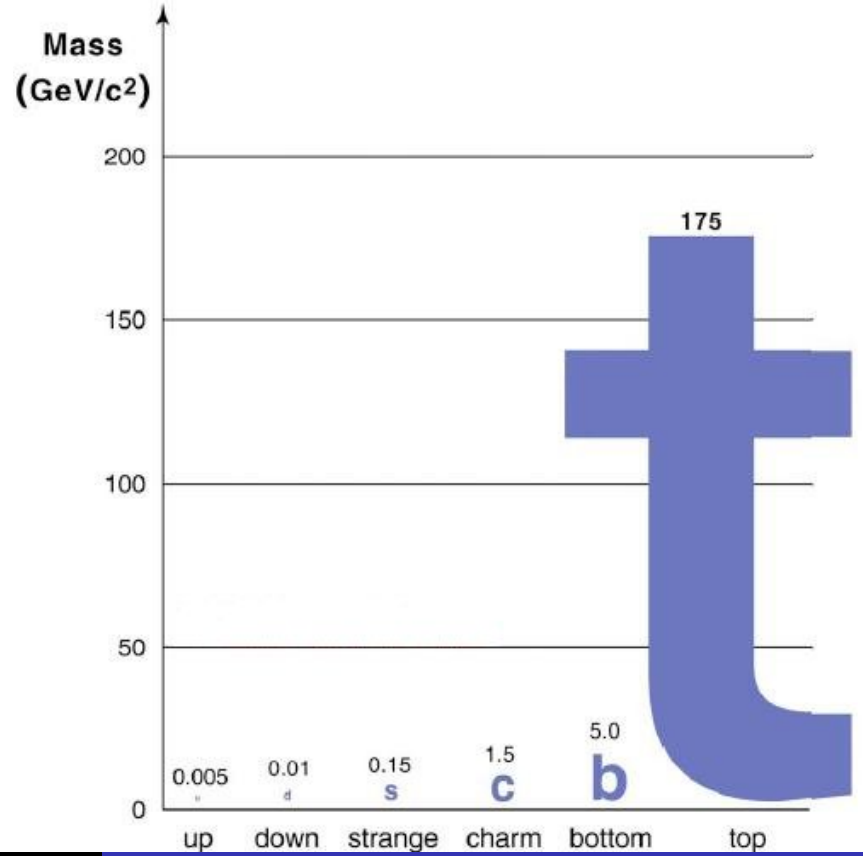

LHC collision process

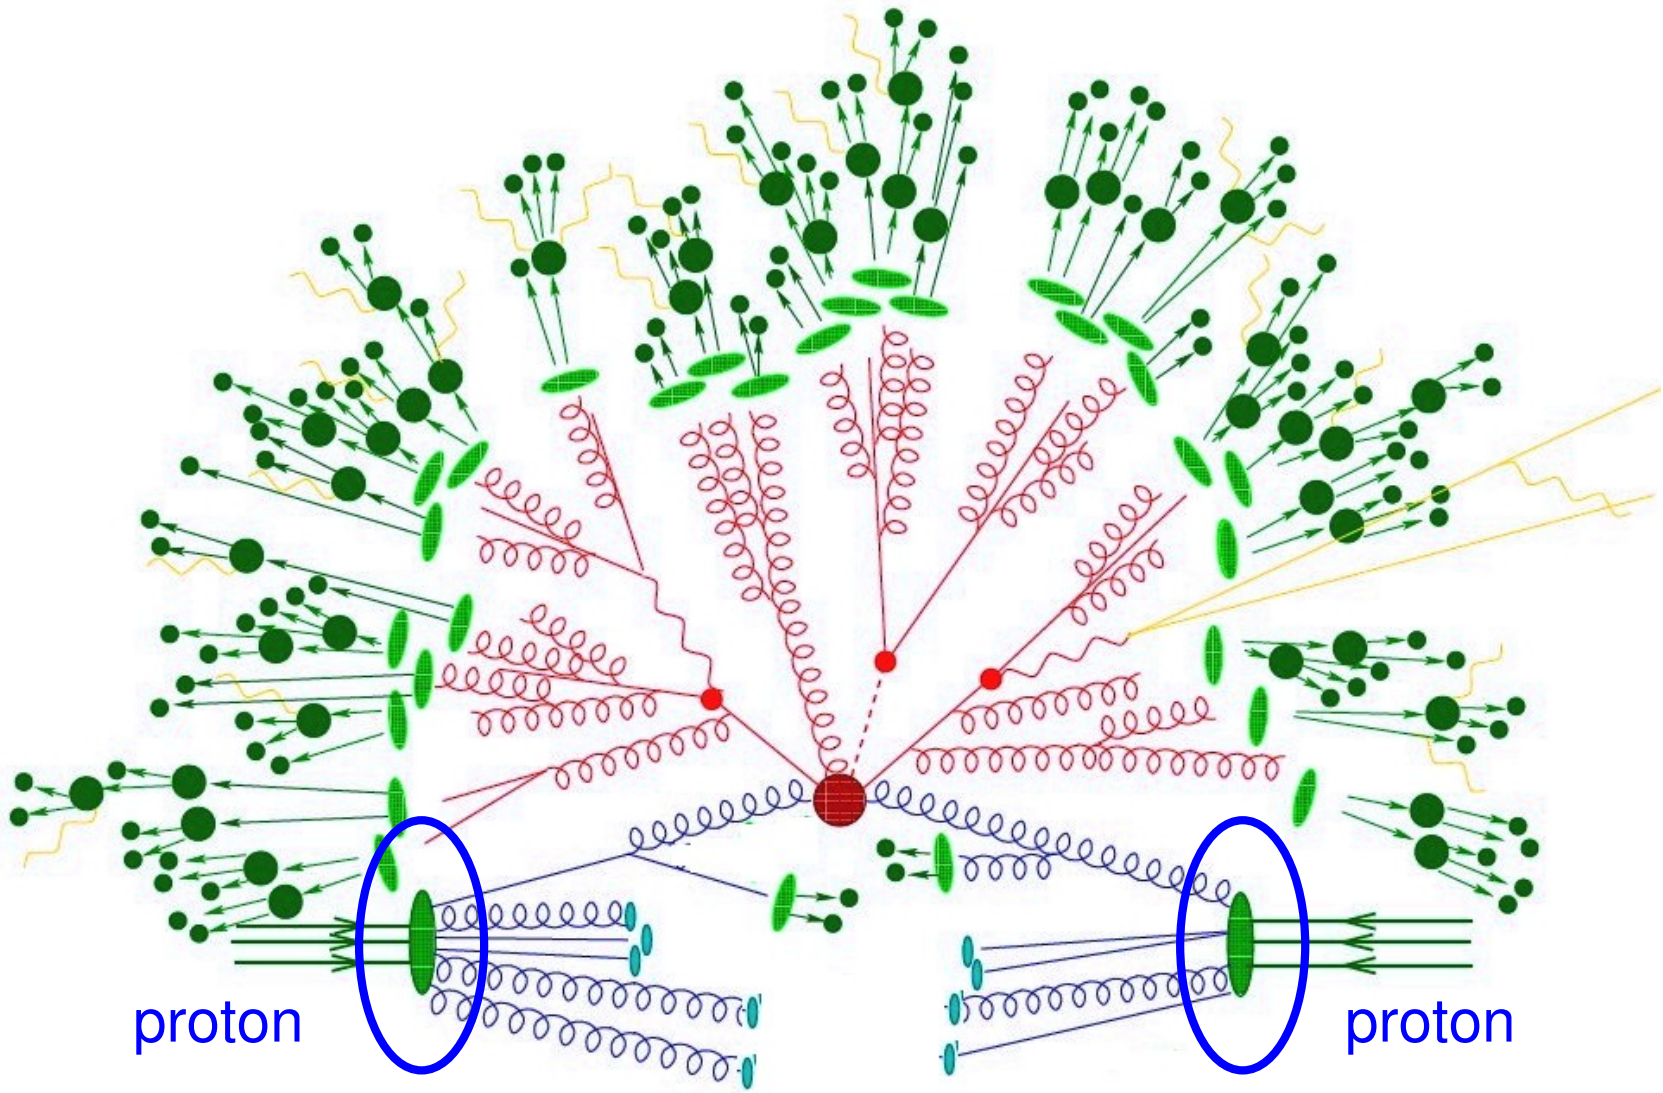

LHC collision process

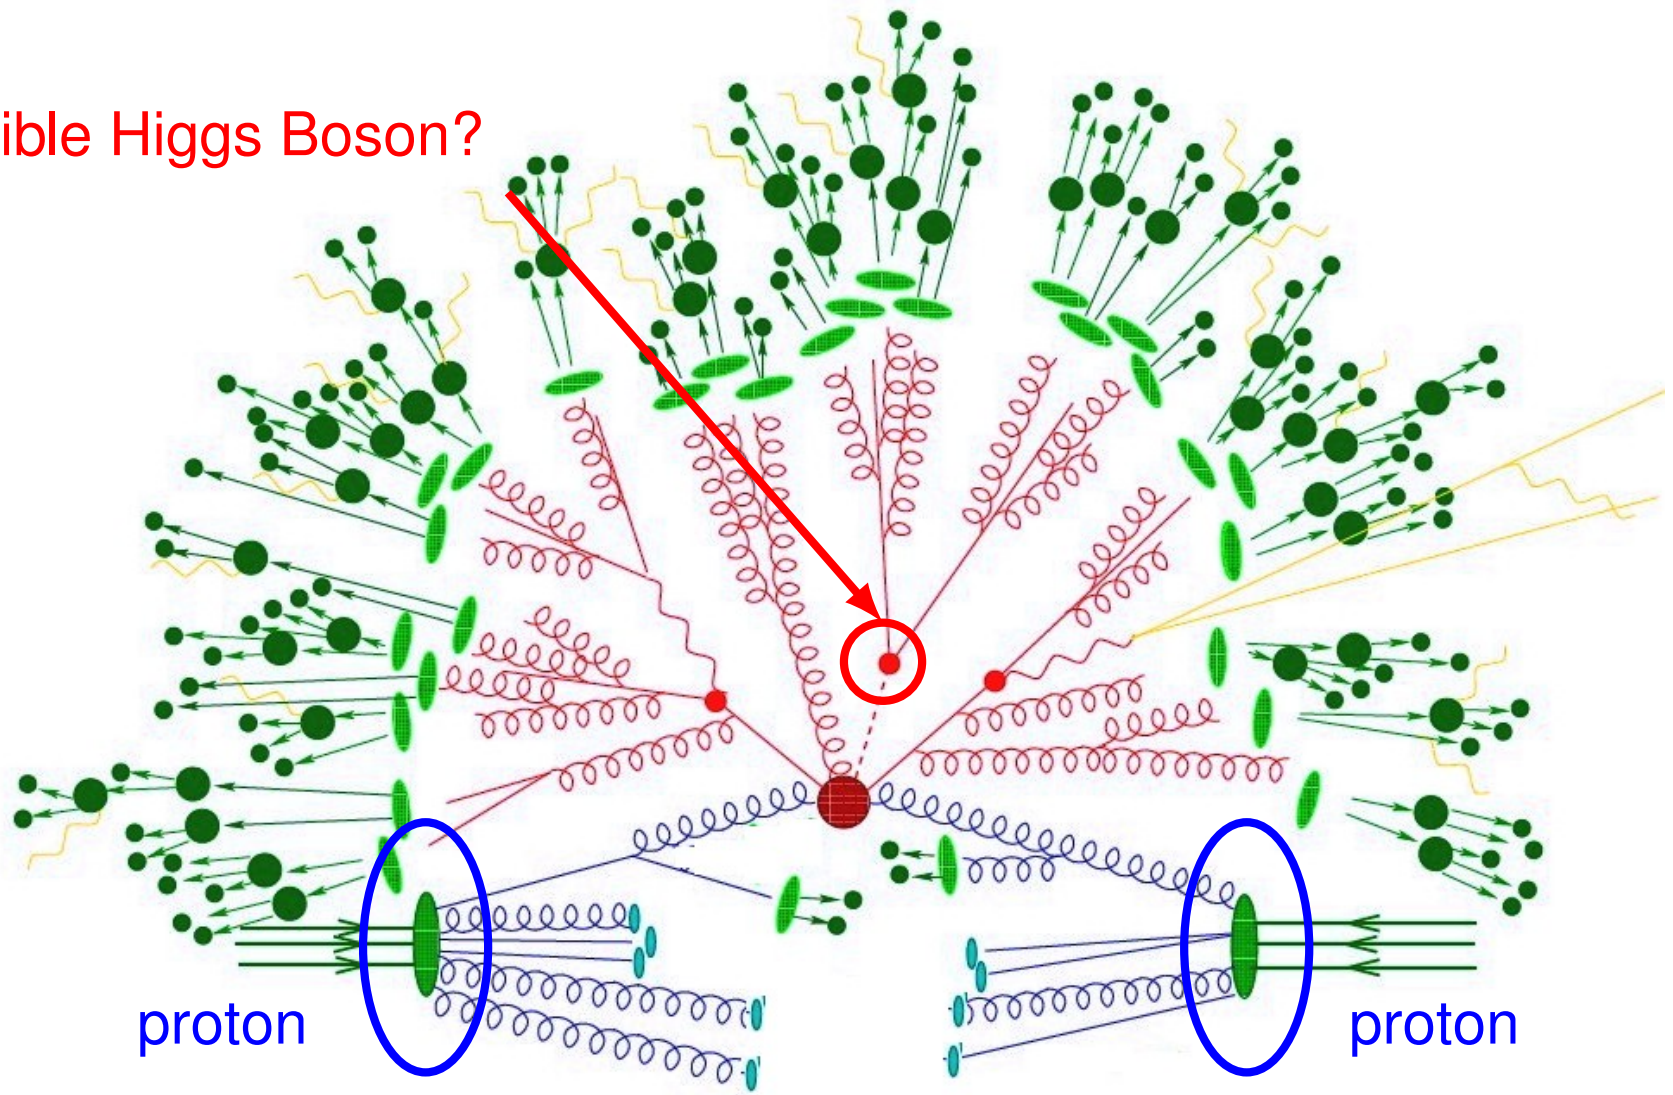

A simulated Higgs boson event

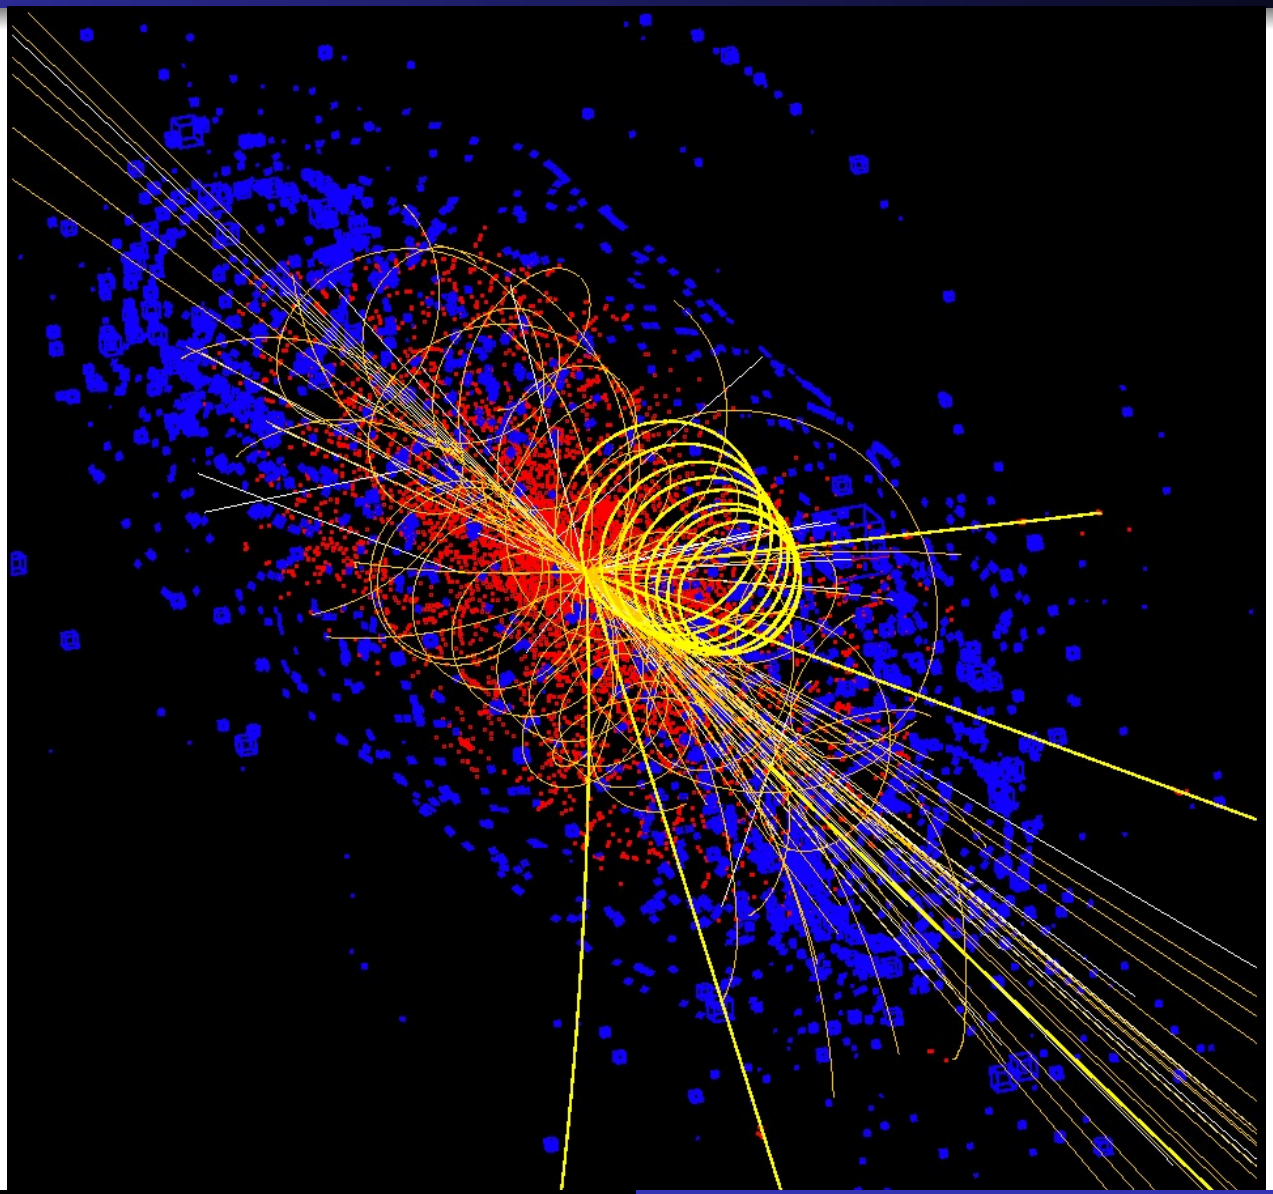

Outline

1. Introduction to the LHC and ATLAS
   - The Large Hadron Collider (LHC)
   - The ATLAS detector
   - The Higgs Boson
2. Areas of study for GPGPU adoption
   - Particle tracking in a magnetic field
   - The ATLAS trigger and data acquisition
   - The worldwide LHC computing grid
3. Using GPUs in the high level trigger
   - The Level 2 Trigger
   - The Z finder algorithm
   - The Kalman Filter

Particle tracking in a magnetic field

Preliminary GPGPU test case study

Particle tracking in a magnetic field

Preliminary GPGPU test case study

- Charged particles bend in the magnetic field

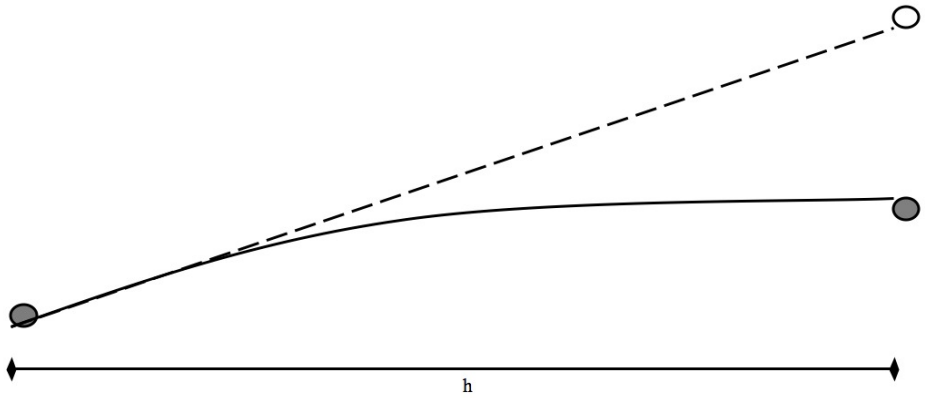

Particle tracking in a magnetic field

Preliminary GPGPU test case study

- Charged particles bend in the magnetic field
- Lorentz force (perpendicular to plane of magnetic field)

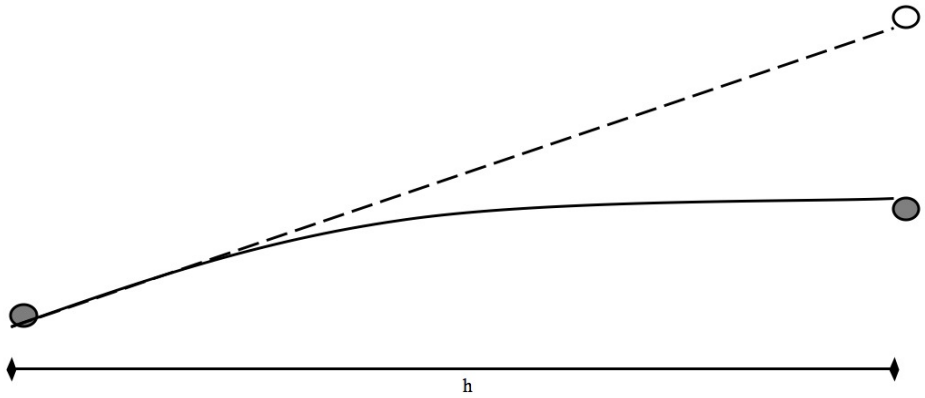

$$\mathbf{F} = m\mathbf{a} = q \cdot (\mathbf{E} + \mathbf{v} \times \mathbf{B})$$

$$\frac{d\mathbf{v}}{dt} = \mathbf{a} = \frac{q}{m} \cdot (\mathbf{E} + \mathbf{v} \times \mathbf{B})$$

Particle tracking in a magnetic field

Preliminary GPGPU test case study

- Charged particles bend in the magnetic field
- Lorentz force (perpendicular to plane of magnetic field)

$$\mathbf{F} = m\mathbf{a} = q \cdot (\mathbf{E} + \mathbf{v} \times \mathbf{B})$$

$$\frac{d\mathbf{v}}{dt} = \mathbf{a} = \frac{q}{m} \cdot (\mathbf{E} + \mathbf{v} \times \mathbf{B})$$

- Solve the differential equation with 4th order Runga Kutta Integration

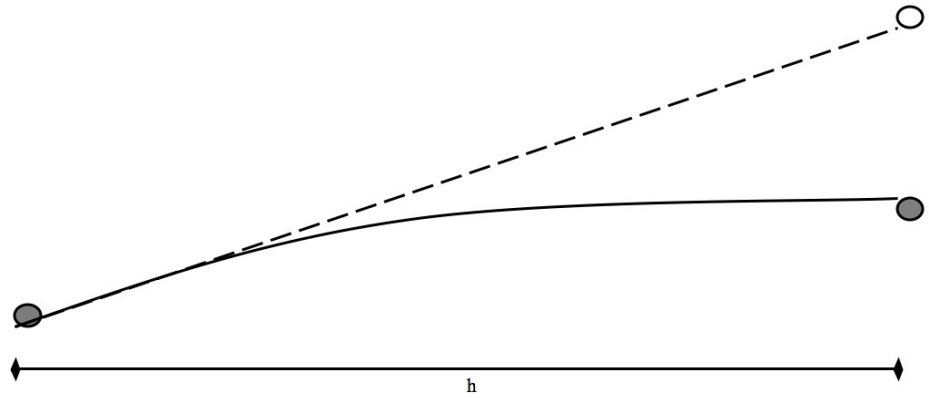

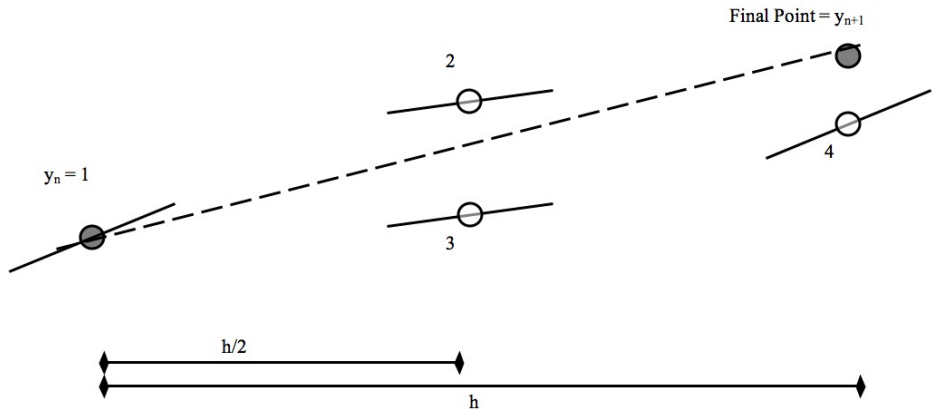

Acceleration with GPGPUs

1. Using the GPGPU, pre-calculated a “look-up” table of derivative calculations for a space point matrix
