Charged Particle Tracking: From the LHC to the Muon Collider

Ben Rosser

University of Chicago

December 2, 2025

Introduction and Outline

- I (mostly) work on ATLAS; topics I won't cover today:
  - Higgs boson decay to invisible particles
  - Long-lived particles (displaced vertices and $E_{\mathrm{T}}^{\mathrm{miss}}$, displaced taus).
  - Axion physics interpretation of LLP searches (2511.07224).
  - Front-end electronics for HL-LHC Inner Tracker Strip detector.

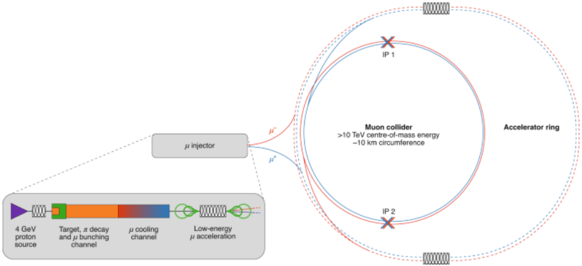

- This talk: can we **reconstruct tracks** at a $10\,\mathrm{TeV}$ **muon collider**?
  - (Brief) introduction to muon colliders and detector challenges.
  - Why tracking is particularly important.
  - General overview of tracking at colliders.
  - Current ATLAS tracking R&D for the **Event Filter** trigger upgrade.
  - Overview of various tracking studies in the muon collider community.

Why Collide Muons?

- Combine advantages of e^+e^- with pp: best of both worlds?
  - Like electrons, **fundamental** particle: all beam energy goes into collision.
  - Muons **200x** heavier than electrons: less synchrotron radiation.
  - $10\,\mathrm{TeV}$ muon collider roughly **equivalent** to $100\,\mathrm{TeV}$ proton collider (FCC).
  - But much smaller footprint, close to LHC: lower cost, lower power.
- ...except they are **unstable**: $2.2\,\mu\mathrm{s}$ lifetime.
- Lots of R&D needed to demonstrate feasibility; especially on the accelerator.

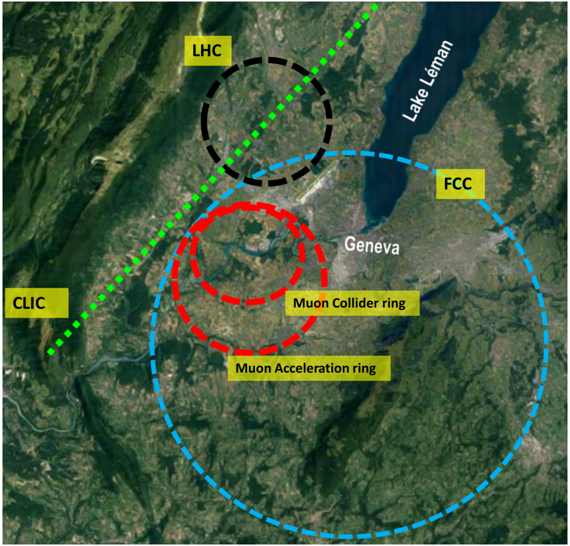

2024 IMCC Annual Meeting

- LHC had Higgs boson "no-lose theorem". No such case for future colliders; why 10 TeV?
  - Can nail down shape of the Higgs potential, achieve strong Higgs precision (2206.08326).
  - 5σ discovery potential for some minimal WIMP dark matter models at correct thermal target.
- Muon colliders **become VBF colliders**: notion of "electroweak PDF" emerges.
  - s-channel interactions (dashed lines) fall with $\sqrt{s}$; electroweak interactions become dominant.

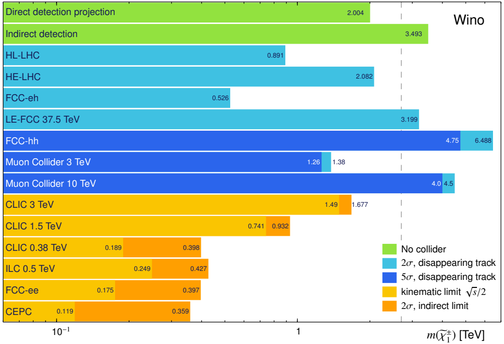

R. Capdevilla et al. (10.1007/JHEP06(2021)133)

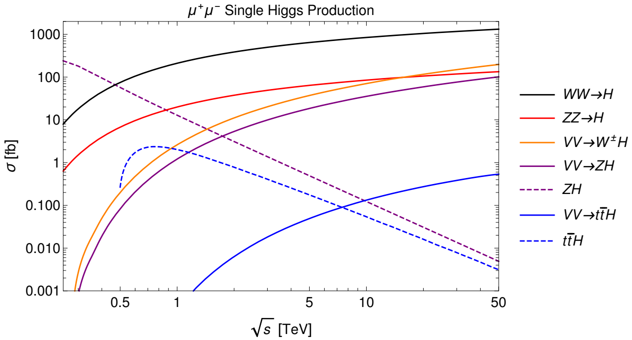

M. Forslund, P. Meade (10.1007/JHEP08(2022)185)

Beam Induced Background

- Main detector challenge: muons are unstable!
  - $5\,\mathrm{TeV}$ muons produce multi-TeV electrons.
  - Time dilation: $\tau \propto E$.

$$E_{\text{decay}} = (13\,\mathrm{EeV})\left(\frac{n_\mu/\text{bunch}}{2\times 10^{12}}\right)$$

- Must shield detector:
  - Optimize accelerator lattice, **deflect** particles.
  - Surround beam pipe in **tungsten nozzles**.
  - High energy electrons shower: 10^8 low-energy γ, n, $e^\pm$ enter detector.

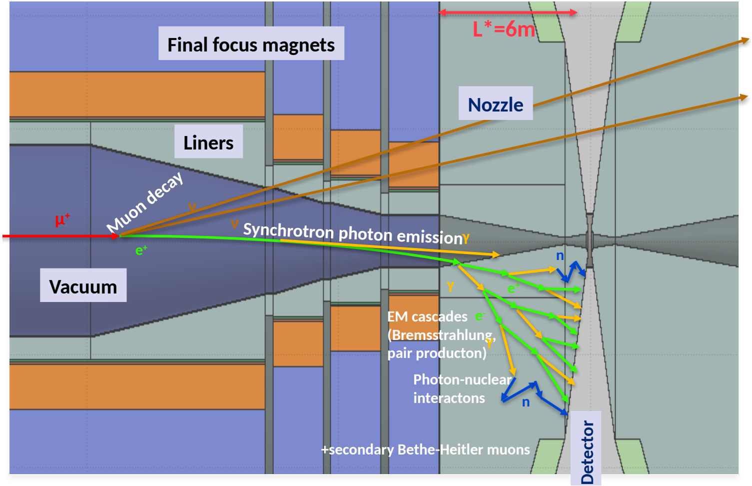

Daniele Calzolari

BIB after Shielding

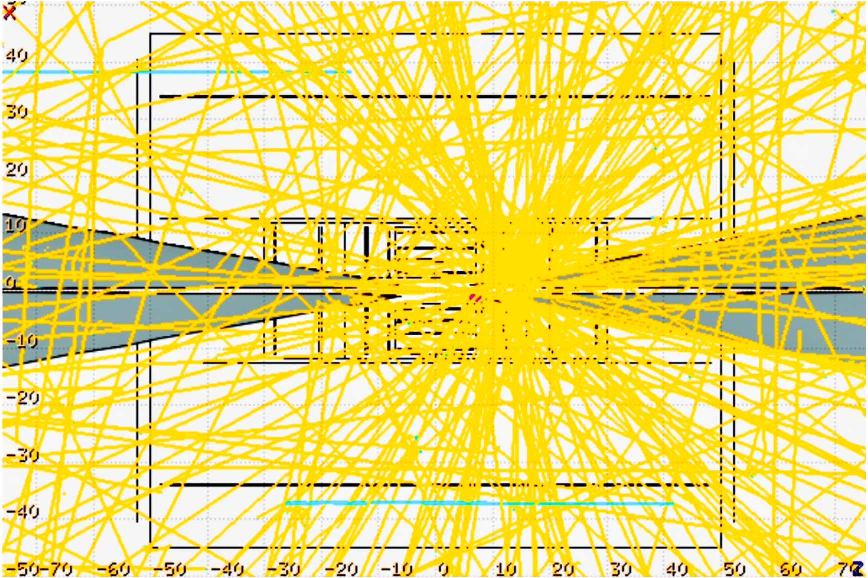

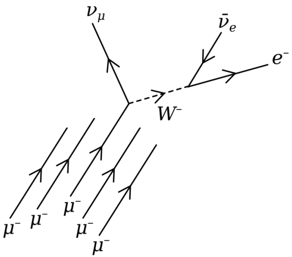

- Photons
- Electrons
- Positrons
- **0.0003%** of one BIB event!

- Machine-detector interface, lattice design **completely determines** BIB in detector.
- Lattice needs to meet key challenges:
  - Very small $\beta^* = 1.5$ (beam size at IP).
  - Large relative energy spread.
  - Short bunch length ($\sigma_z = 1.5$).
  - Radiation damage to magnets from BIB.
  - Limited straight sections to mitigate **neutrino radiation**.
- Current target: collider lifetime of **5-10 years**; aim to collect $10\,\mathrm{ab}^{-1}$ in 5 years.
- Studies performed with two designs:
  - v0.4: initial design (2502.00181).
  - v0.8: **more realistic**, now baseline for 2025 European Strategy.

