NEW APPROACHES FOR FAST AND EFFICIENT GRAPH CONSTRUCTION ON CPU / GPU AND HETEROGENEOUS ARCHITECTURES FOR THE ATLAS EVENT RECONSTRUCTION

ATLAS detector

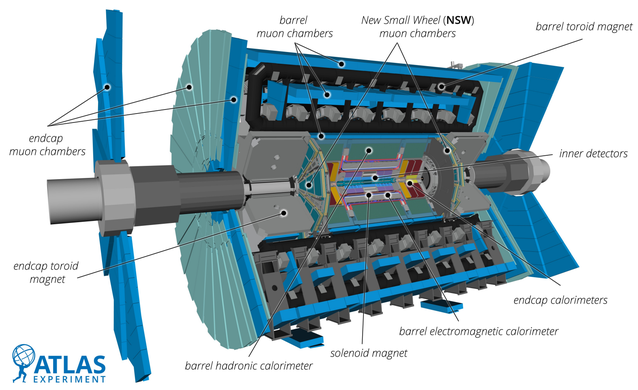

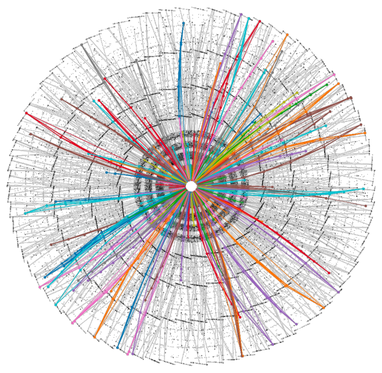

18108 modules in the ATLAS ITK detector (for HL-LHC 2029)
average number of space points / event: O(300k)

Track reconstruction challenge for HL-LHC

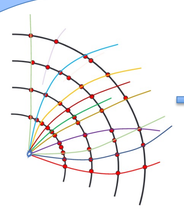

Charged particles leave hits in the detector

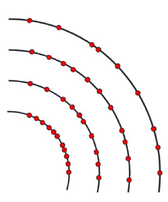

Hits are collected by Atlas Tracker

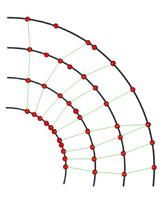

Goal: rebuild particle tracks

- large amount of data
- need fast reconstruction
- machine learning ?

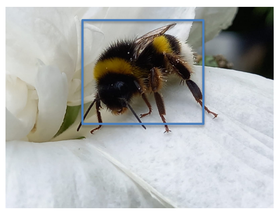

1600 * 1200 pixels
a large fraction of the image carries information of the subject

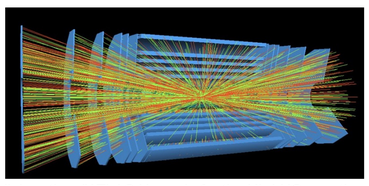

Atlas Tracker for HL-LHC
5 * 10^9 readout channels
~ 3 * 10^5 3D space points / event
data are sparse

The ATLAS collaboration, "ATLAS Upgrade for the HL-LHC: meeting the challenges of a five-fold increase in collision rate", ATL-UPGRADE-PROC-2012-003, EPJ Web Conf. 28 (2012) 12069

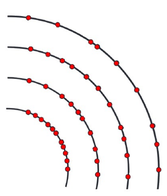

Hits are collected by Atlas Tracker

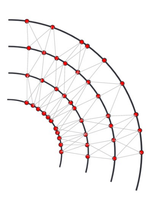

Represent data using a graph

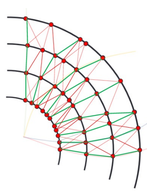

Classify the edges of the graph

- High classification score
  likely edge is part of the track
- Low classification score
  edge is probably NOT part of the track

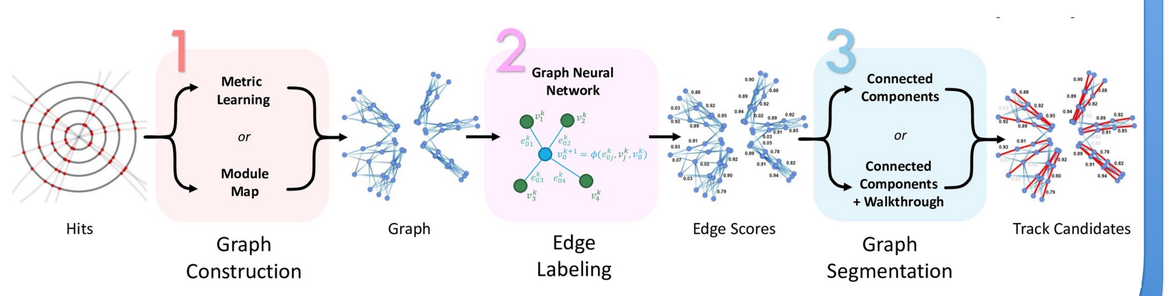

GNN4ITK Full pipeline

The ATLAS Collaboration, "Track finding performance plots for a Graph Neural Network pipeline on ATLAS ITk Simulated Data", IDTR-2022-01, https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/IDTR-2022-01/ (2022)

Module Maps

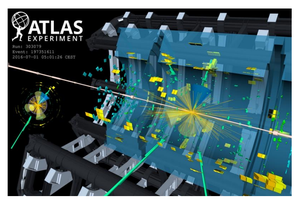

Module Maps are computed from hits (energy deposits) left in the detector by particles based on several events