   - Calculation time not a limiting factor (abandoned this idea)
   - Also lost accuracy due to rounding to nearest look up point

Acceleration with GPGPUs

1. Using the GPGPU, pre-calculated a "look-up" table of derivative calculations for a space point matrix
   - Calculation time not a limiting factor (abandoned this idea)
   - Also lost accuracy due to rounding to nearest look up point
2. Increased calculation complexity to use adaptive stepping
   - Adjusting step size to be within an error tolerance
   - Still slower than the CPU...

Acceleration with GPGPUs

1. Using the GPGPU, pre-calculated a "look-up" table of derivative calculations for a space point matrix
   - Calculation time not a limiting factor (abandoned this idea)
   - Also lost accuracy due to rounding to nearest look up point
2. Increased calculation complexity to use adaptive stepping
   - Adjusting step size to be within an error tolerance
   - Still slower than the CPU...
3. Treated x,y,z coordinates in parallel (3 threads in block)
   - Cross-product ($\mathbf{v} \times \mathbf{B}$) calculation needs perp. coordinates
   - Set up the threads in the block to use shared memory
   - Speed was now closer to CPU

Acceleration with GPGPUs

1. Using the GPGPU, pre-calculated a "look-up" table of derivative calculations for a space point matrix
   - Calculation time not a limiting factor (abandoned this idea)
   - Also lost accuracy due to rounding to nearest look up point
2. Increased calculation complexity to use adaptive stepping
   - Adjusting step size to be within an error tolerance
   - Still slower than the CPU...
3. Treated x,y,z coordinates in parallel (3 threads in block)
   - Cross-product ($\mathbf{v} \times \mathbf{B}$) calculation needs perp. coordinates
   - Set up the threads in the block to use shared memory
   - Speed was now closer to CPU
4. Next stage was to do many particle tracks in parallel...