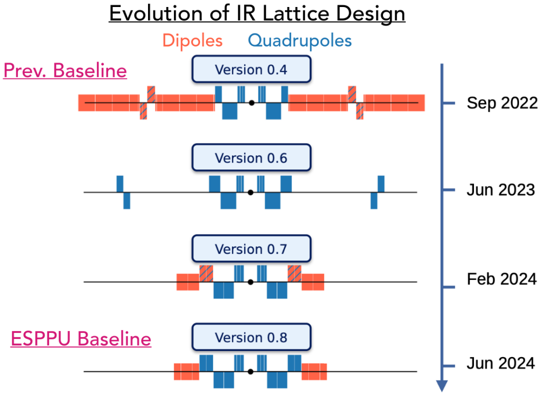

Kiley Kennedy (CPAD 2024), Marion Vanwelde

MAIA Detector Concept

- Detectors must be designed to handle BIB.
- Community studying three designs:
    - $3\,\mathrm{TeV}$ CLIC baseline; two $10\,\mathrm{TeV}$ variants.
    - MUon System for Interesting Collisions: **MUSIC** (2511.23273).
    - Muon Accelerator Instrumented Apparatus: **MAIA** (2502.00181).
- Main difference: solenoid placement.
    - MAIA: $5\,\mathrm{T}$ solenoid **inside calorimeters**.
    - $4X_0$ of material; shielding from BIB.
    - MUSIC: solenoid between ECAL, HCAL.
- Other updates for $10\,\mathrm{TeV}$:
    - Partially **reoptimized** tracker geometry.
    - Calorimeter **depth** increased.

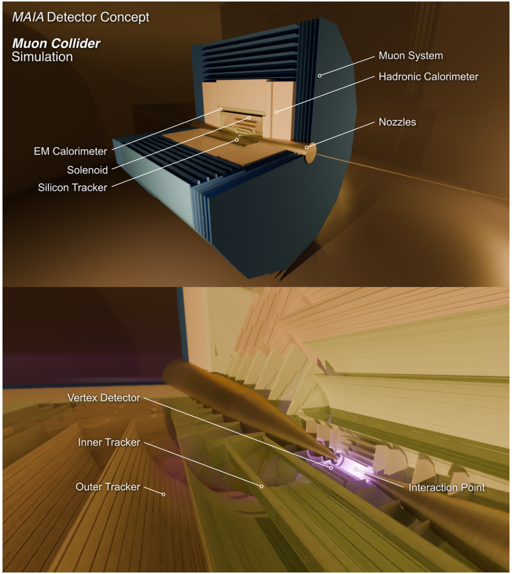

MAIA Performance Studies

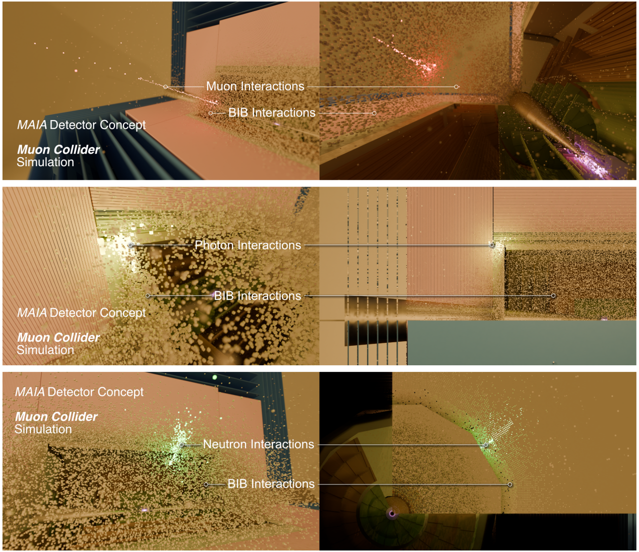

- Detector performance evaluated using **single particle** samples:
  - Muons for the tracker: flat in p_T.
  - Photons the ECal: flat in E.
  - Neutrons for the HCal: flat in E.
- Reconstruction using key4hep-based MuonColliderSoft framework.
- BIB overlaid during digitization.
- Reco uses **Pandora** particle flow:
  - Optimized for e^+e^- environment; needs tuning for $\mu^+\mu^-$.

Studies with Beam Induced Background

- BIB simulated in FLUKA: reports particle flux **entering** detector.
  - Very low-energy; many particles **out of time**: timing important to suppress BIB.
- Need to run GEANT to model BIB interactions with detector:
  - Extremely **resource intensive**: one BX takes a month!
  - Simulate tenths of BX; randomly sample and combine to create independent pseudo-events.

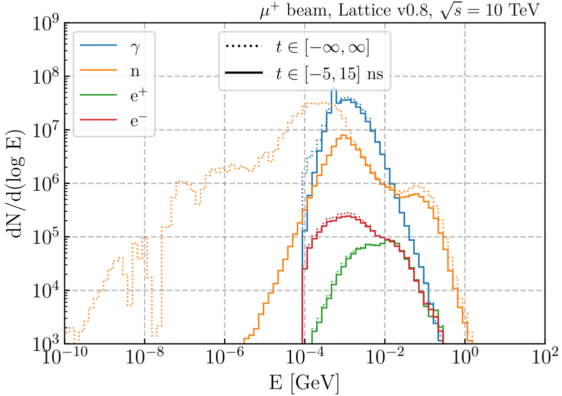

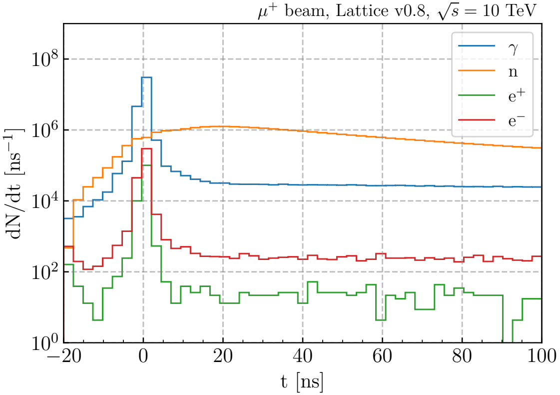

- Radiation at $10\,\mathrm{TeV}$ **comparable to HL-LHC** and previous $3\,\mathrm{TeV}$ muon collider studies; much lower than FCC-hh ($10^{18}\ 1\,\mathrm{MeV}\text{-}\mathrm{n_{eq}}/\mathrm{cm}^2$) (2209.01318, 2105.09116)

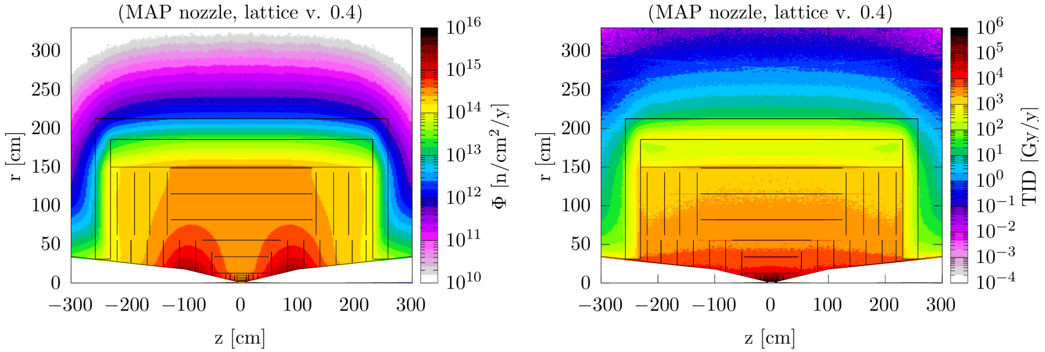

	Maximum Dose (Mrad)		Maximum Fluence (1 MeV-neq/cm^2)	
	R= 22 mm	R= 1500 mm	R= 22 mm	R= 1500 mm
Muon Collider ($3\,\mathrm{TeV}$)	10	0.1	10^{15}	10^{14}
HL-LHC	100	0.1	10^{15}	10^{13}
Muon Collider ($10\,\mathrm{TeV}$)	20	0.2	3×10^{14}	10^{14}

Expected Data Rates

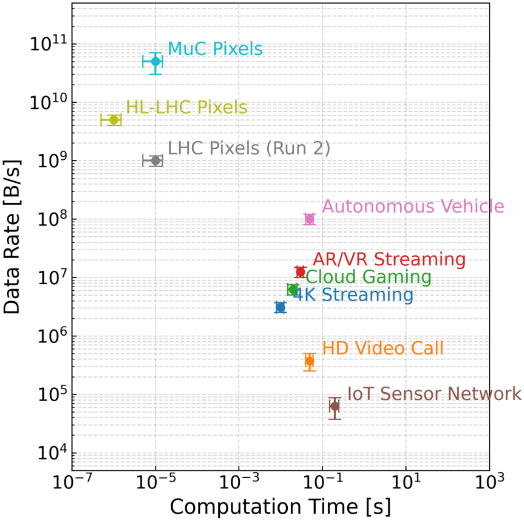

Eliza Howard, Karri Di Petrillo

- Muon collisions 10^3 **slower** than LHC:
  - Only **one bunch** per beam.
  - $30\ \mu s$ collision rate.
  - Read out **every event**: ideally no trigger.
- Much more detector activity **per collision**:
  - ATLAS tracker: O(0.1%) occupancy.
  - Muon collider tracker: **20%** occupancy.
  - Huge uncertainties: accelerator, MDI.
- Multiple challenges:
  - Front-end electronics: in-pixel pileup.
  - Readout link capacity.
  - Ultra-high occupancy track reconstruction.