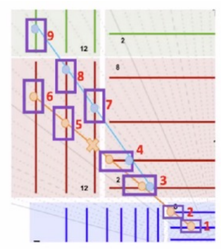

Hits in the active sensors of the detector

For each event

Map built using a large number of simulated events

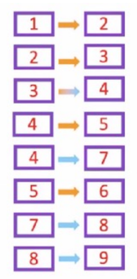

Module Map Doublet

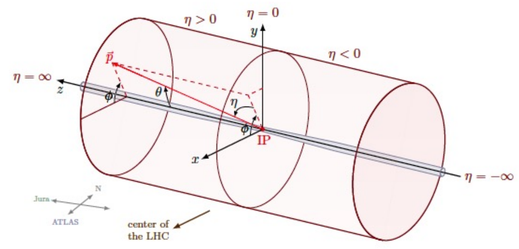

Module Map Triplet

relation hits <-> particles

several events used to create a map of energy deposits in the detector

For each triplet of modules calculate geometrical cuts

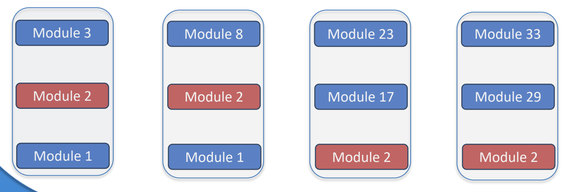

Module Map Doublet

calculate geometrical cuts on doublets as min and max values of triplet cuts
→ *deduced from Module Map Triplet*

Graph creation

Module Map gives interconnected modules

Connect hits

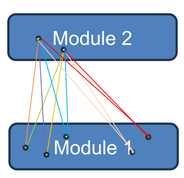

connect all the hits of Module 1 and 2 together

Apply doublet cuts

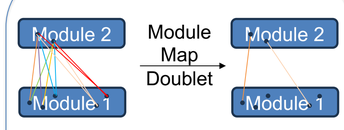

over 90% of triplet possible connections killed

much faster than triplet selection (no combinatorial)

allows to reduce triplet selection time

Apply triplet cuts

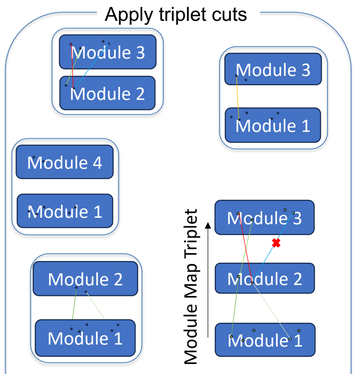

over 50% of remaining connections killed

Graph creation : from CPU to GPU

Step 1: Doublet cuts
- send event data to GPU
- apply doublet cuts
- sort edges on output hit (quick sort)

Step 2: Count edges after doublet cuts and reduce data
- parallel scan
- compact stream

Step 3: Triplet cuts

Step 4: Data reduction
- active edges only
  parallel scan + stream compaction
- copy reduced data to cpu
- free cuda memory

Parallel scan

hits	hit 01	hit 02	hit 03	hit 04	hit 05	hit 06	hit 07	hit 08	hit 09	hit 10
tag value before cuts	1	1	1	1	1	1	1	1	1	1
tag value after MMD cuts	1	0	1	1	0	1	0	0	1	1
scan function	1	1	2	3	3	4	4	4	5	6

position of the elements in the reduced vector

Compact stream

hits	hit 01	hit 02	hit 03	hit 04	hit 05	hit 06	hit 07	hit 08	hit 09	hit 10
reduced hits	hit 01	hit 03	hit 04	hit 06	hit 09	hit 10				
reduced R	R01	R03	R04	R06	R09	R10				

+ other data

GPU profiling and computational time

TTree_hits_constants
Section: Occupancy

Metric Name	Metric Unit	Metric Value
Block Limit SM	block	16
Block Limit Registers	block	3
Block Limit Shared Mem	block	16
Block Limit Warps	block	2
Theoretical Active Warps per SM	warp	32
Theoretical Occupancy	%	100
Achieved Occupancy	%	91.23
Achieved Active Warps Per SM	warp	29.19

TTree_hits_constants

Metric Name	Metric Unit	Metric Value
DRAM Frequency	cycle/nsecond	6.02
SM Frequency	cycle/nsecond	1.15
Elapsed Cycles	cycle	703,443
Memory Throughput	%	1.15
DRAM Throughput	%	1.15
Duration	usecond	609.38
L1/TEX Cache Throughput	%	0.56
L2 Cache Throughput	%	0.71
SM Active Cycles	cycle	656,862.85
Compute (SM) Throughput	%	80.70

% GPU time
chart of the 69ms by function

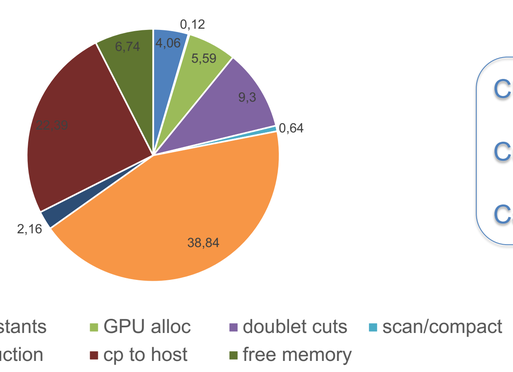

CUDA kernel for graph construction

Currently being integrated in ACTS

Can be integrated to Athena

Comparison CPU vs GPU

differences on 100 events data
nb nodes 22 / 27,384,853
nb of edges 446 / 177,067,905

Memory ressources

10 to 11 Go / event

Average GPU time

69ms / event

based on 100 random events of CTD 2023 (simulated pile up of 200 ttbar events) dataset running on Nvidia A100

COLLARD Christophe: L2IT, Toulouse, France, christophe.collard@l2it.in2p3.fr