Magnetic Field Integration results

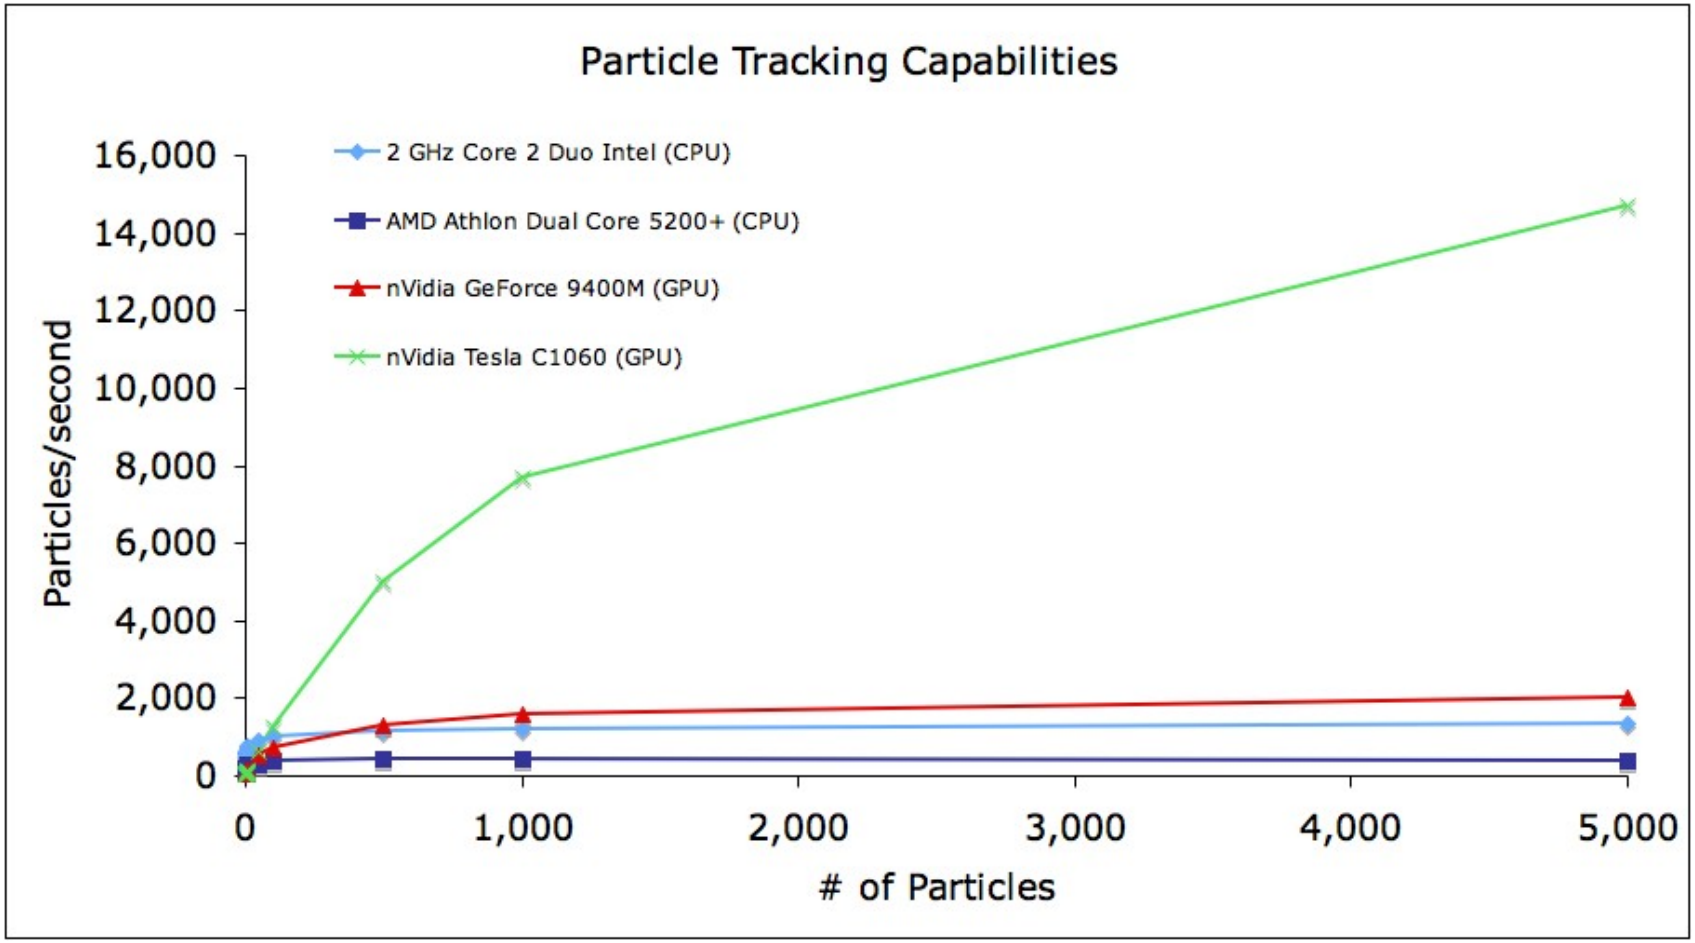

Magnetic Field Integration results

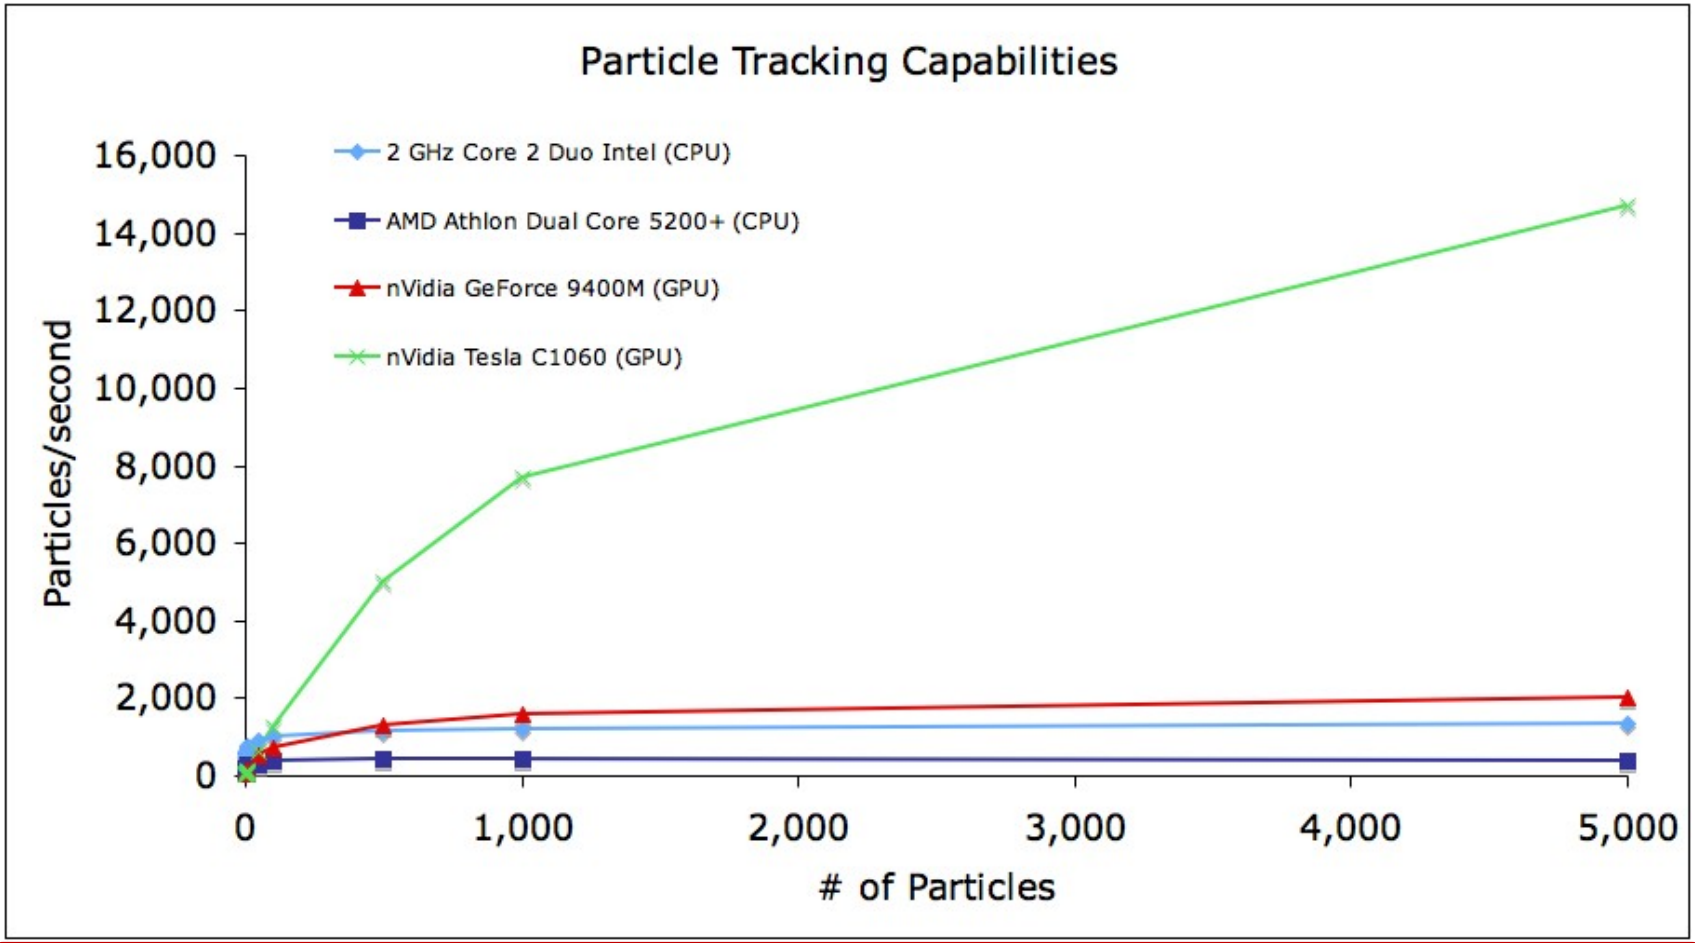

Preliminary results (Tesla C1060)

Rapidly achieved a factor 32 speedup (more in progress)

Outline

1. Introduction to the LHC and ATLAS
   - The Large Hadron Collider (LHC)
   - The ATLAS detector
   - The Higgs Boson
2. Areas of study for GPGPU adoption
   - Particle tracking in a magnetic field
   - The ATLAS trigger and data acquisition
   - The worldwide LHC computing grid
3. Using GPUs in the high level trigger
   - The Level 2 Trigger
   - The Z finder algorithm
   - The Kalman Filter

ATLAS Trigger and data acquisition

- Two bunches of protons cross head-on at 40 MHz (25 ns)

ATLAS Trigger and data acquisition

- Two bunches of protons cross head-on at 40 MHz (25 ns)
- Each bunch contains 100 billion protons

ATLAS Trigger and data acquisition

- Two bunches of protons cross head-on at 40 MHz (25 ns)
- Each bunch contains 100 billion protons
- On average there are 23 proton collisions per crossing

ATLAS Trigger and data acquisition

- Two bunches of protons cross head-on at 40 MHz (25 ns)
- Each bunch contains 100 billion protons
- On average there are 23 proton collisions per crossing

$\Rightarrow$ Approx. 1 billion proton collisions in detector per second

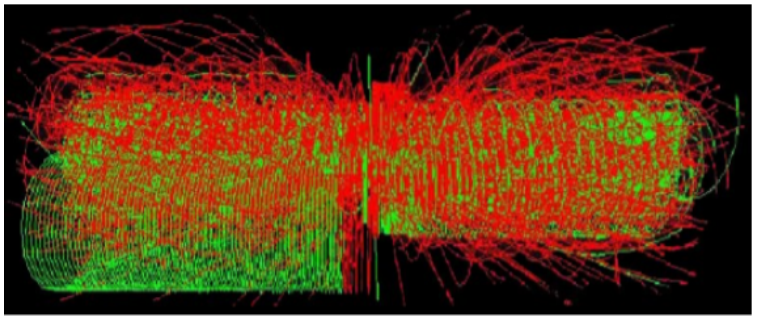

ATLAS Trigger and data acquisition

- Two bunches of protons cross head-on at 40 MHz (25 ns)
- Each bunch contains 100 billion protons
- On average there are 23 proton collisions per crossing

$\Rightarrow$ Approx. 1 billion proton collisions in detector per second

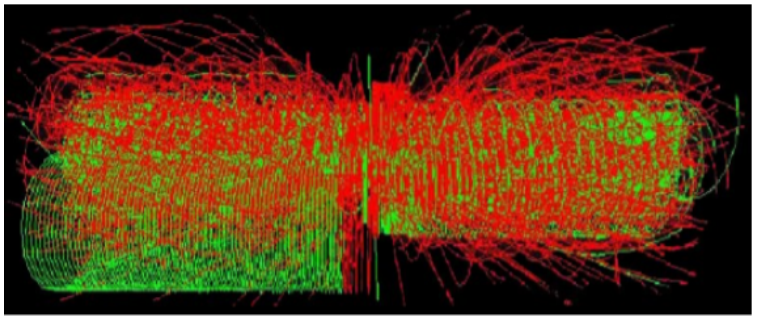

- The ATLAS detector has 140 million electronic channels

ATLAS Trigger and data acquisition

- Two bunches of protons cross head-on at 40 MHz (25 ns)
- Each bunch contains 100 billion protons
- On average there are 23 proton collisions per crossing

$\Rightarrow$ Approx. 1 billion proton collisions in detector per second

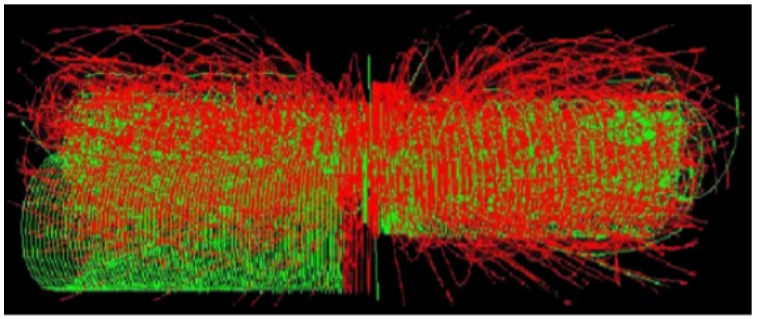

- The ATLAS detector has 140 million electronic channels

The ATLAS Data Challenge

If we recorded everything it would be Petabytes per second

ATLAS Trigger and data acquisition

- Two bunches of protons cross head-on at 40 MHz (25 ns)
- Each bunch contains 100 billion protons
- On average there are 23 proton collisions per crossing

$\Rightarrow$ Approx. 1 billion proton collisions in detector per second

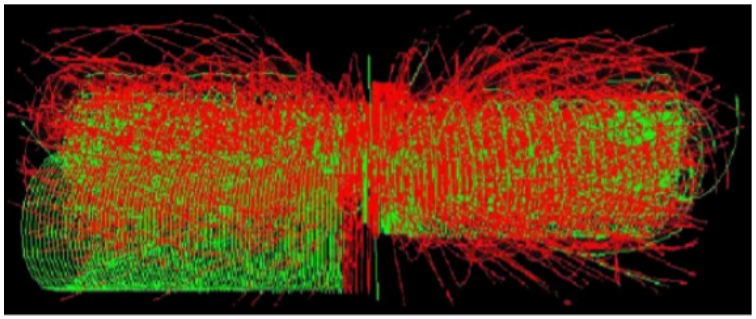

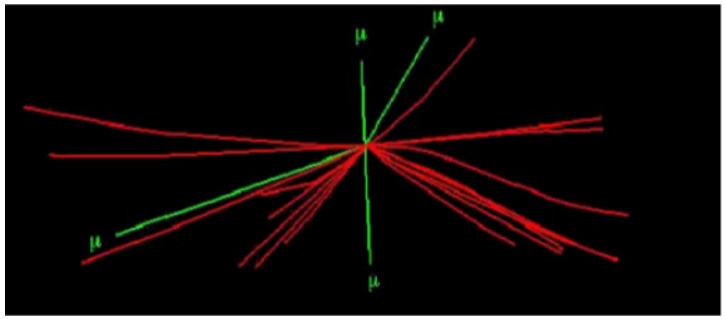

- The ATLAS detector has 140 million electronic channels

The ATLAS Data Challenge

If we recorded everything it would be Petabytes per second

The ATLAS Trigger

The solution is to select (trigger) events of interest

The ATLAS trigger and software

- Selection algorithms must be very fast & massively parallel

The ATLAS trigger and software

- Selection algorithms must be very fast & massively parallel
- Must be accurate & reliable (LHC is a $9 billion machine)

The ATLAS trigger and software

- Selection algorithms must be very fast & massively parallel
- Must be accurate & reliable (LHC is a $9 billion machine)

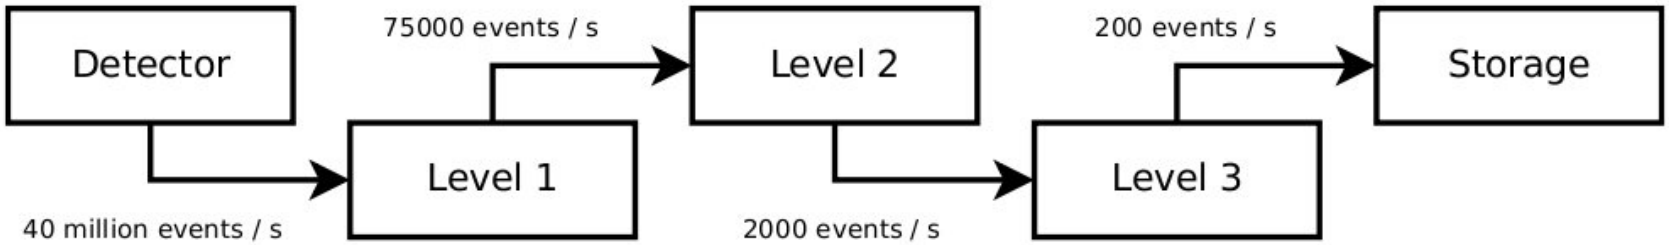

Level 1 Custom built hardware with special processor units (ASICs, FPGAs)

The ATLAS trigger and software

- Selection algorithms must be very fast & massively parallel
- Must be accurate & reliable (LHC is a $9 billion machine)

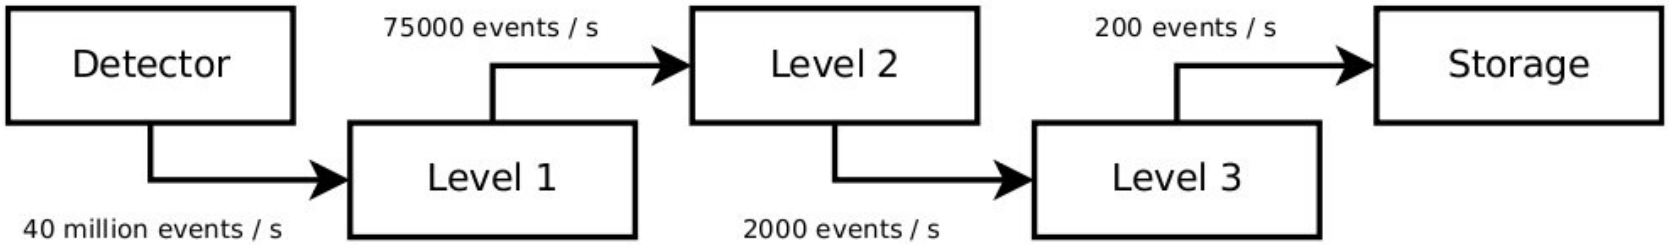

Level 1 Custom built hardware with special processor units (ASICs, FPGAs)

Level 2 Software based trigger operating on detector regions of interest (RoIs) in parallel

The ATLAS trigger and software

- Selection algorithms must be very fast & massively parallel
- Must be accurate & reliable (LHC is a $9 billion machine)

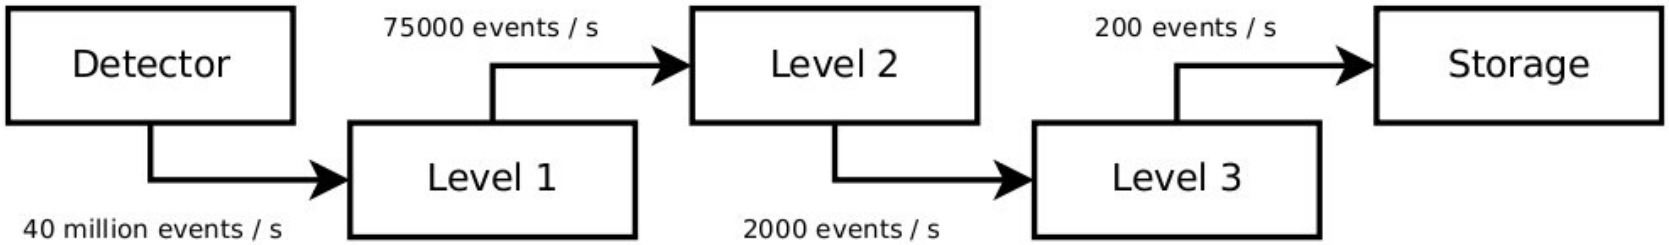

Level 1 Custom built hardware with special processor units (ASICs, FPGAs)

Level 2 Software based trigger operating on detector regions of interest (RoIs) in parallel

Level 3 Software based trigger analysing whole event signatures

The ATLAS trigger and software

- Selection algorithms must be very fast & massively parallel
- Must be accurate & reliable (LHC is a $9 billion machine)