Tracking at ATLAS

Tracking at Colliders

- Measure trajectory of particles:
  - Charged particles bend in B field.
  - Circular motion in x-y; linear motion in z-r.
  - **Helical** trajectory: five parameters.
  - Many different parametrizations; ATLAS commonly uses "perigee": $(d_0, z_0, p_\mathrm{T}, \eta, \phi)$.
- Make many measurements of position; fit helix.

- **ACTS**: A Common Track Reconstruction Software (2106.13593).
  - Based on ATLAS.
  - Experiment agnostic.
  - Used for future collider studies.

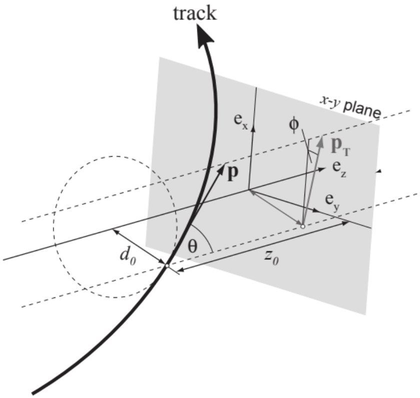

ATLAS Tracking Tutorial

Silicon Sensors

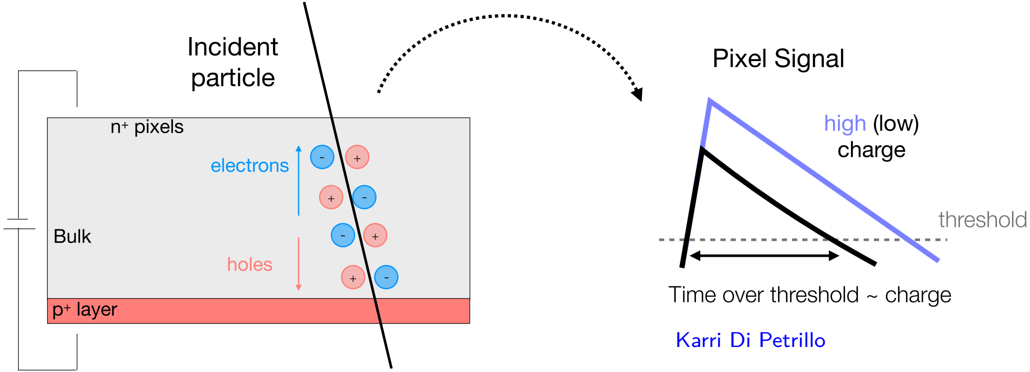

- Silicon trackers preferred due to spatial resolution, timing, radiation hardness.
- Charged particles create e-h pairs; typically 80/micron for minimum ionizing particle.
- Absorb electrons, measure time over charge threshold; digitize as a "hit".

Tracking Detectors

- Basic design constraints: maximize granularity, minimize occupancy.
  - Prioritize precision in x-y plane; maximize "lever arm" for p_T reconstruction.
  - Minimize distance to beam line for d_0 / z_0 resolution and vertexing.

ATLAS Tracking Tutorial

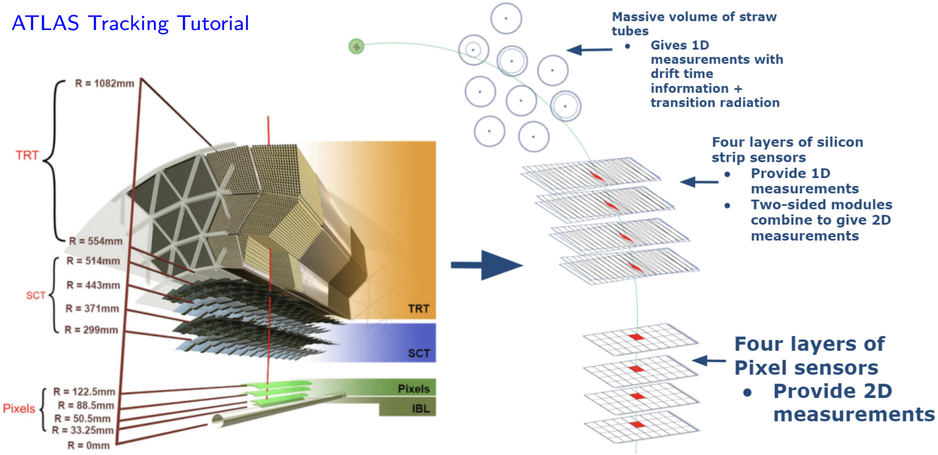

Track Seeding and Pattern Recognition

- ATLAS approach for tracking: key steps:
  - Cluster digitized hits into **spacepoints**.
  - Find **triplet seeds** from three spacepoints.
  - Extend using combinatorial **Kalman filter**.
- Seed-finding strategies:
  - Standard is to perform simple circle fit.
  - Other approaches, e.g. machine learning (**G**raph **B**ased **T**rack **S**eeding).
  - Remove any **overlapping seeds**.
  - Impose quality requirements (e.g. on d_0, z_0).
- Find **search roads**:
  - Preselect hits for use in extrapolation.

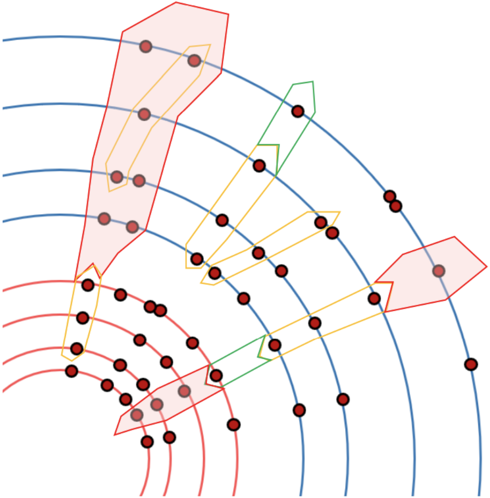

ATLAS Tracking Tutorial

Tracking with Kalman Filters

- Combine multiple discrete noisy measurements.
- Given a model of a system and initial seed:
  - Predict state vector (track) in next layer.
  - Add compatible hits, recalculate track.
  - Repeat until all layers have been searched!
  - Propagate backwards to smooth trajectory.
- Combinatorial Kalman Filter: branching.

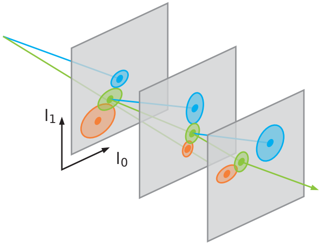

ACTS Common Tracking Software

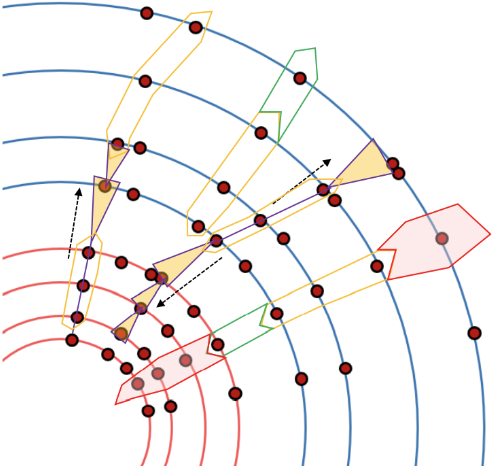

ATLAS Tracking Tutorial

Track Fitting and Ambiguity Resolution

- Rejecting **fake tracks**: common issue:
- Different strategies for ambiguity resolution:
  - Perform global χ^2 fit; prefer lower χ^2.
  - Remove outlier measurements and refit.
  - Overlap removal for tracks sharing hits.
  - Prefer tracks that have more hits (i.e. aren't missing a hit in a given layer).
  - Prefer tracks with higher p_{T}, pass basic quality requirements on parameters.

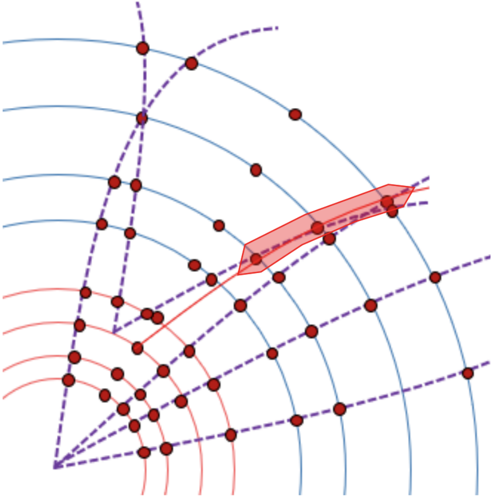

ATLAS Tracking Tutorial

Benchmarking at ATLAS

- Most expensive reconstruction step at the LHC! Worse at high pileup: more activity.
- Two questions for muon collider: will tracking work; will it work **in finite time**?