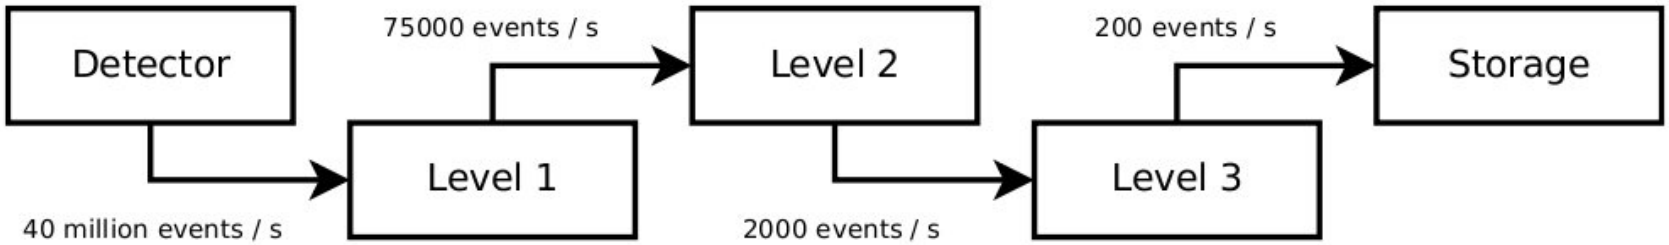

Level 1 Custom built hardware with special processor units (ASICs, FPGAs)

Level 2 Software based trigger operating on detector regions of interest (RoIs) in parallel Ideal for GPGPUs

Level 3 Software based trigger analysing whole event signatures Ideal for GPGPUs

The software (high-level) trigger farm

- Level 2 and Level 3 triggers collectively called the high-level trigger (HLT)

Figure: L2 supervisors, event builder, data logger

The software (high-level) trigger farm

- Level 2 and Level 3 triggers collectively called the high-level trigger (HLT)
- Around 1000 PCs (XPU: Interchangeable processing unit (i.e. Level 2 or Level 3)

Figure: L2 supervisors, event builder, data logger

The software (high-level) trigger farm

- Level 2 and Level 3 triggers collectively called the high-level trigger (HLT)
- Around 1000 PCs (XPU: Interchangeable processing unit (i.e. Level 2 or Level 3)
- For our GPGPU studies we decided to study algorithms that are run in the Level 2 (Z finder and Kalman filter)

Figure: L2 supervisors, event builder, data logger

Outline

1. Introduction to the LHC and ATLAS
   - The Large Hadron Collider (LHC)
   - The ATLAS detector
   - The Higgs Boson
2. Areas of study for GPGPU adoption
   - Particle tracking in a magnetic field
   - The ATLAS trigger and data acquisition
   - The worldwide LHC computing grid
3. Using GPUs in the high level trigger
   - The Level 2 Trigger
   - The Z finder algorithm
   - The Kalman Filter

The LHC computing Grid

After triggering the LHC experiments still produce vast amounts of data!

The LHC computing Grid

After triggering the LHC experiments still produce vast amounts of data!
We developed worldwide LHC computing grid infrastructure

- Approximately 15 PB of data recorded per annum
- Currently >100,000 processors across Grid
- 130 sites in 34 countries

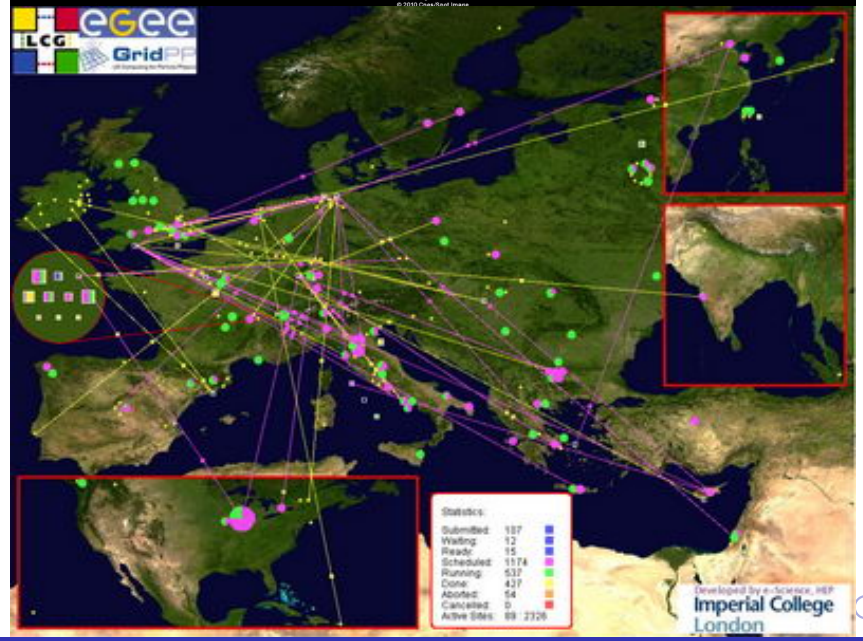

The LHC computing Grid

After triggering the LHC experiments still produce vast amounts of data!
We developed worldwide LHC computing grid infrastructure

- Approximately 15 PB of data recorded per annum
- Currently >100,000 processors across Grid
- 130 sites in 34 countries

We also simulate the physics events (∼ 1000 cpu seconds per event)

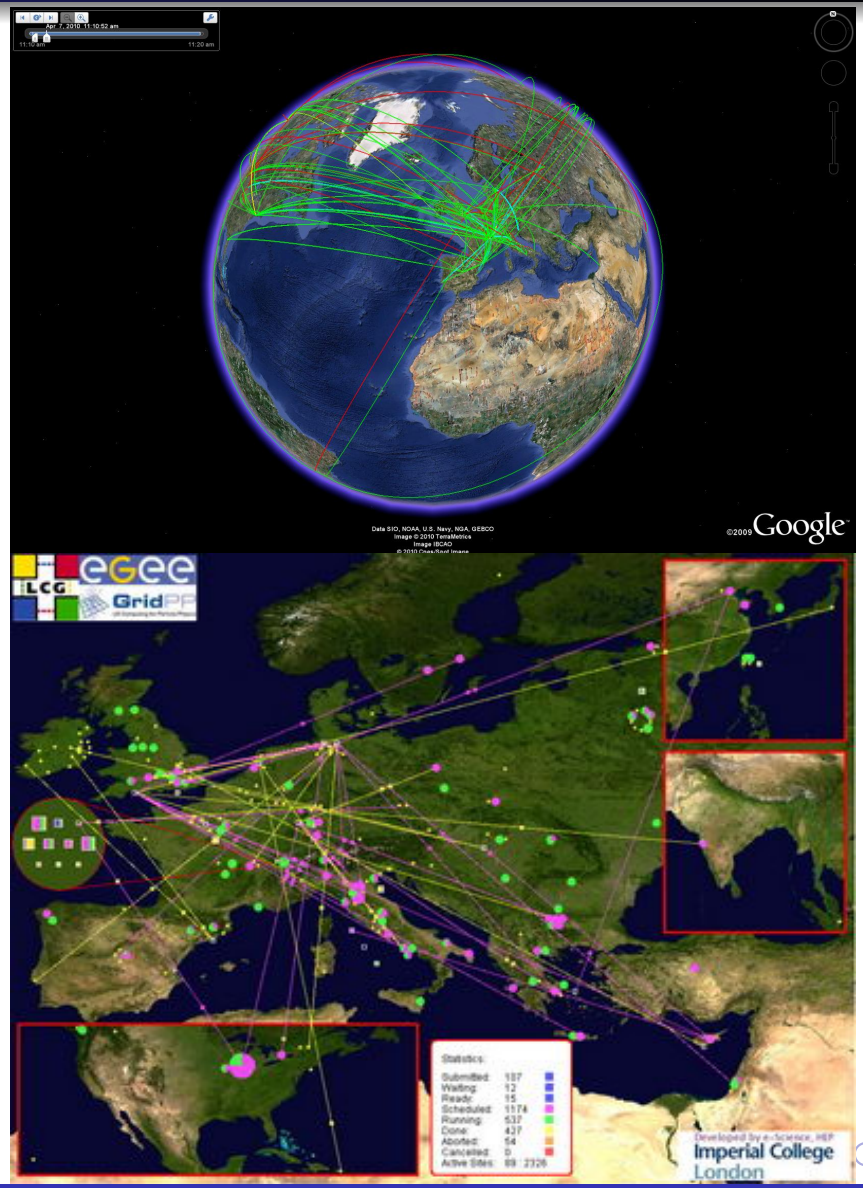

The LHC computing Grid

After triggering the LHC experiments still produce vast amounts of data!
We developed worldwide LHC computing grid infrastructure

- Approximately 15 PB of data recorded per annum
- Currently >100,000 processors across Grid
- 130 sites in 34 countries

We also simulate the physics events (∼ 1000 cpu seconds per event)

Up to eight million events simulated daily

Failure rate is less than 10^{-6}

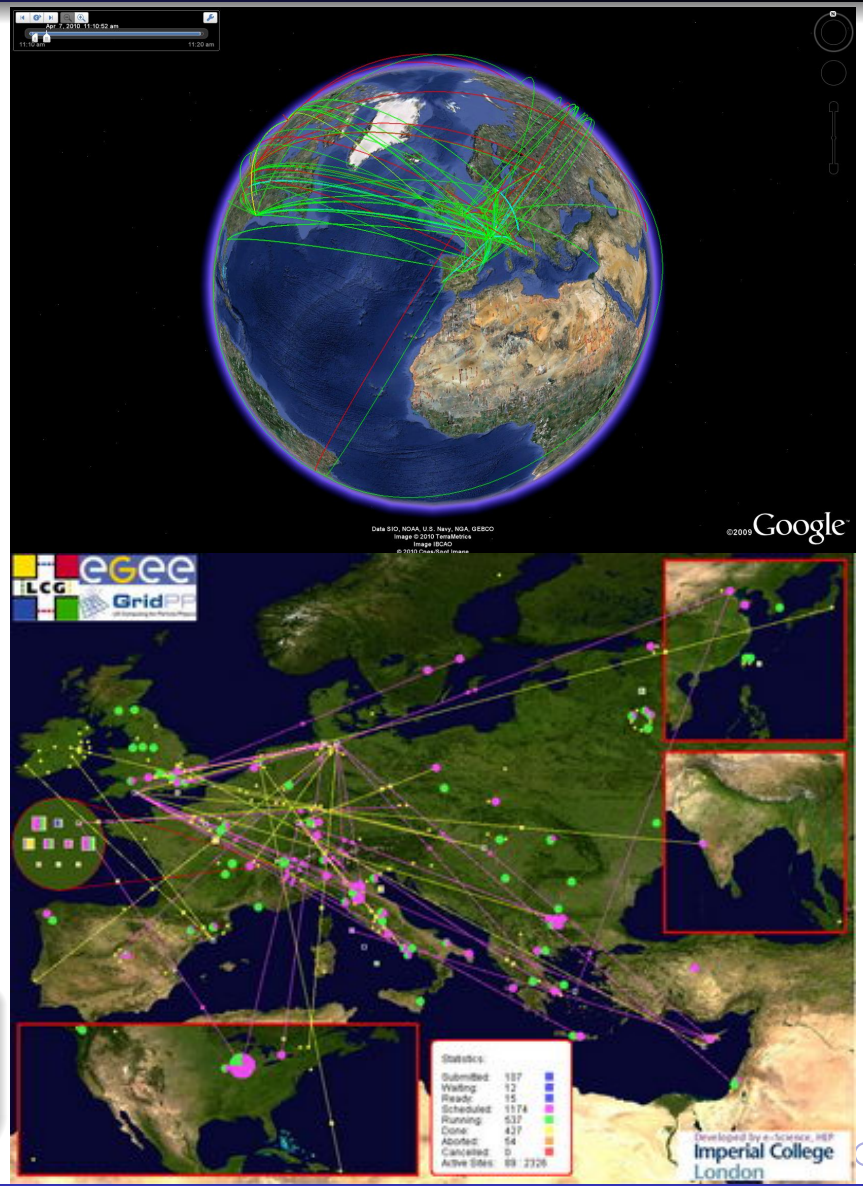

Outline

1. Introduction to the LHC and ATLAS
   - The Large Hadron Collider (LHC)
   - The ATLAS detector
   - The Higgs Boson
2. Areas of study for GPGPU adoption
   - Particle tracking in a magnetic field
   - The ATLAS trigger and data acquisition
   - The worldwide LHC computing grid
3. Using GPUs in the high level trigger
   - The Level 2 Trigger
   - The Z finder algorithm
   - The Kalman Filter

Level 2 Trigger Routines

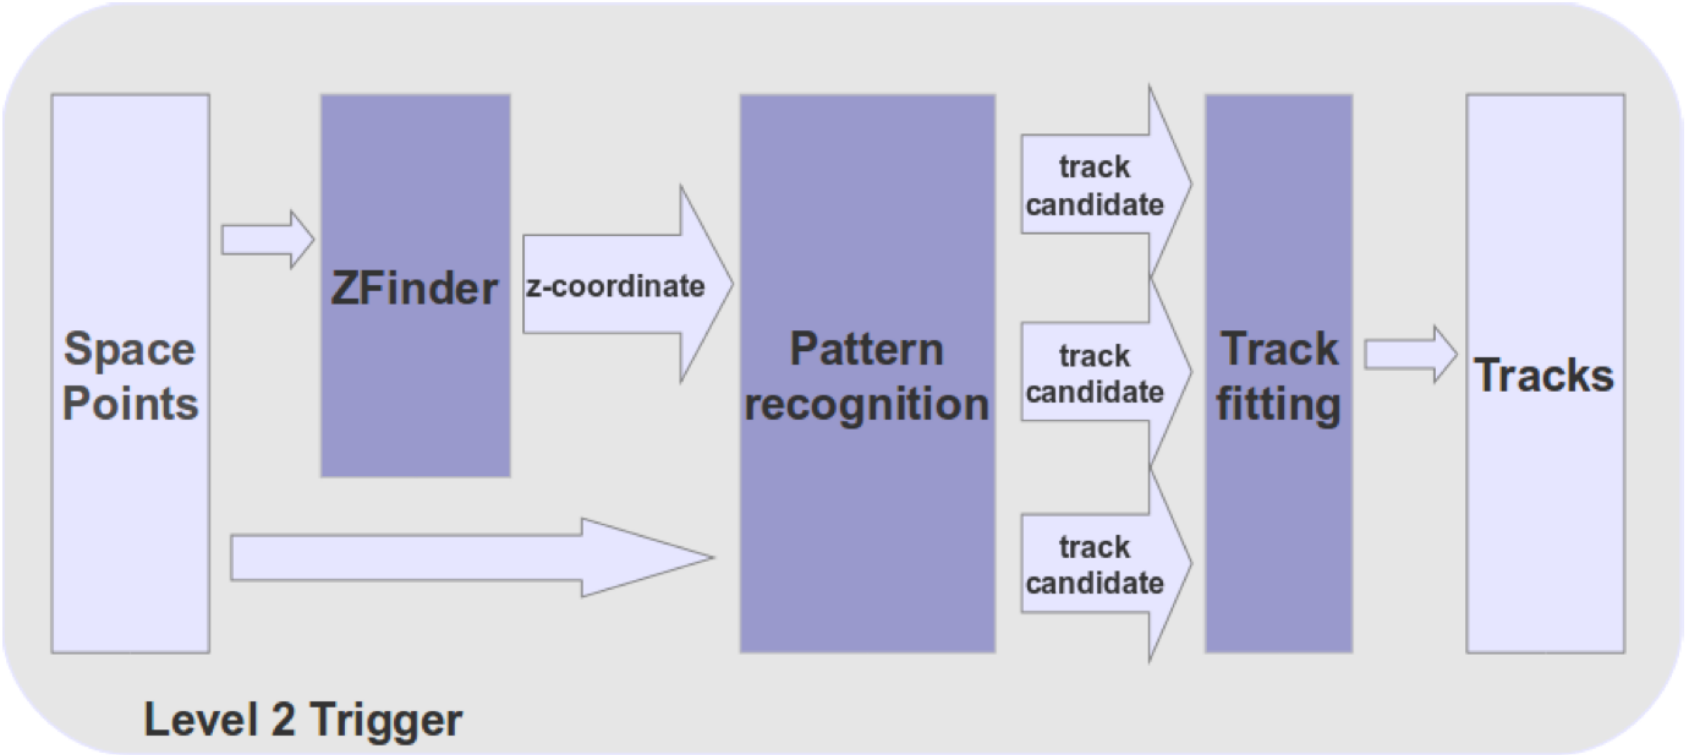

The Level 2 regions of interest (ROIs)

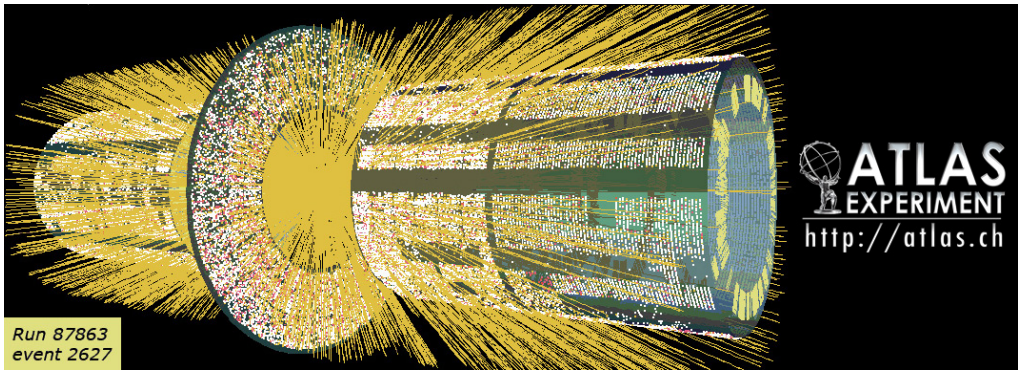

- We take a cross-section view of the detector
- Break it up into regions of interest (ROIs)
  i.e. “phi slices” (ϕ coordinate)
- Candidate for parallelisation using GPUs

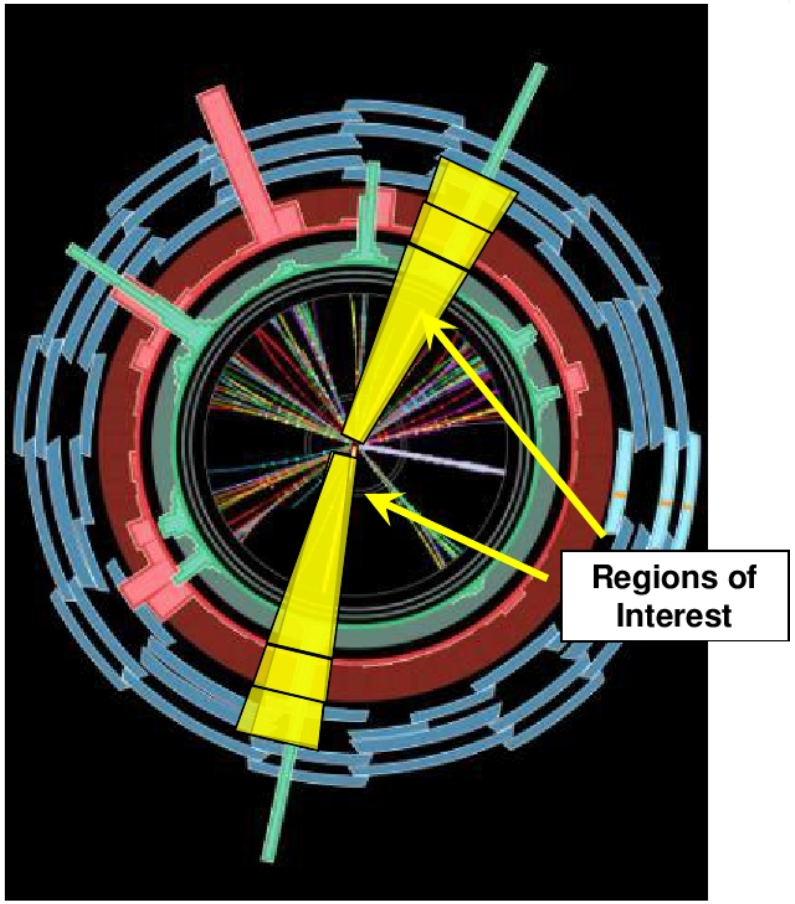

Cross section view of the ATLAS detector

Outline

1. Introduction to the LHC and ATLAS
   - The Large Hadron Collider (LHC)
   - The ATLAS detector
   - The Higgs Boson
2. Areas of study for GPGPU adoption
   - Particle tracking in a magnetic field
   - The ATLAS trigger and data acquisition
   - The worldwide LHC computing grid
3. Using GPUs in the high level trigger
   - The Level 2 Trigger
   - The Z finder algorithm
   - The Kalman Filter

The Z Finder Algorithm

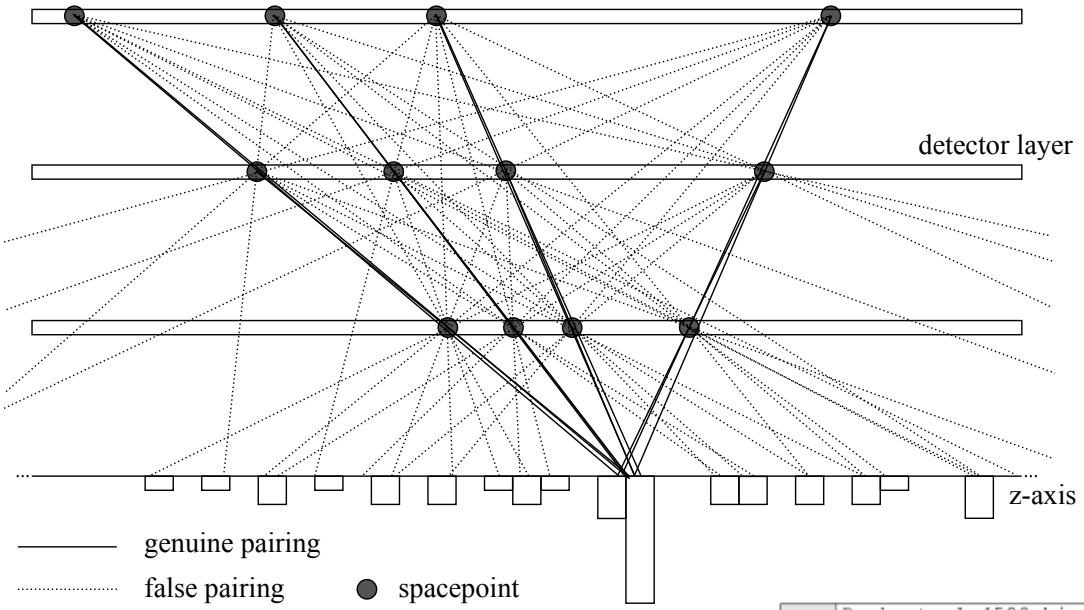

- Process each combination of detector hits ("spacepoints") and extrapolate back to the beam line.
- The histogram peak is the chosen interaction point.

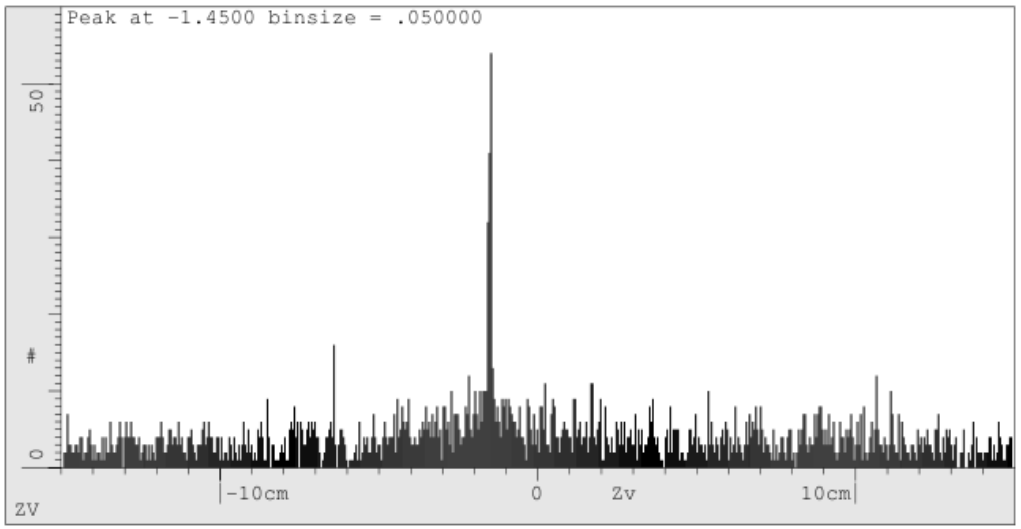

Z Finder Algorithm Test Case

- Z Finder code extracted from ATLAS framework for feasibility studies with CUDA.
- Timing performance measured using two samples of simulated events (low and high *luminosity*).
- Comparison of Tesla and Fermi architectures for each code iteration.

	lowlum	**highlum**
Spacepoints	333	8104

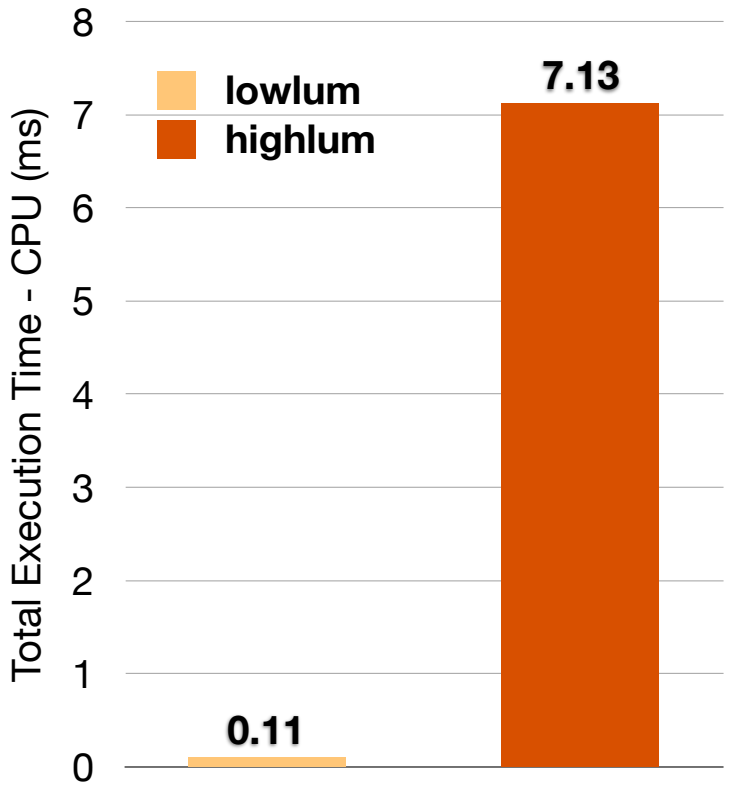

Z Finder Kernel: Histogram Summation

Code Iterations

- Single thread per ϕ slice.
- Thread block per ϕ slice.
- Histogram per thread block in shared memory.
- Improve spacepoint pair allocation method.

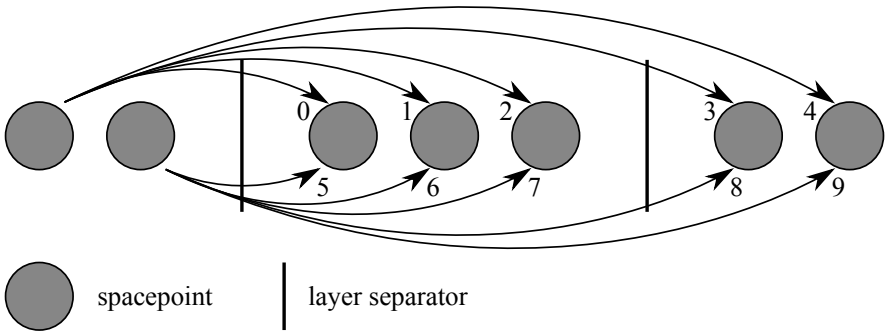

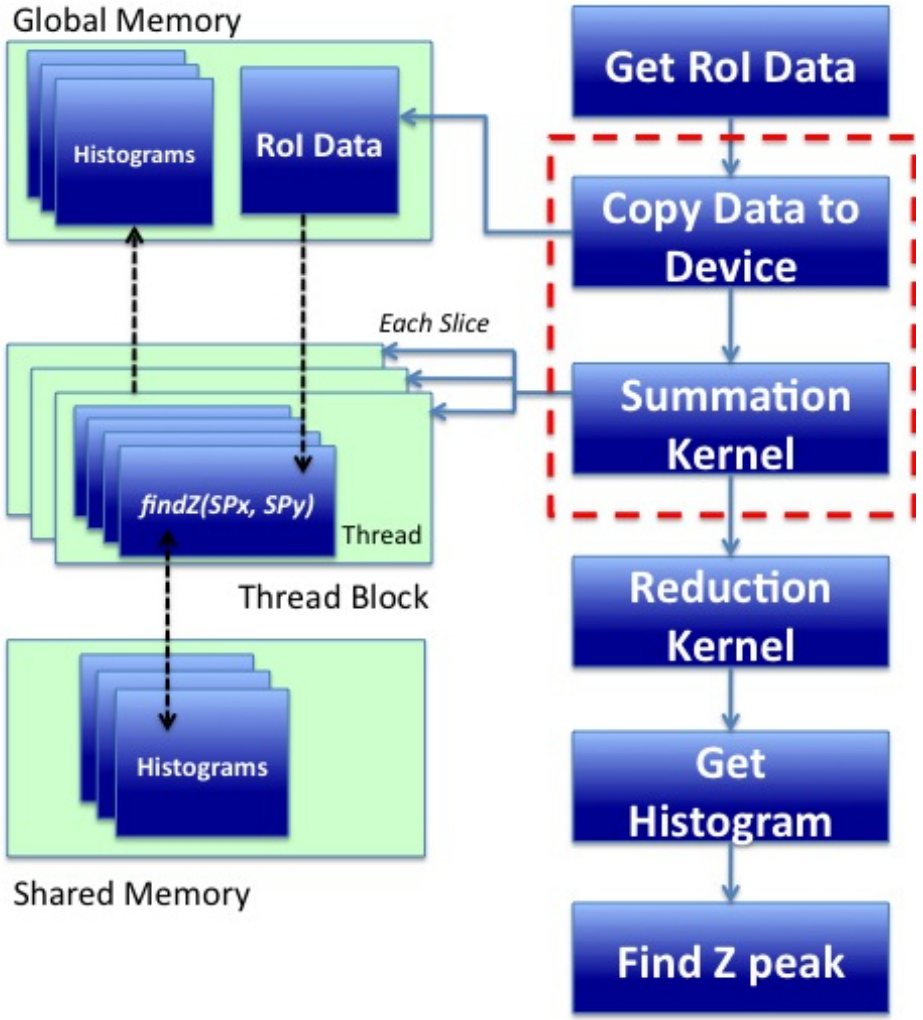

Z Finder Kernel: Histogram Summation

Code Iterations

- Single thread per ϕ slice.
- Thread block per ϕ slice.
- Histogram per thread block in shared memory.
- Improve spacepoint pair allocation method.