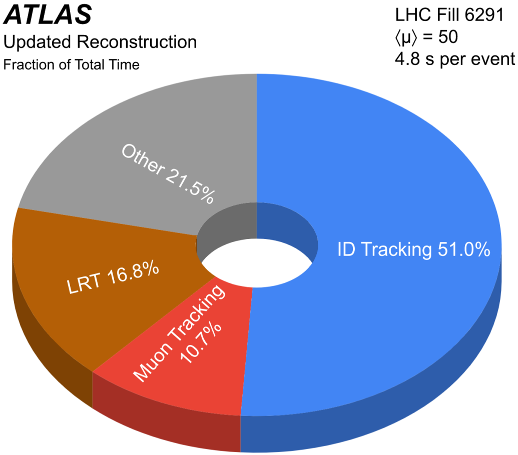

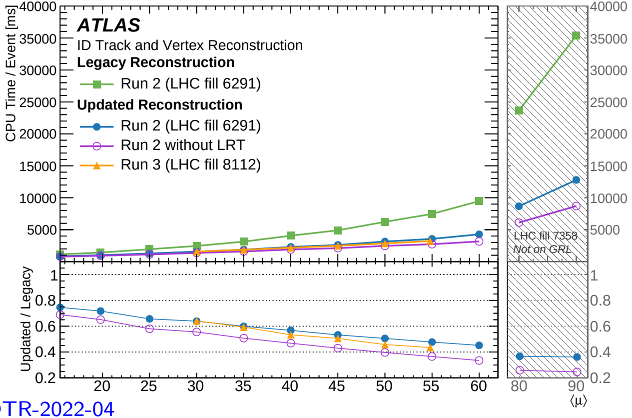

IDTR-2022-04

HL-LHC Collisions

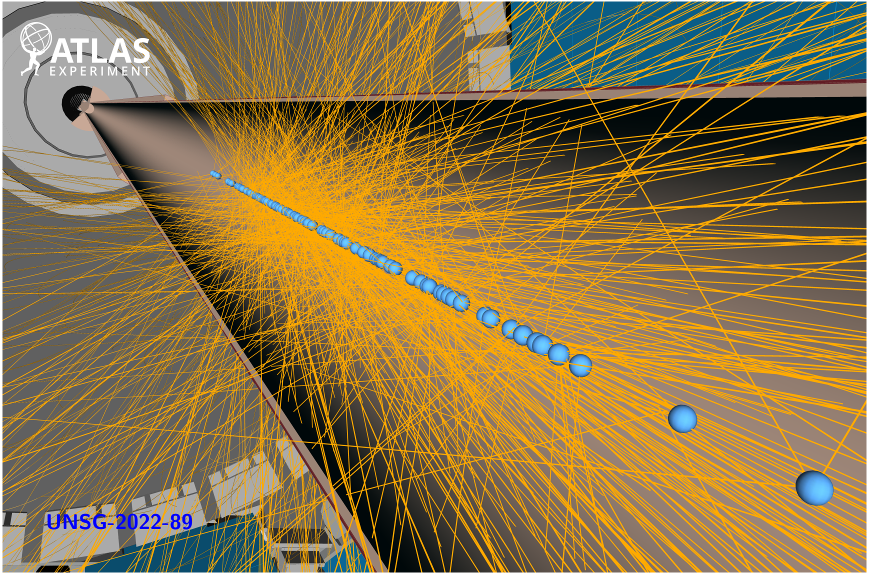

- $t\bar{t}$ event:
  - $\langle\mu\rangle = 200$ pileup.
  - 88 primary vertices.
  - **2000+** tracks.
- ITk simulation.

High-μ Events in Run 3

- Real event from high-μ run:
  - 150 pp interactions.
  - 92 primary vertices.

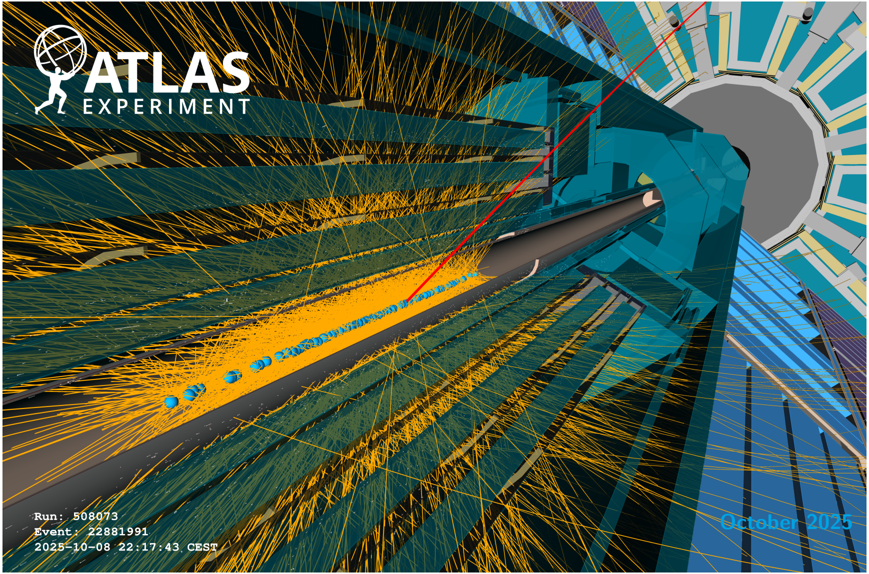

Tracking in the ATLAS Trigger

- Compute time already challenge in **ATLAS trigger**.
- For HL-LHC, target trigger rates:
  - $40\,\mathrm{MHz} \rightarrow 1\,\mathrm{MHz}$ **Level 0** (hardware)
  - $1\,\mathrm{MHz} \rightarrow 10\,\mathrm{kHz}$ **Event Filter** (CPU)
  - **3-10x** increase in rate compared to today.
  - $O(\mathrm{ms})$ to process events in EF; $O(\mu\mathrm{s})$ in L0.
- ATLAS only planning on tracking in EF:
  - Concern about **power consumption**.
  - Available power for EF server farm: $2.5\,\mathrm{MW}$ total; only $1.8\,\mathrm{MW}$ for tracking.
  - Running at required rates **risk exceeding limit**.
- Solutions:
  - More power: financially and ecologically not ideal.
  - Alternative algorithms?
  - Heterogeneous computing?

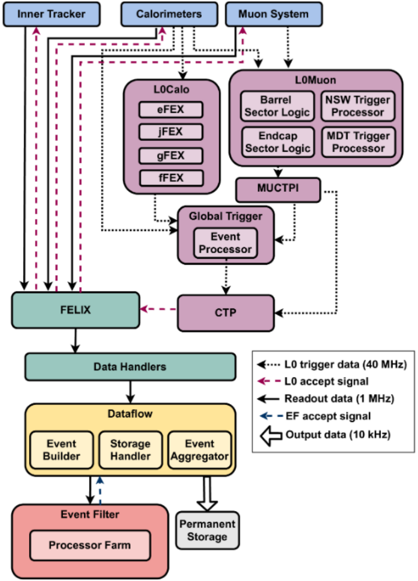

Event Filter Tracking

- ATLAS set up R&D project, running since **2021**, to study options:
  - Compare optimized CPU-based tracking to **accelerator** cards.
  - Focusing on commodity hardware: CPU-only vs CPU plus GPU/FPGA.
- Wide range of options explored.
  - Minimal pipelines: only run **clustering** on **on accelerator** as proof-of-concept, rest **on CPU**.
  - Maximal pipelines: run as much as possible on GPU or FPGA.
- FPGAs: lower **power and latency**, but harder to program, wrt CPUs.

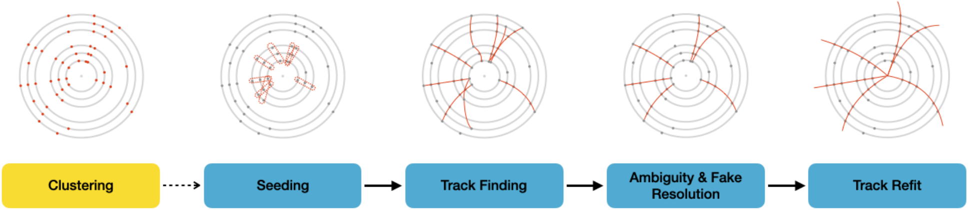

PUB-TDAQ-2025-02

FPGA Pipelines

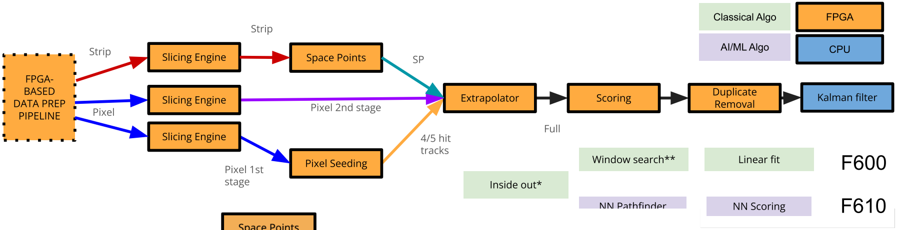

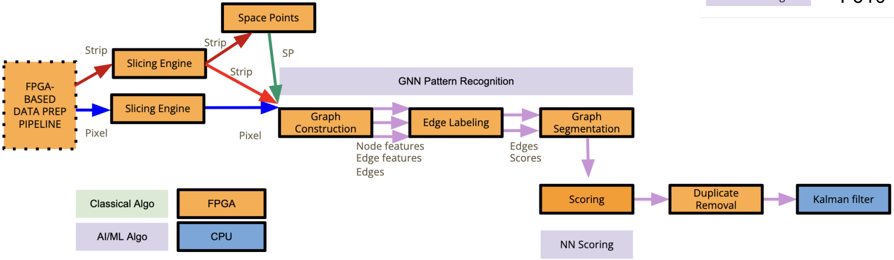

- EF Tracking process selecting **technology**, not final trigger system.
  - For FPGAs: three maximal pipelines built on top of minimal "data preparation" (F-100).
  - Classical (F-600); semi-classical (F-610); and full ML (F-410).

Requirements for Pipelines

- All pipelines must meet **physics** and **technical** requirements to be considered.

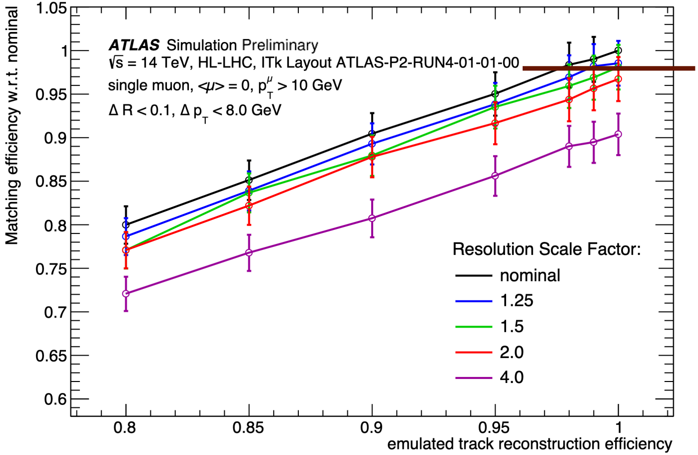

ATL-COM-DAQ-2023-071

- Set **minimum** requirements:
  - Established goal of **98%** single muon efficiency, match run 3.
  - Looked at impact of tracking on muon trigger efficiency.
  - Set 98% tracking efficiency **relative to offline** as requirement.
  - Stable even if track p_T significantly mismeasured.
- Focus on conventional (**prompt**) signatures for technology choice.

Data Preparation Performance

- Clustering followed by ACTS based **fast tracking**. Implemented in firmware!
- Efficiency not perfect compared to offline, but matches CPU-only fast tracking.

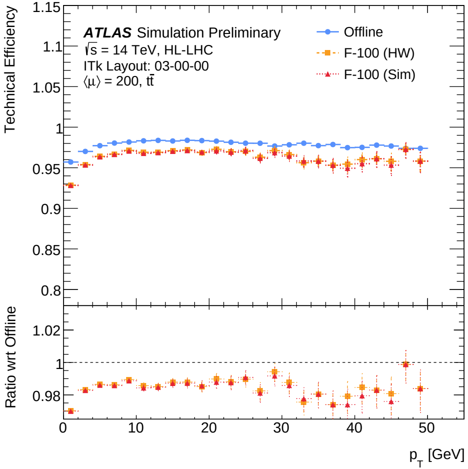

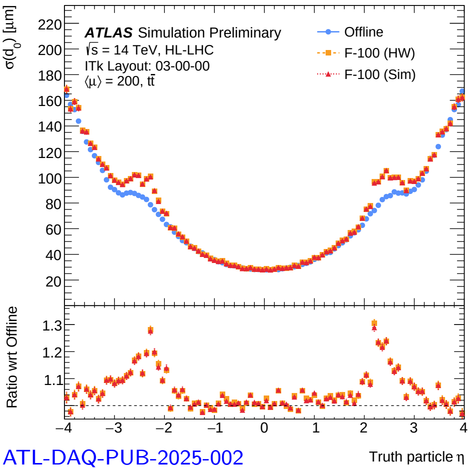

ATL-DAQ-PUB-2025-002

Maximal Pipeline Performance

- Maximal FPGA pipelines: challenging to implement! Only tested in simulation.
  - Classical pipeline: great efficiency, at cost of completely **unoptimized** fake rate.
  - Pure-ML pipeline: great fake rate, but worse efficiency, especially in the barrel.

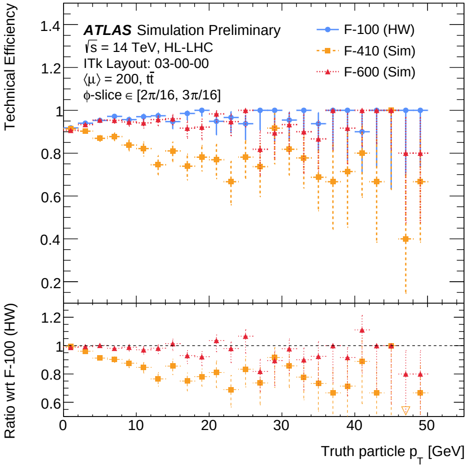

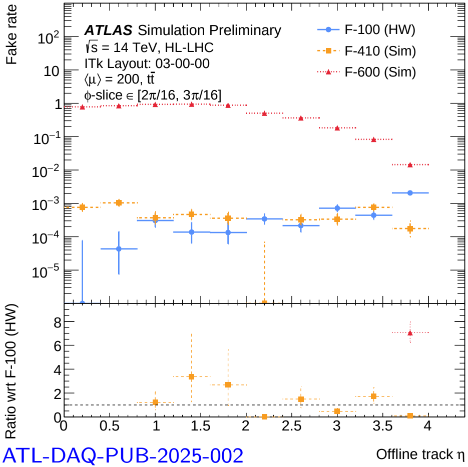

ATL-DAQ-PUB-2025-002

Current Status and Lessons Learned

- Currently preparing **final** inputs to ATLAS technology choice decision process:
  - Due December 12th!
  - Refocused on intermediate FPGA option: seeding on FPGA, CKF on CPU.
  - Decision expected in **February 2026**.
- Heterogeneous computing for tracking: it works, but many challenges to overcome!
  - Challenges with both GPUs and FPGAs.
  - GPU-based data preparation worked; also issues with end-to-end GPU pipeline.
  - Can't match decades of work on CPU algorithms in just a few years.

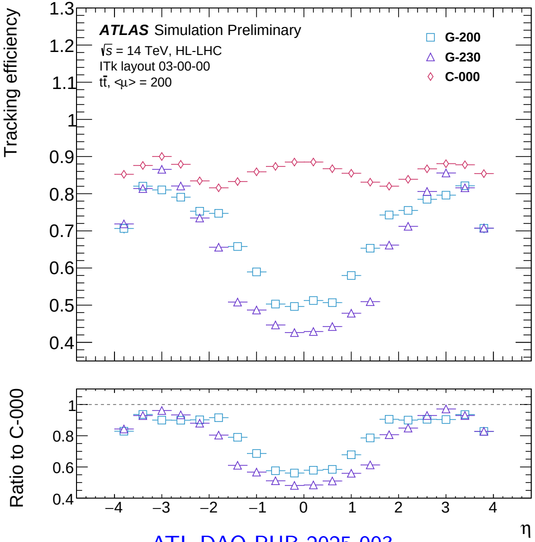

ATL-DAQ-PUB-2025-003

Muon Collider Studies

MAIA Tracker Layout

	Vertex Detector	**Inner Tracker**	**Outer Tracker**
Sensor Type	pixels	macropixels	macropixels
Layers, Barrel (Endcap)	4 (4)	3 (7)	3 (4)
Cell Size	$25\mu m \times 25\mu m$	$50\mu m \times 1mm$	$50\mu m \times 10mm$
Sensor Thickness	$50\mu m$	$100\mu mm$	$100\mu mm$
Time Resolution	30ps	60ps	60ps
Spatial Resolution	$5\mu m \times 5\mu m$	$7\mu m \times 90\mu m$	$7\mu m \times 90\mu m$

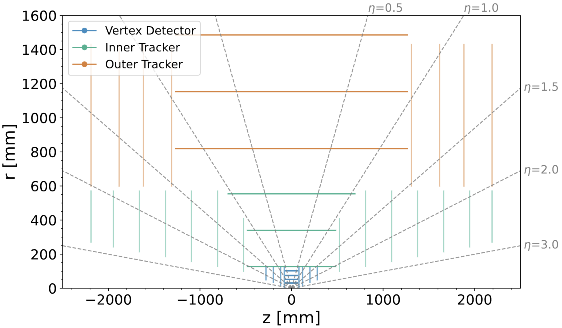

- Baseline tracker inherited from CLIC.
  - Original design had double-sided layers.
  - All but one removed to save power, reduce inefficiencies for LLPs.
- Conservative technology assumptions:
  - Compare to HL-LHC LGADs: same timing resolution but 1mm × 1mm.
  - Outer tracker **could** be silicon strips if technically feasible.