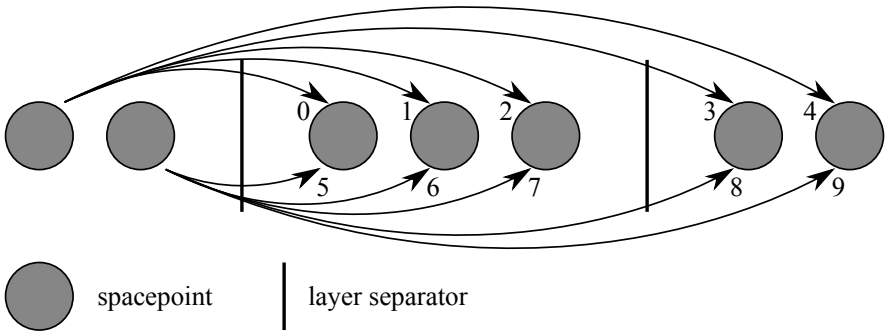

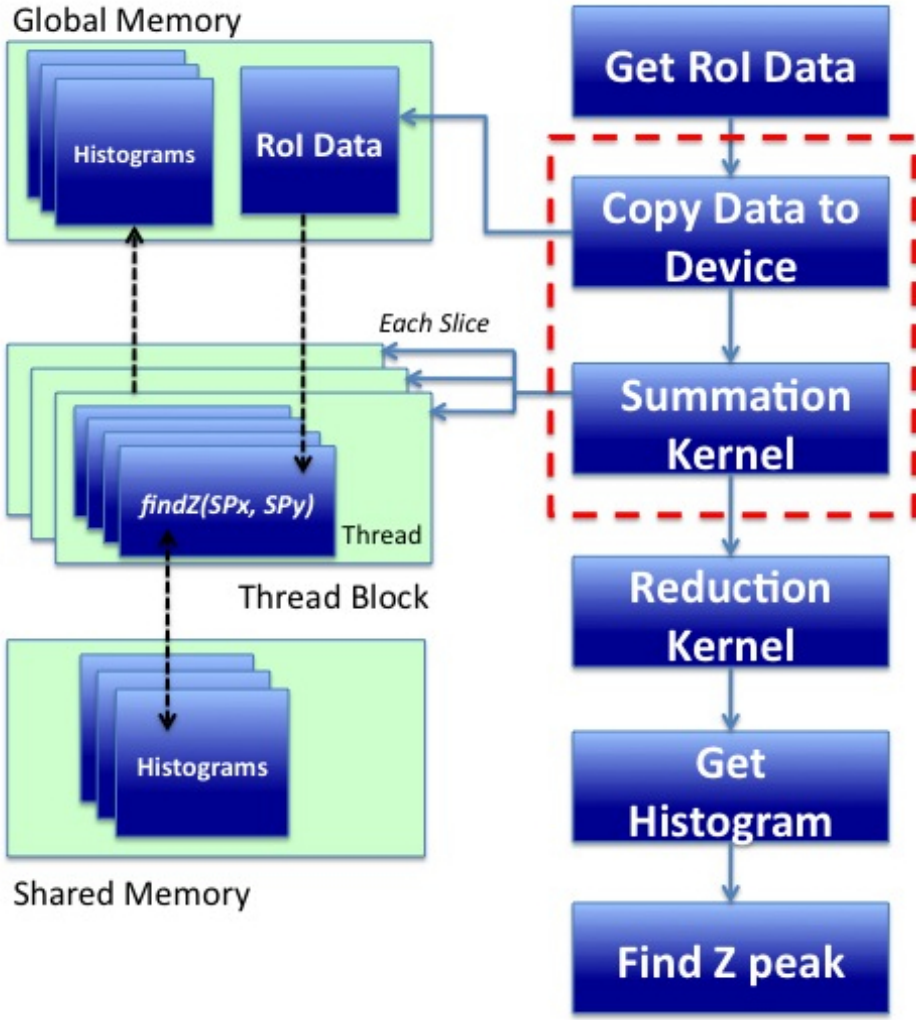

Z Finder Kernel: Histogram Summation

Code Iterations

- Single thread per ϕ slice.
- Thread block per ϕ slice.
- Histogram per thread block in shared memory.
- Improve spacepoint pair allocation method.

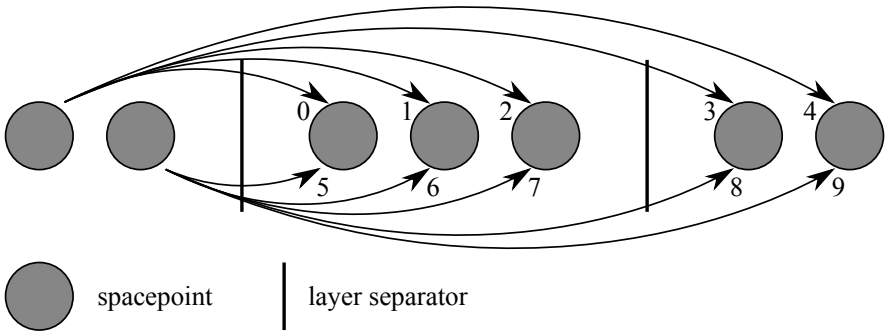

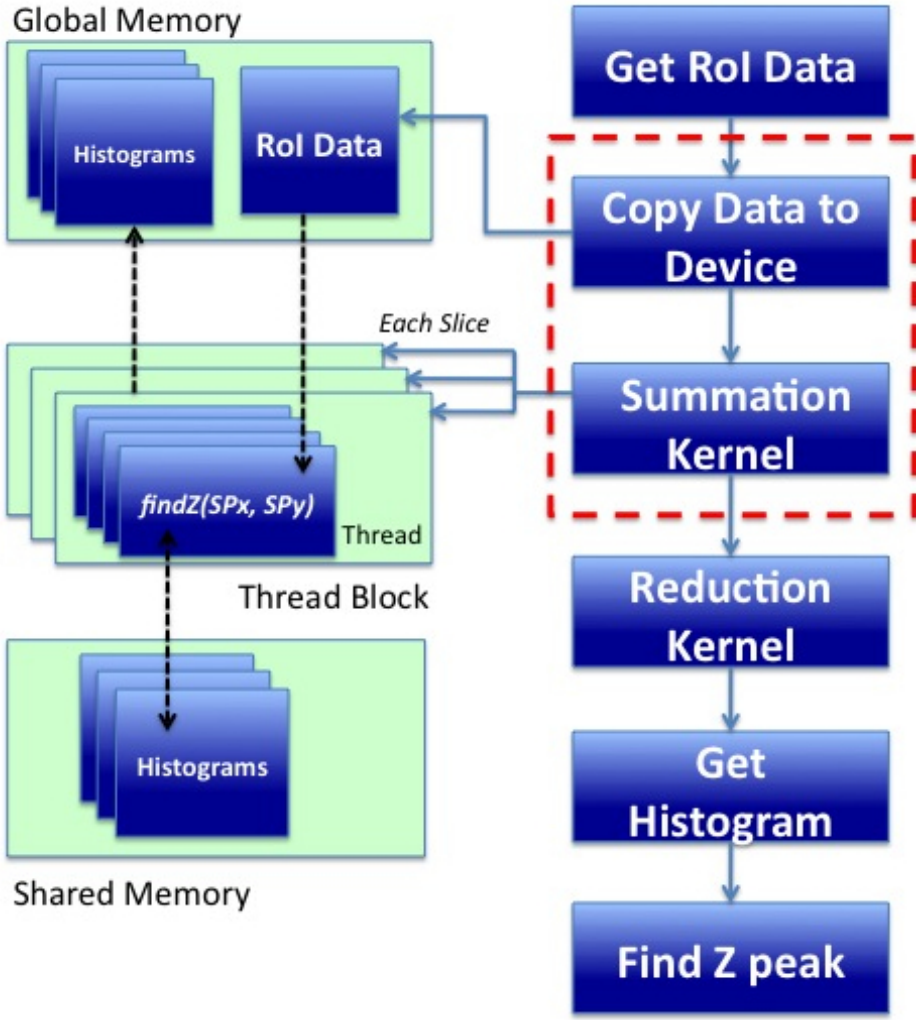

Z Finder Kernel: Histogram Summation

Code Iterations

- Single thread per ϕ slice.
- Thread block per ϕ slice.
- Histogram per thread block in shared memory.
- Improve spacepoint pair allocation method.

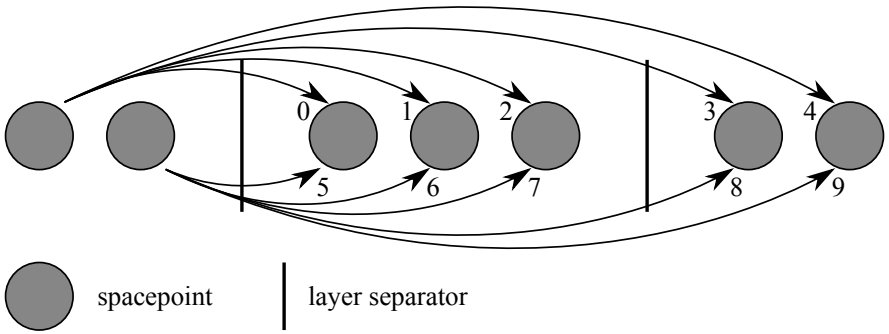

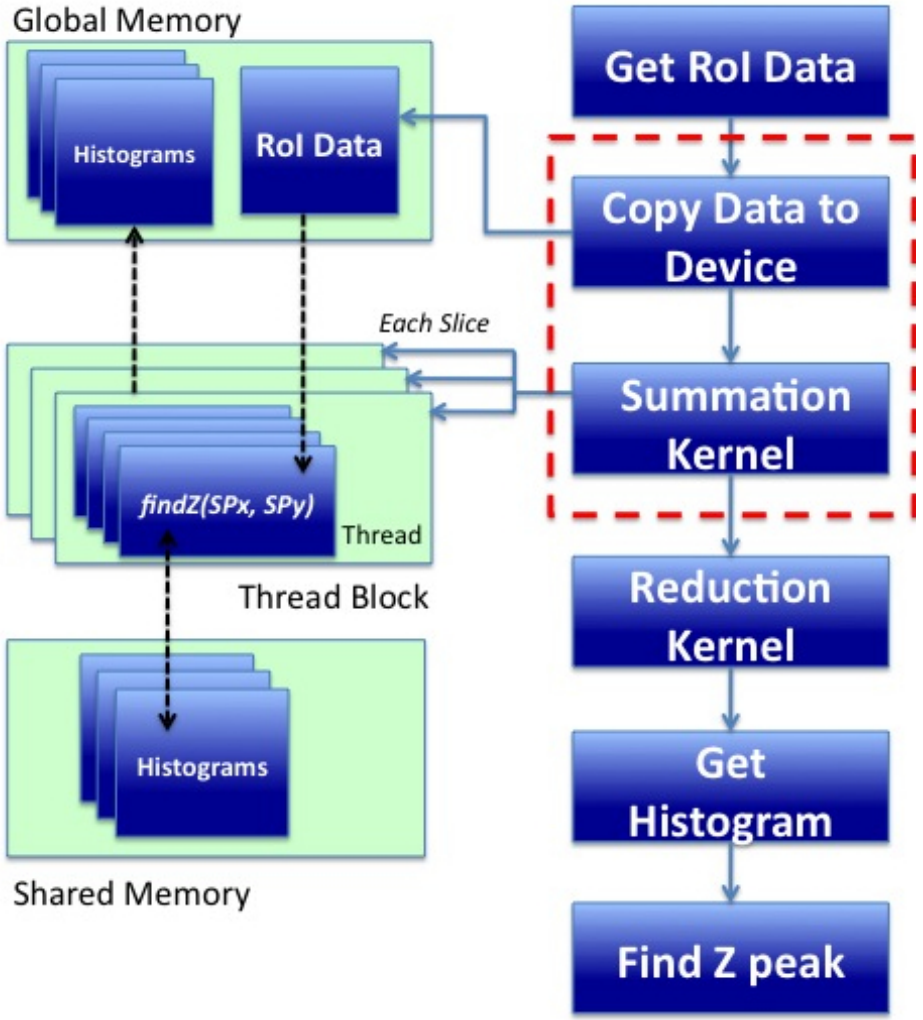

ZFinder Kernel: Histogram Summation Results

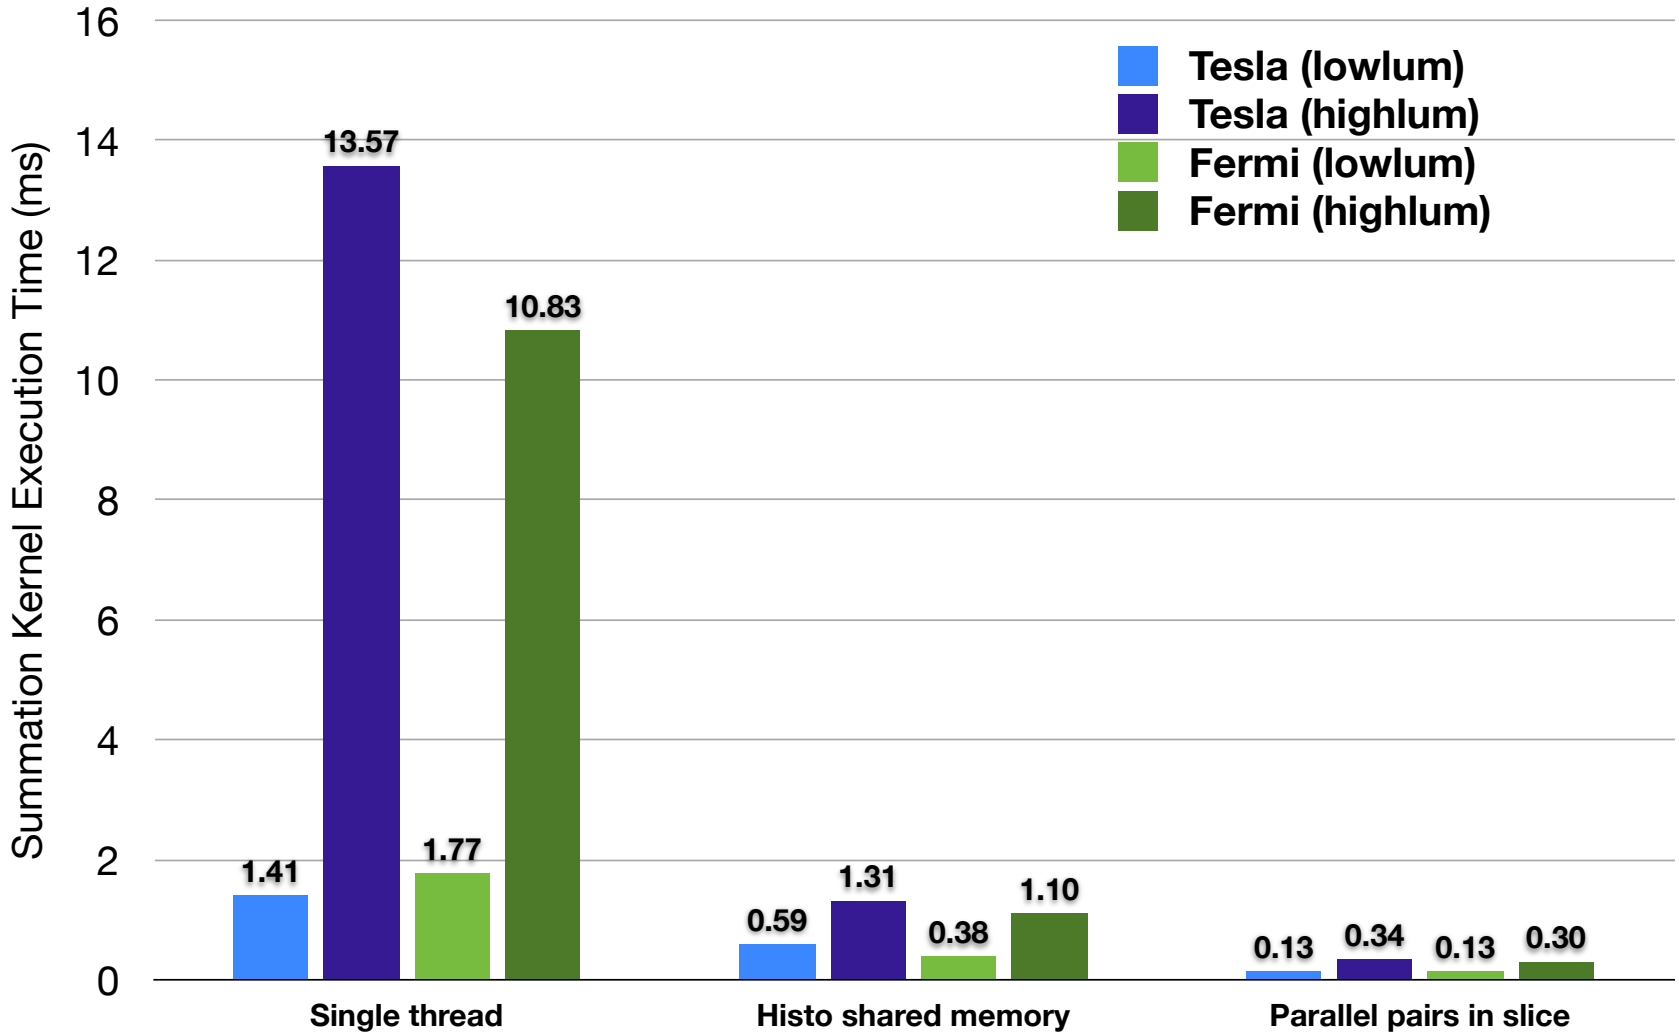

ZFinder Kernel: Histogram Combination

Code Iterations

- Combine histograms on the GPU $\Rightarrow$ *reduce data transfer by* $\sim$*500x*
- Reduce the data to a single histogram in multiple steps.