Tracker Occupancy and Reconstruction Time

- Very high hit densities in vertex: 1600 hits/cm^2 in innermost with original lattice (v4).
  - Higher with new lattice (v8); incoherent e^+e^- pairs also now included, but only 20% effect.
  - $[-3\sigma, 5\sigma]$ beam crossing timing window eliminates much of the BIB.
- Track reconstruction time has ranged from minutes (v4) to **a few hours** (v8) per event.
  - Has proven necessary to only run tracking **near truth particles**: not feasible long-term!

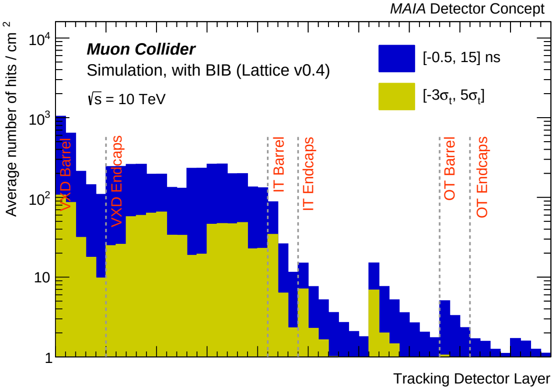

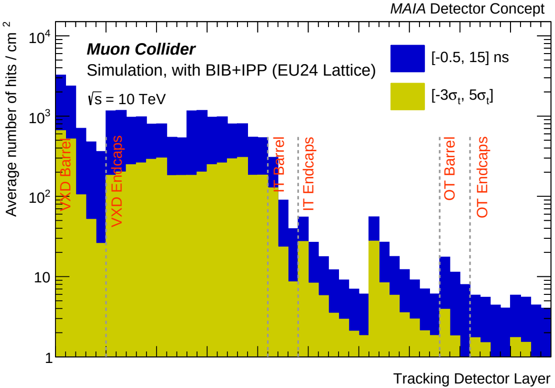

Tracking in MAIA

- Study tracking with single muon samples.
  - Efficiency with BIB very high, except in forward region around nozzles: needs optimizing.
  - Apply cleaning selection ($p_\mathrm{T} > 1\,\mathrm{GeV}$, $N_\mathsf{hits} > 5$, $\chi^2/\mathrm{DOF} < 3$) to reject fakes.
  - End up with around 90-95% efficiency in the barrel; promising, but room for improvement.

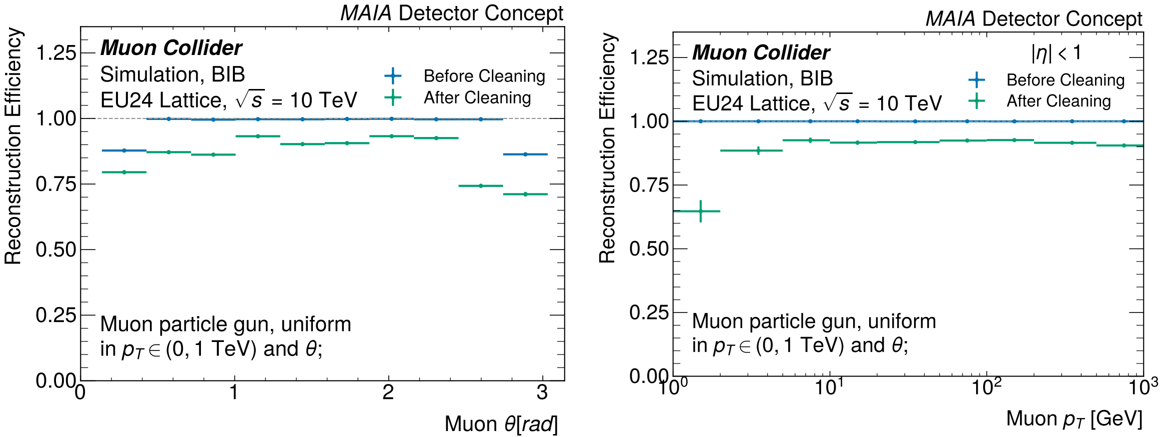

Track Parameter Resolution

- p_{T}, d_0 resolution shown as function of θ for different p_{T} muons.
- Generally good parameter resolution for both low- and high- p_{T} tracks, even with BIB!

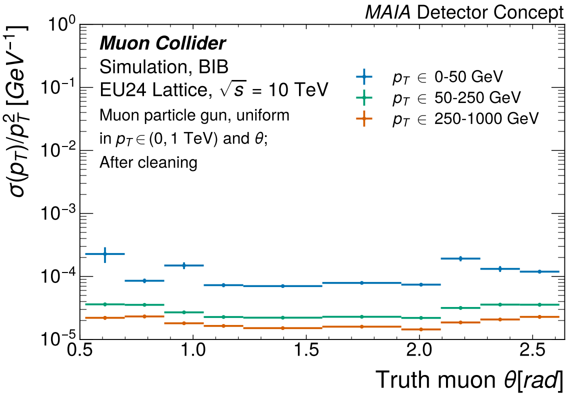

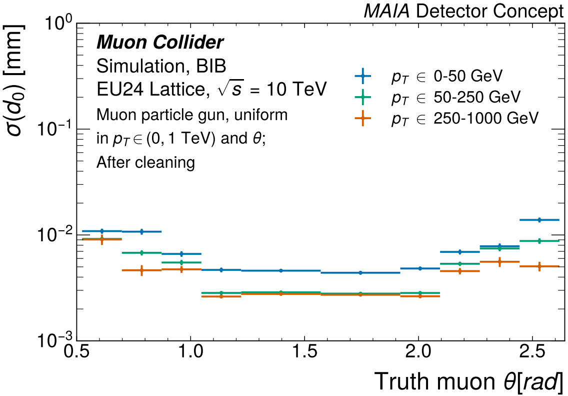

Track Quality Issues

- Despite good performance: some problems to debug with current setup.

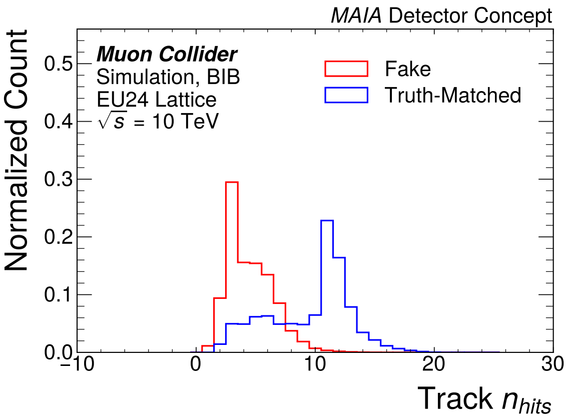

- ACTS finds tracks with **less than 3** hits:
  - Can't fit a helix to fewer than three hits.
  - Seeds start out as three hits.
  - May suggest we fail to extrapolate some seeds, and hits get marked as outliers.
- Outlier removal in general needs work:
  - Cases where clear outliers get added to a track, not removed.
- Reconstruction actively evolving!
  - Currently using older ACTS v32; upgrade to latest v44 in progress.
  - Efforts being driven by Simone Pagan Griso (LBNL); Rocky Bala Garg (Stanford); Federico Meloni (DESY); et al.

Long Lived Particles

- Focus has been on getting **standard** tracking to work.
- Must ensure we can search for LLPs at a muon collider as well!
- One benchmark model: gauge-mediated symmetry breaking.

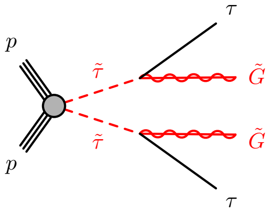

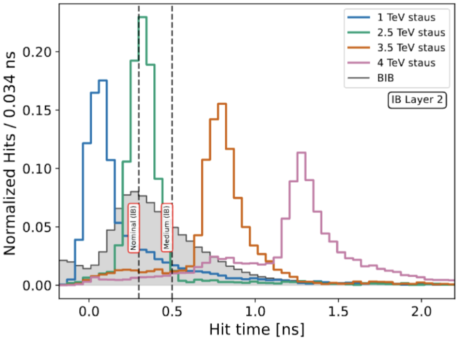

Mira Littmann, Mark Larson (Chicago)

- Massive stau NLSPs slowly moving: **heavy stable charged particles**.
- Problem for muon collider:
  - Timing cuts being used to suppress BIB.
  - $4\,\mathrm{TeV}$ stau won't reach end of inner tracker within time window.
  - Need to either relax timing cuts or find another BIB suppression method.
- Similar problems for **displaced** tracking.

Adjusting Time Windows

- Studied impact of looser timing windows on BIB occupancy, reconstruction efficiency:
  - Define "loose" windows that capture 90% of $4\,\mathrm{TeV}$ stau hits in all three subdetectors.
  - "Medium" windows roughly middle ground between "loose" and "tight" (default) window.
  - Looser windows let in about 2x more BIB, but necessary for heavy stable charged particles!
- Investigating **sliding** β-dependent timing windows: adjust minimum and maximum bound.

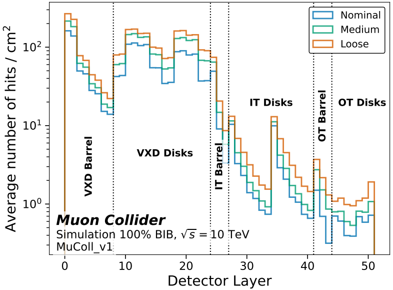

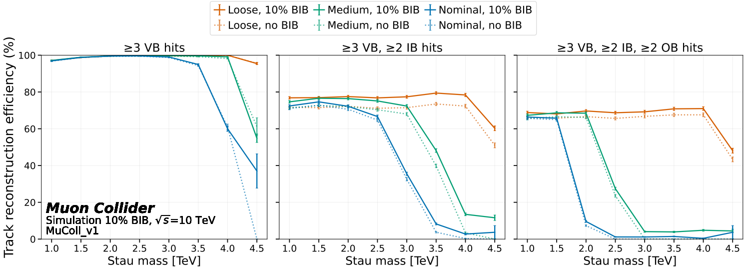

Mira Littmann, Mark Larson (Chicago)

Realistic Digitization

- Digitization step currently very simple:
  - Hit positions taken from GEANT, averaged with some Gaussian smearing.
  - Realistic digitization: take into account silicon sensors, front-end electronics.
- With realistic digitization model: access to additional information.
  - BIB mostly collinear to the beam; very different incident angle.
  - Using cluster shape information: can achieve 20-30% BIB rejection, 95% signal efficiency.
  - Alternative to timing cuts!

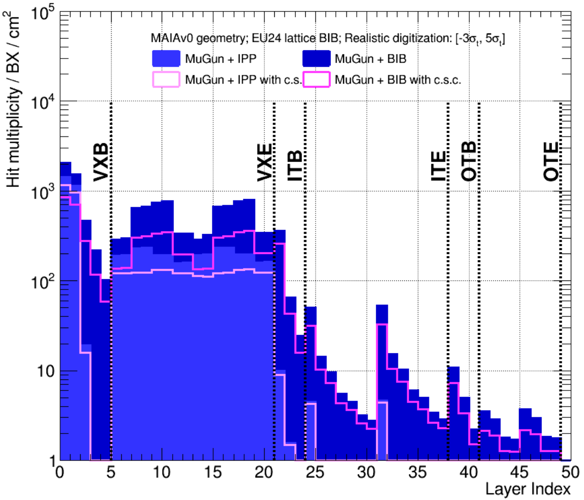

Angira Rastogi

- Very similar approach being studied as part of the Smart Pixels project:
  - Train neural network on pixel cluster information, embed in **front-end electronics**.
  - **Simplest** network (model 1): 95% signal efficiency with **80%** background rejection.

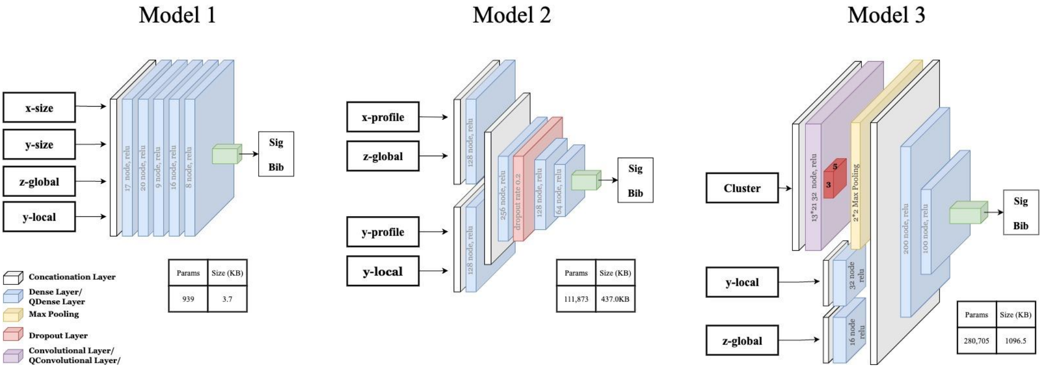

Eliza Howard

Line Segment Tracking

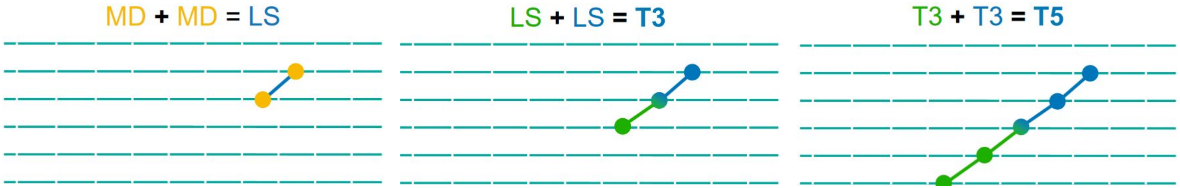

- New method from CMS:
  - Double-sided sensors in CMS outer tracker for HL-LHC treated as **single module**.
  - Form **mini-doublets** with hits from both sides.
  - LST: iteratively link longer and longer segments together.
- Fermilab, UCSD exploring muon collider applications.

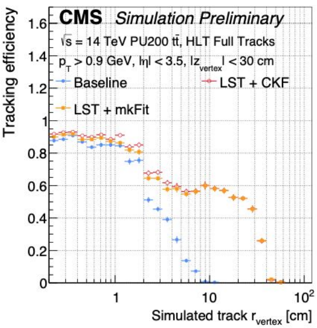

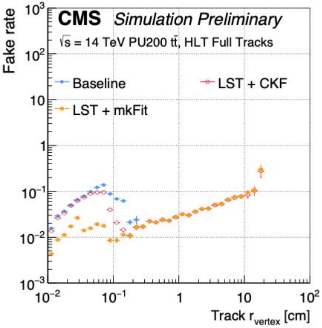

Yanxi Gu, Mario Masciovecchio

4D Tracking and Quantum Computing

- **4D** tracking: use timing instead of applying cuts.
  - Can also measure angles: "6D" (Victor Turbiner).
- Studies at DESY: **quantum** 4D tracking.
  - Quadratic unconstrained binary optimization.
  - Propagate 3D seeds using 1D, 3D, or 4D information.

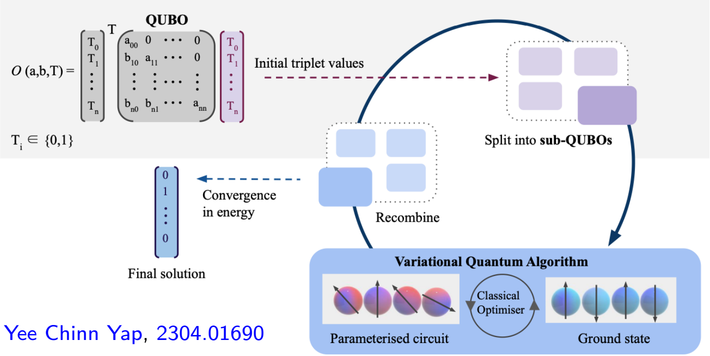

Yee Chinn Yap, 2304.01690

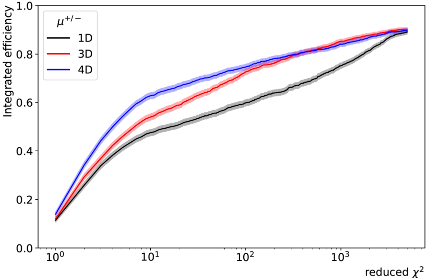

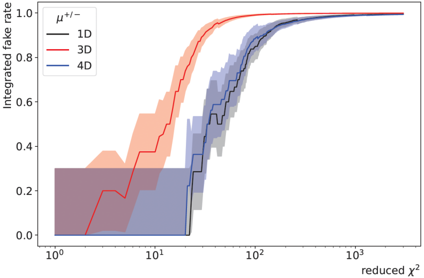

Conclusion

- Presented overview of track reconstruction challenges for a future muon collider.
- Lots of room for future improvement:
  - Sensor technologies.
  - (Endcap) geometry optimization.
  - New ideas for track finding.
  - 4D+ tracking.
  - Heterogeneous computing.
- Further reading: 2504.21417
- Muon collider **software tutorials**:
  - `https://mcd-wiki.web.cern.ch/software/tutorials/fermilab2024/`
  - Contact info for USMCC: `https://www.muoncollider.us/`.
- Thanks for your attention!

MAIA: A new detector concept for a 10 TeV muon collider

Charles Bell,[1] Daniele Calzolari,[2] Christian Carli,[2] Karri Folan Di Petrillo,[3] Micah Hillman,[1] Tova R. Holmes,[1] Sergo Jindariani,[4] Kiley E. Kennedy,[5] Ka Hei Martin Kwok,[4] Anton Lechner,[2] Lawrence Lee,[1] Thomas Madlener,[6] Federico Meloni,[6] Isobel Ojalvo,[5] Priscilla Pani,[6] Rose Powers,[5] Benjamin Rosser,[3] Leo Rozanov,[3] Kyriacos Skoufaris,[2] Elise Sledge,[7] Alexander Tuna,[1] and Junjia Zhang[5]

[1]*University of Tennessee, Knoxville, TN, USA*
[2]*European Organization for Nuclear Research (CERN), Switzerland*
[3]*University of Chicago, IL, USA*
[4]*Fermi National Accelerator Laboratory, IL, USA*
[5]*Princeton University, NJ, USA*
[6]*Deutsches Elektronen-Synchrotron DESY, Germany*
[7]*California Institute of Technology, CA, USA*

(Dated: November 10, 2024)

arXiv:2502.00181

Backup

Snowmass Implementation Task Force

- Smaller footprint leads to **lower cost** and **power** compared to hadron collider alternatives.
- Technical readiness and risk evaluated in detail during 2021 Snowmass process.