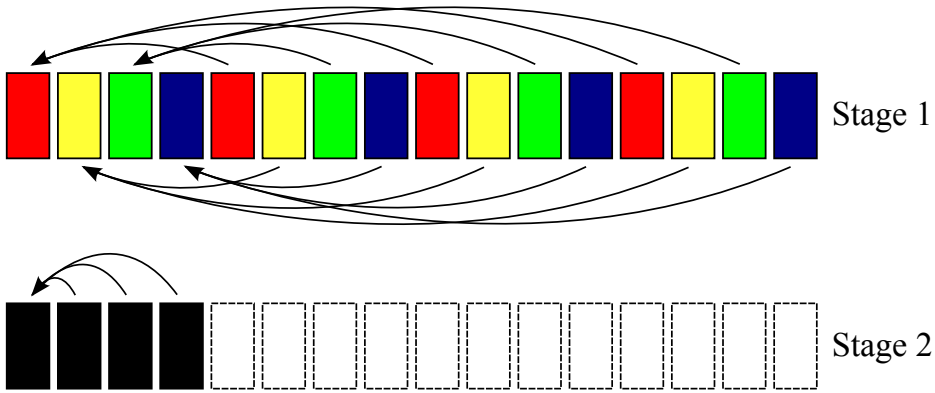

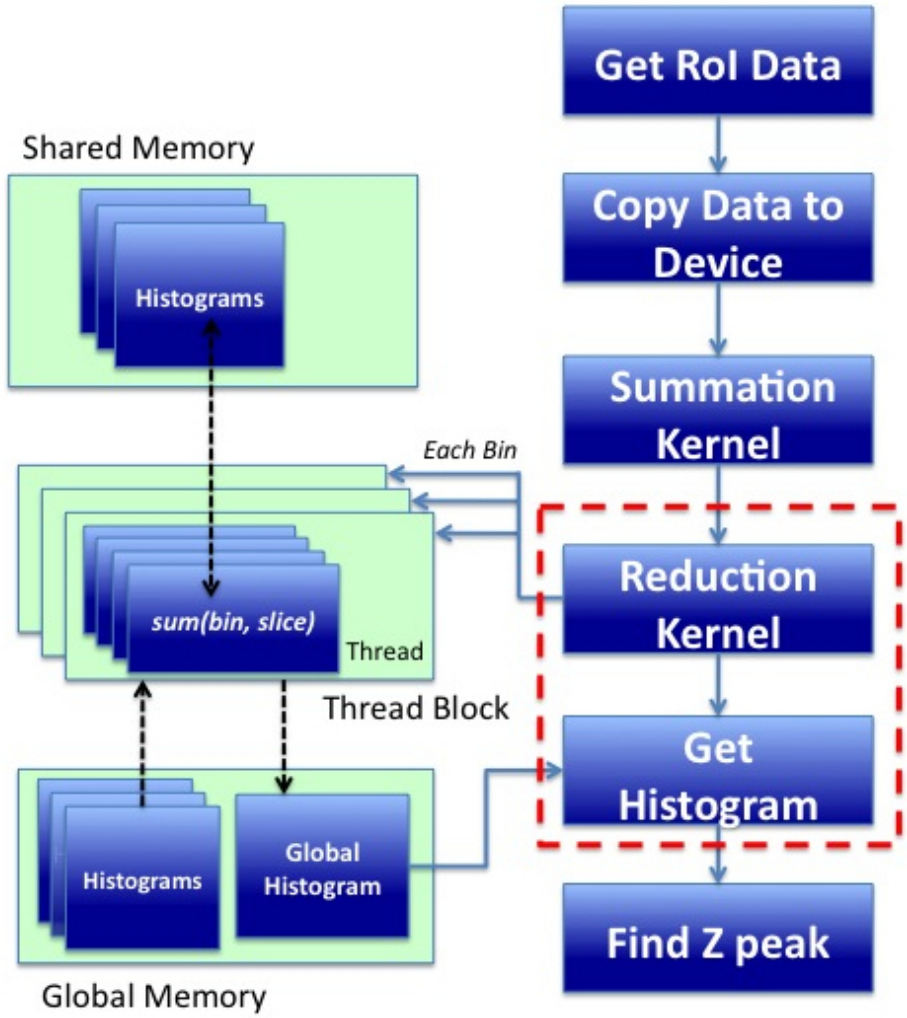

ZFinder Kernel: Histogram Combination

Code Iterations

- Combine histograms on the GPU $\Rightarrow$ *reduce data transfer by* $\sim$*500x*
- Reduce the data to a single histogram in multiple steps.

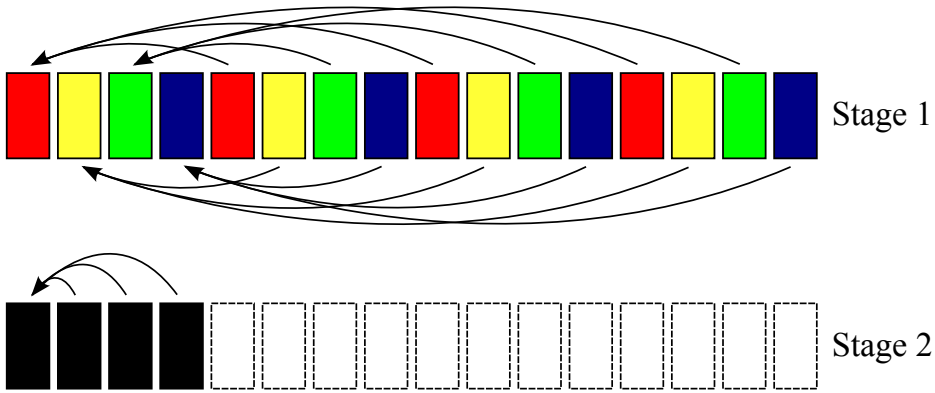

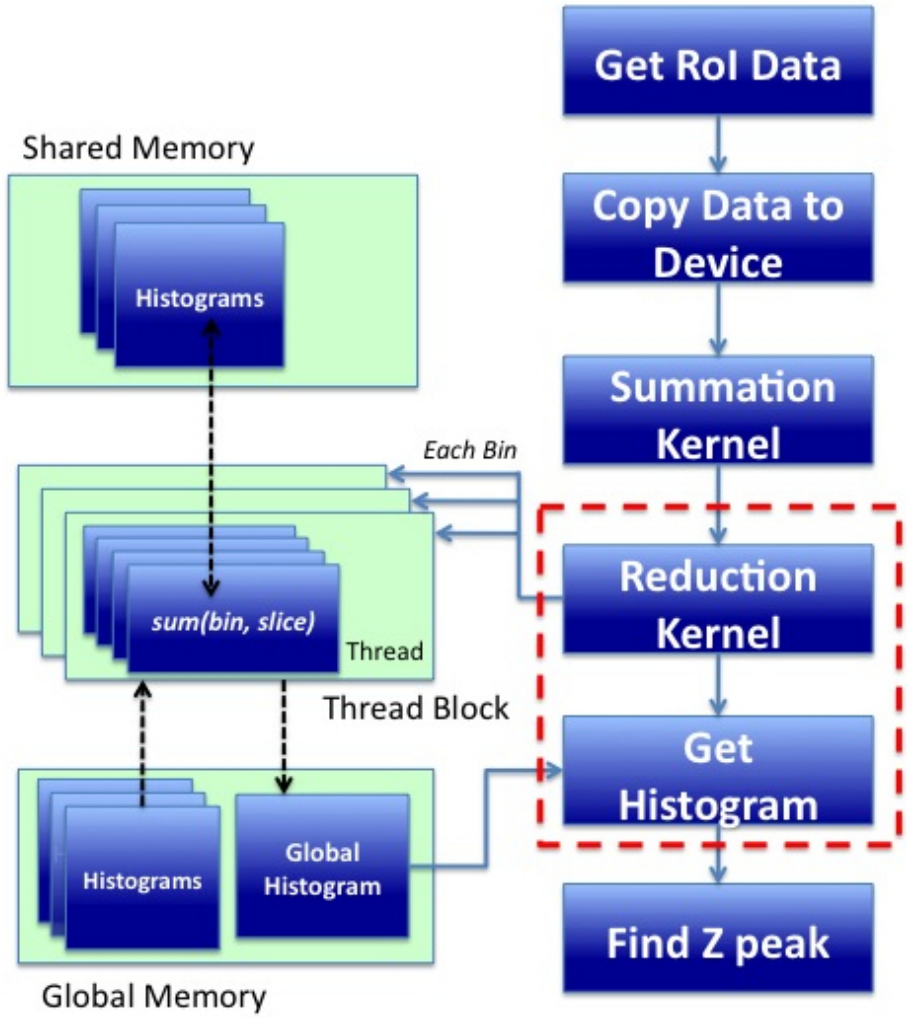

Z Finder Kernel: Streaming

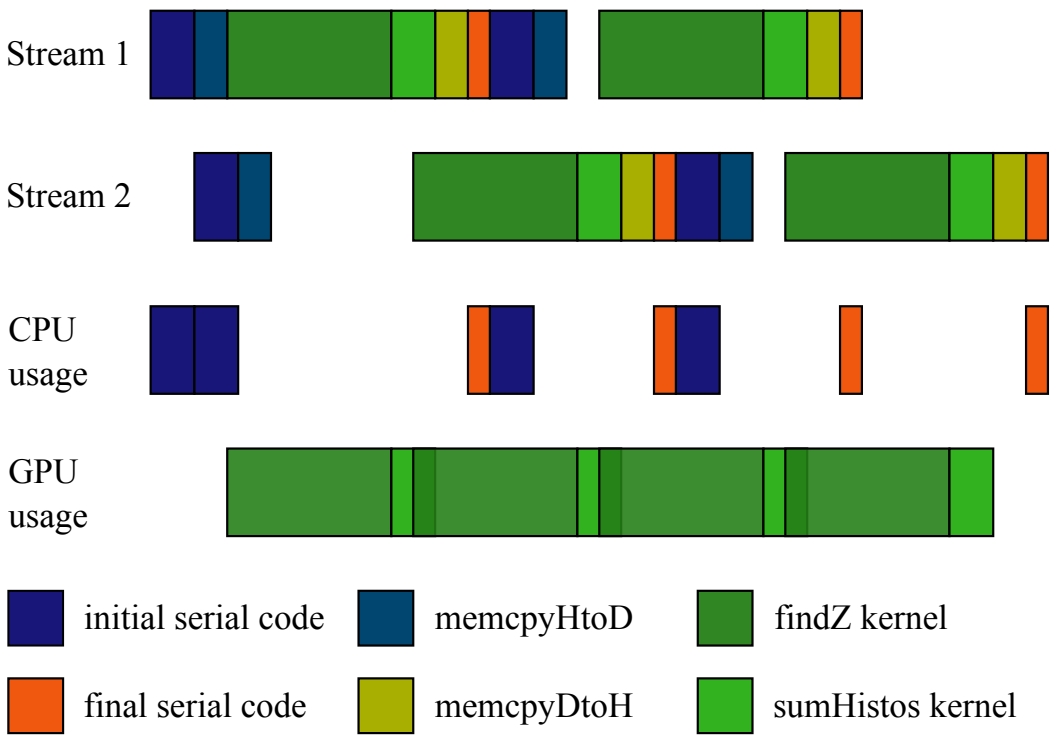

- Each RoI calculation independent $\Rightarrow$ use CUDA streams.
- Successful in disguising any host to device transfer latency.

Z Finder Kernel Results

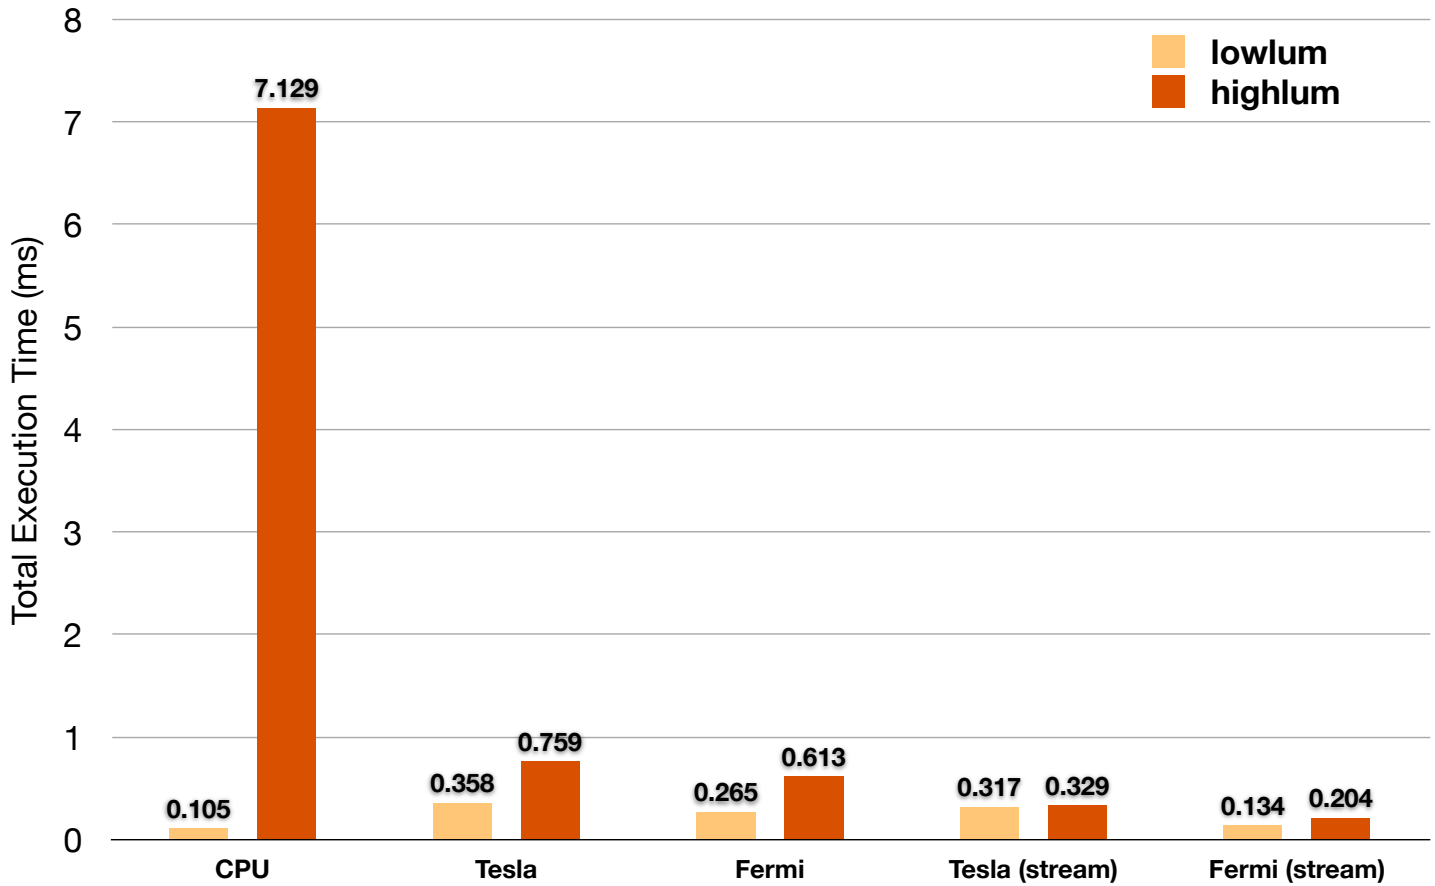

- Initial timing results show up to *35x* speed up (Fermi).
- Performance studies continuing with *triplets* of spacepoints.

Outline

1. Introduction to the LHC and ATLAS
   - The Large Hadron Collider (LHC)
   - The ATLAS detector
   - The Higgs Boson
2. Areas of study for GPGPU adoption
   - Particle tracking in a magnetic field
   - The ATLAS trigger and data acquisition
   - The worldwide LHC computing grid
3. Using GPUs in the high level trigger
   - The Level 2 Trigger
   - The Z finder algorithm
   - The Kalman Filter

The Kalman Filter

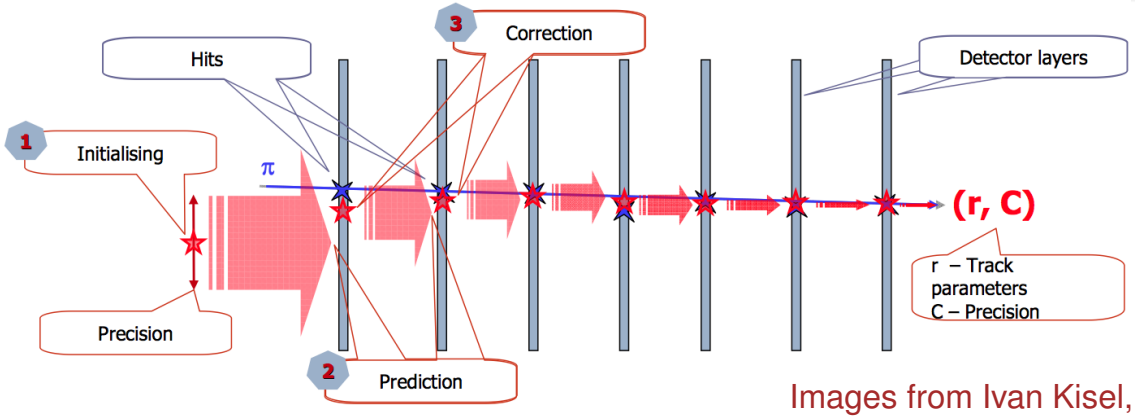

Images from Ivan Kisel, GSI

- Particle tracks reconstructed using the Kalman filter method.
- The trajectory of a track is predicted using detector hits as input.
- A backward smoothing filter is applied after the final Kalman Filter estimation.

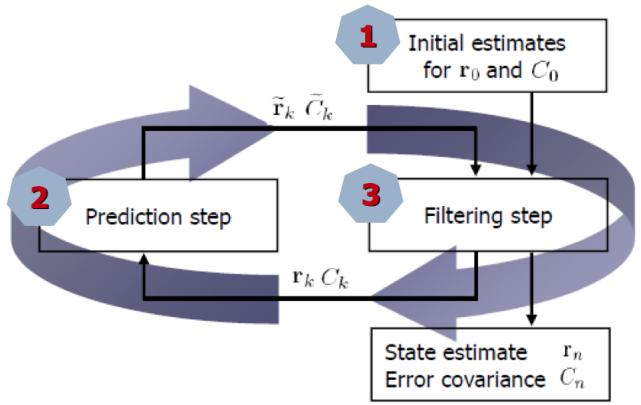

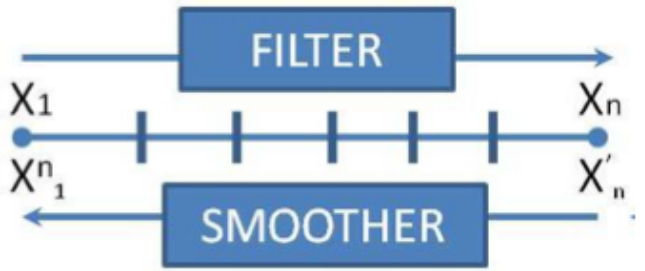

GPU Motivation for Track Reconstruction

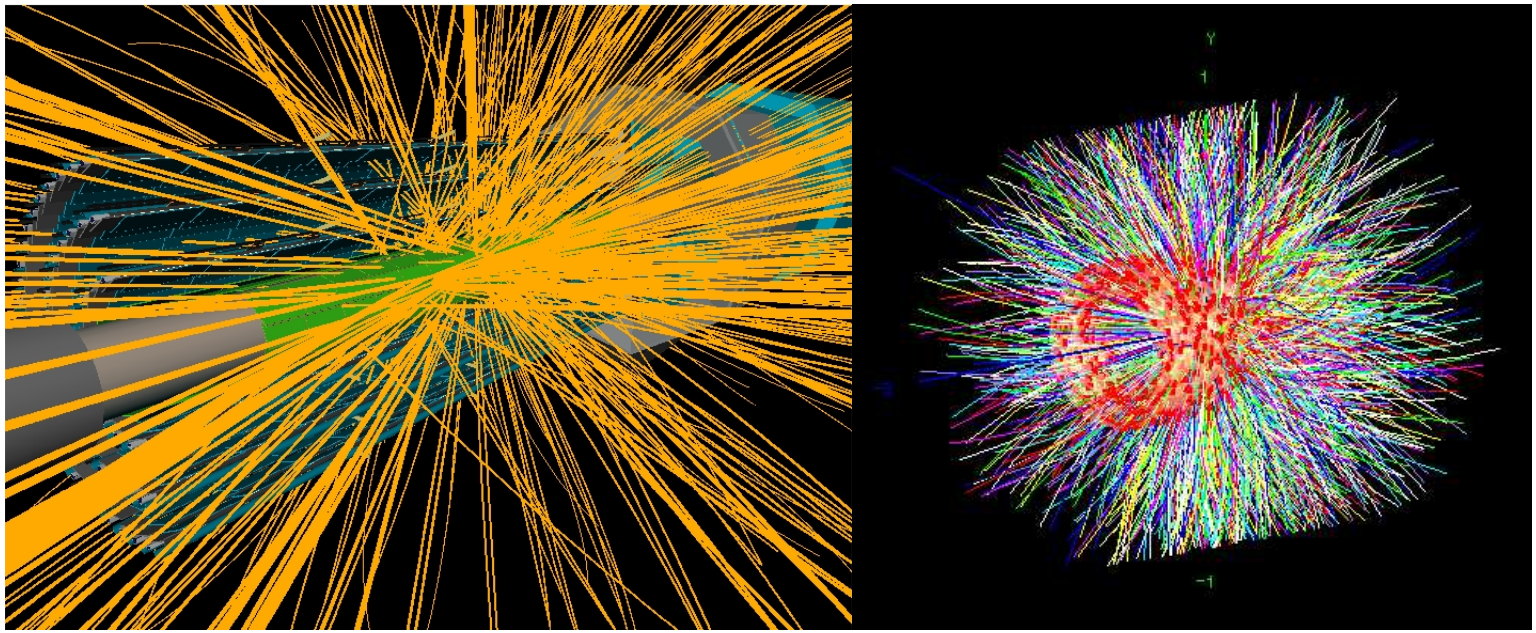

ATLAS simulations of high luminosity events

- Potentially *thousands* of tracks to process for every event.
- Significant acceleration possible by reconstructing one track per GPU thread.

ATLAS Kalman Filter Framework

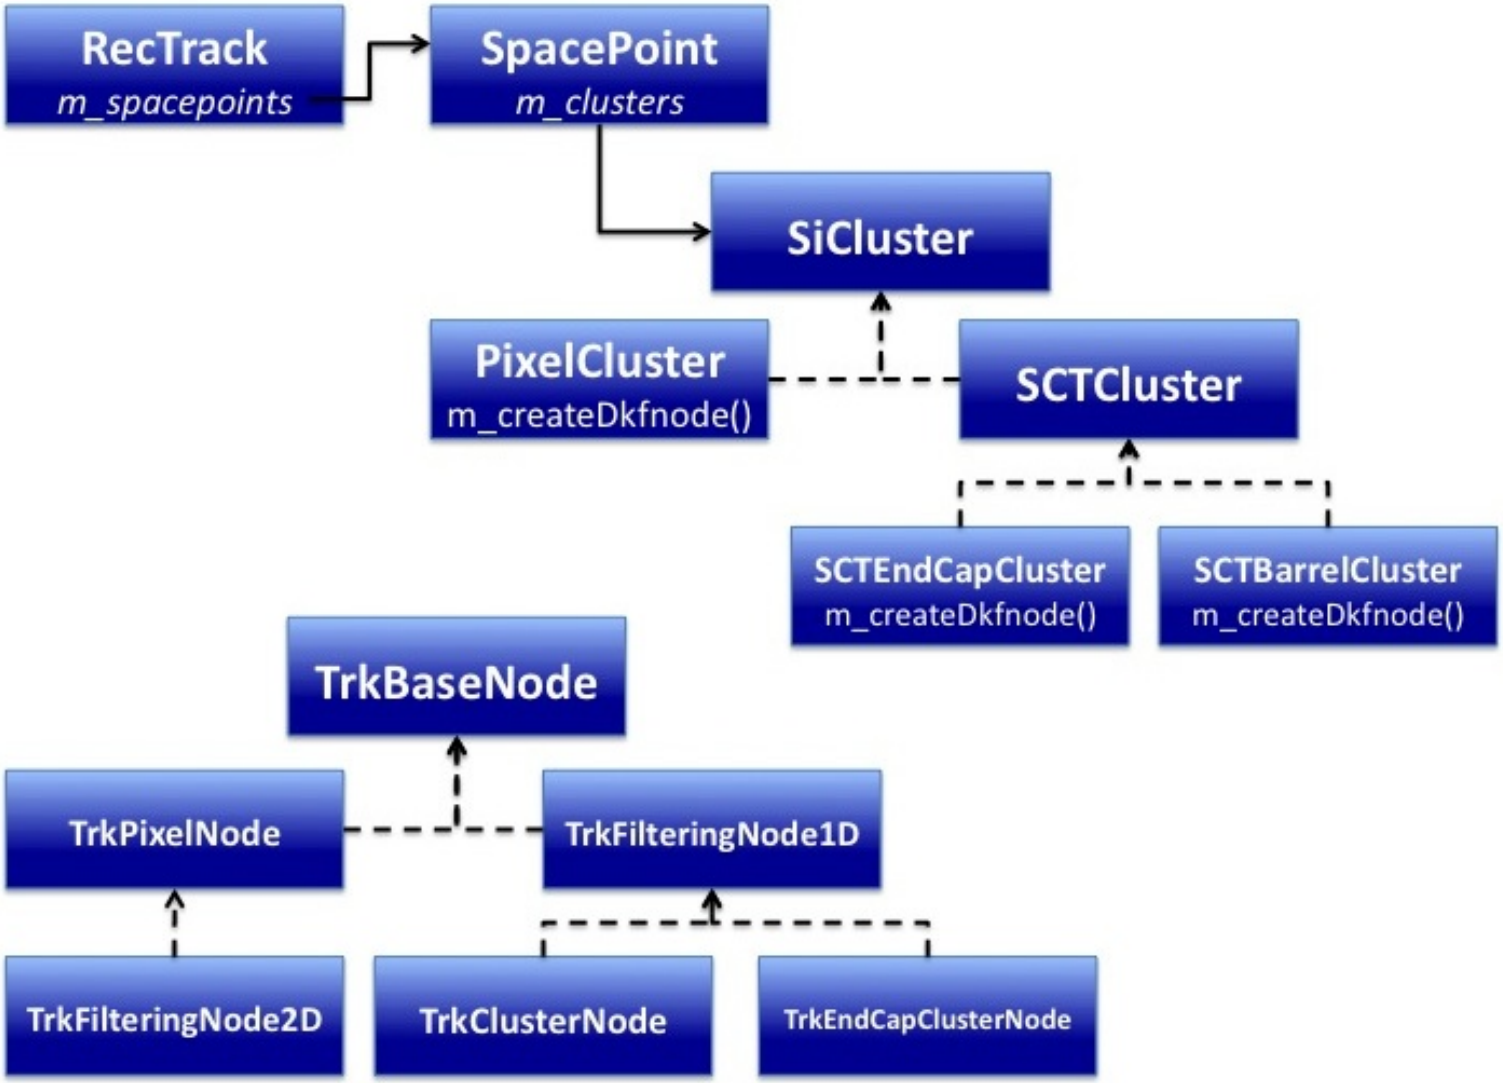

CUDA Challenges

Initial Complications

- Class inheritance structure captures filter specialism for each sub-detector.
- Dynamic creation of objects in the main routine.
- Track state retained at each filtering step.
- Main routine has over 2000+ lines of code with multiple branches.

Feasibility Studies

- Standalone version successfully ported to C.
- Pre-allocated memory needed for track objects.
- Promising results $\Rightarrow$ need to reduce memory usage.

CUDA Challenges

Initial Complications

- Class inheritance structure captures filter specialism for each sub-detector.
- Dynamic creation of objects in the main routine.
- Track state retained at each filtering step.
- Main routine has over 2000+ lines of code with multiple branches.

Feasibility Studies

- Standalone version successfully ported to C.
- Pre-allocated memory needed for track objects.
- Promising results $\Rightarrow$ need to reduce memory usage.

CUDA Challenges

Initial Complications

- Class inheritance structure captures filter specialism for each sub-detector.
- Dynamic creation of objects in the main routine.
- Track state retained at each filtering step.
- Main routine has over 2000+ lines of code with multiple branches.

Feasibility Studies

- Standalone version successfully ported to C.
- Pre-allocated memory needed for track objects.
- Promising results $\Rightarrow$ need to reduce memory usage.

CUDA Challenges

Initial Complications

- Class inheritance structure captures filter specialism for each sub-detector.
- Dynamic creation of objects in the main routine.
- Track state retained at each filtering step.
- Main routine has over 2000+ lines of code with multiple branches.

Feasibility Studies

- Standalone version successfully ported to C.
- Pre-allocated memory needed for track objects.
- Promising results $\Rightarrow$ need to reduce memory usage.

CUDA Challenges

Initial Complications

- Class inheritance structure captures filter specialism for each sub-detector.
- Dynamic creation of objects in the main routine.
- Track state retained at each filtering step.
- Main routine has over 2000+ lines of code with multiple branches.

Feasibility Studies

- Standalone version successfully ported to C.
- Pre-allocated memory needed for track objects.
- Promising results $\Rightarrow$ need to reduce memory usage.

CUDA Challenges

Initial Complications

- Class inheritance structure captures filter specialism for each sub-detector.
- Dynamic creation of objects in the main routine.
- Track state retained at each filtering step.
- Main routine has over 2000+ lines of code with multiple branches.

Feasibility Studies

- Standalone version successfully ported to C.
- Pre-allocated memory needed for track objects.
- Promising results $\Rightarrow$ need to reduce memory usage.

CUDA Challenges

Initial Complications

- Class inheritance structure captures filter specialism for each sub-detector.
- Dynamic creation of objects in the main routine.
- Track state retained at each filtering step.
- Main routine has over 2000+ lines of code with multiple branches.

Feasibility Studies

- Standalone version successfully ported to C.
- Pre-allocated memory needed for track objects.
- Promising results $\Rightarrow$ need to reduce memory usage.

Kalman Filter Potential

- Our present Kalman Filter could be modified.

GPU benefits at other experiments

- Kalman Filter port to CUDA (GSI Scientific Report 2008, FAIR-EXPERIMENTS-38)
- ALICE TPC HLT code GPU based / Future PANDA TPC code
- GPUs to be used for STS (Silicon Tracking System) within CBM (Compressed Baryonic Matter) experiment at FAIR/GSI.

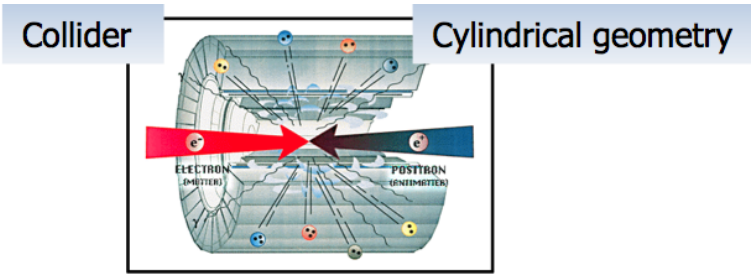

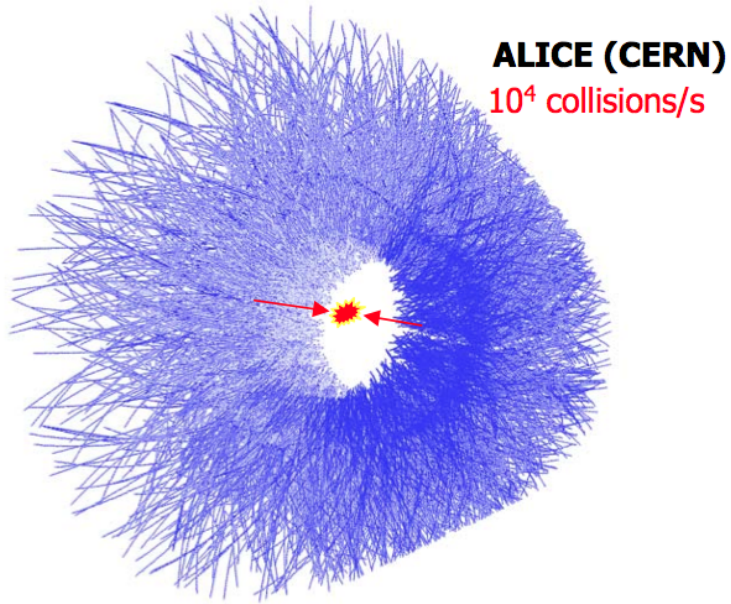

Summary

- The ATLAS trigger, particle tracking & simulation algorithms are key places where GPUs can be used to improve performance.
- Preliminary results show substantial performance.
  - Initial *32x* speed-up for parallel RK4 integration.
  - With optimisation up to *35x* speed up for Level2 Z Finder.
  - Initial port of OO based Kalman Filter algorithm.

Further information


SIMT design of the High Level Trigger Kalman Fitter


Porting the Z-finder algorithm to GPU
ATLAS Edinburgh GPU Computing


LHC and ATLAS papers
2008 JINST 3 S08003

Thanks to Peter Jenni, Iain Longstaff for material.
Thanks to NVIDIA for their support