Proposal Name	CM energy nom. (range) [TeV]	Lum./IP @ nom. CME [10^{34} cm^{-2}s^{-1}]	Years of pre-project R&D	Years to first physics	Construction cost range [2021 B$]	Est. operating electric power [MW]
Muon Collider	10 (1.5-14)	20 (40)	>10	>25	12-18	$\sim$300
LWFA - LC (Laser-driven)	15 (1-15)	50	>10	>25	18-80	$\sim$1030
PWFA - LC (Beam-driven)	15 (1-15)	50	>10	>25	18-50	$\sim$620
Structure WFA (Beam-driven)	15 (1-15)	50	>10	>25	18-50	$\sim$450
FCC-hh	100	30 (60)	>10	>25	30-50	$\sim$560
SPPC	125 (75-125)	13 (26)	>10	>25	30-80	$\sim$400

Snowmass ITF Report (2208.06030)

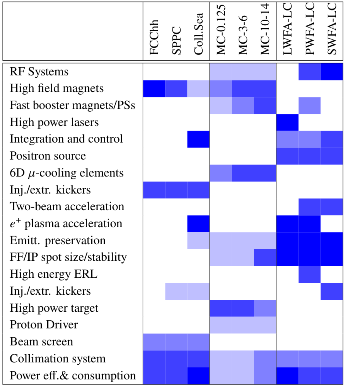

How to Collide Muons

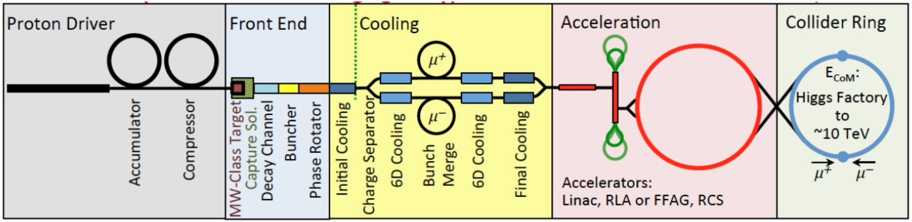

- Main muon collider challenges all **accelerator** related:
  - 1-4 MW target for proton driver: alternatives to liquid mercury needed.
  - 6D ionization cooling: must focus beam as quickly as possible, reduce transverse emittance.
  - Fast ramping magnets to inject, accelerate beam: need $1000\,\mathrm{T/s}$, plus $16\,\mathrm{T}$ DC magnet.
  - Collider ring: 12-$16\,\mathrm{T}$ large aperture dipoles, 15-$20\,\mathrm{T}$ quadrupoles; **similar to FCC-hh**.
  - Neutrino radiation flux from decaying muons.
- Baseline design from US Muon Accelerator Program, updated by IMCC (2303.08533)

Siting Options

- Options for Fermilab or CERN: fit **on current land**, reuse **existing** accelerators.

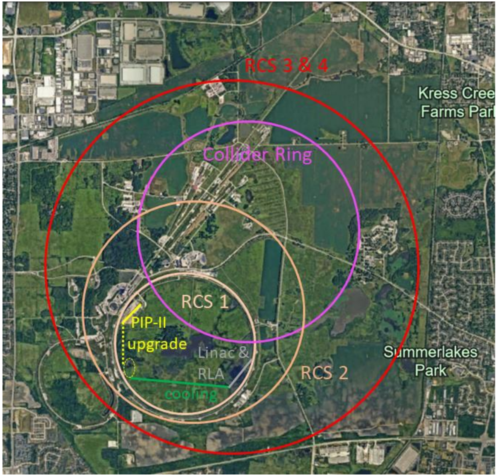

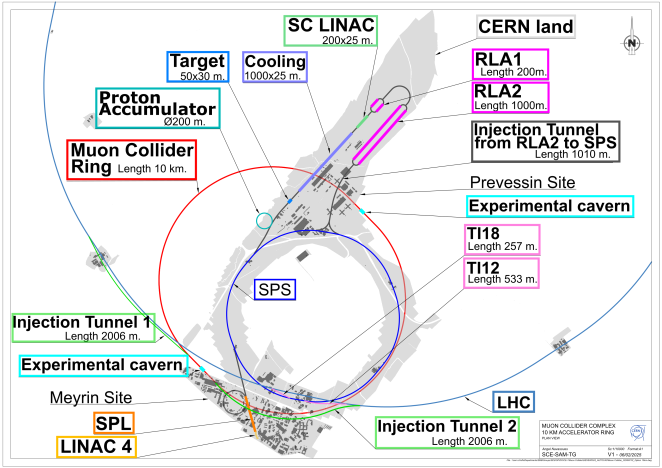

IMCC 2025 ESPPU Input

Detector Challenges

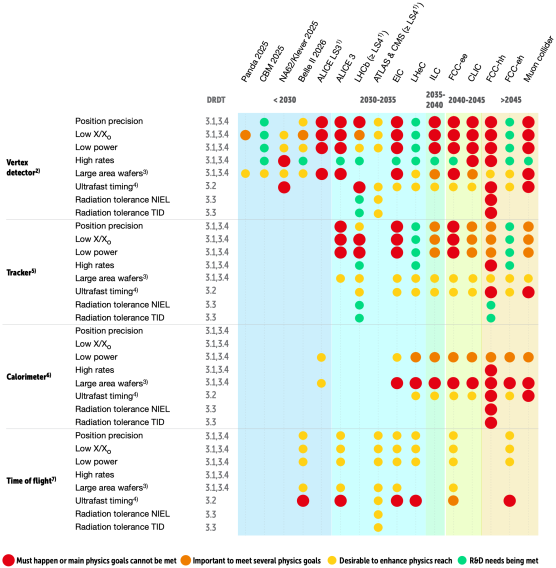

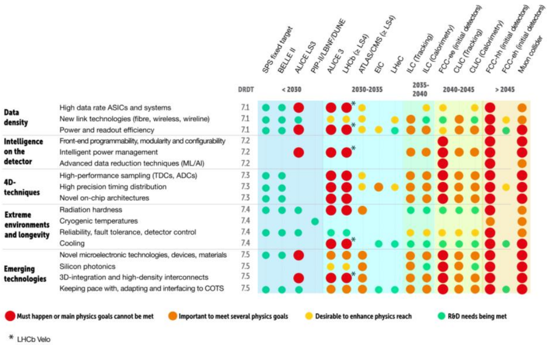

- ECFA Roadmap for future colliders:
  - Many common needs with e^+e^-, pp: work can benefit multiple projects!
- What is unique about muon colliders?
  - **Beam induced background** (BIB) from muon decay: huge challenge!

Calorimetry

	ECAL	HCAL
Cell type	Silicon - Tungsten	Iron - Scintillator
Cell Size	5.1mm × 5.1mm	30.0mm × 30.0mm
Sensor Thickness	0.5mm	3.0mm
Absorber Thickness	2.2mm	20.0mm
Number of layers	50	100

- Current proposals comparable to existing calorimeter technology:
  - ECal very similar to CMS high granularity calorimeter upgrade.
  - HCal similar to ATLAS TileCal (10x smaller sensors).
- Impact of solenoid shielding (v4):
  - Approximately $4X_0$ worth of material.
  - Reduces incoming BIB flux by a factor of **10** compared to previous design.

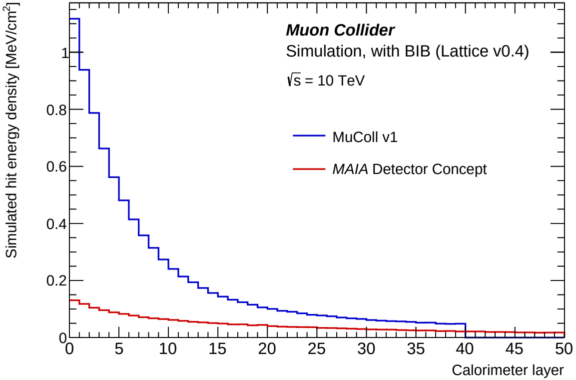

Variable Cell Thresholds

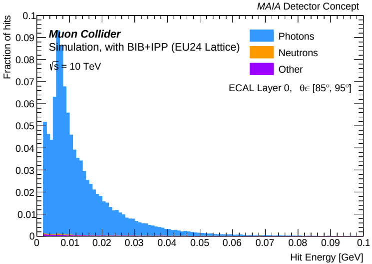

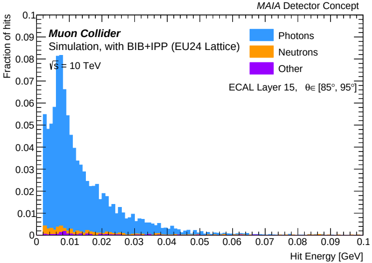

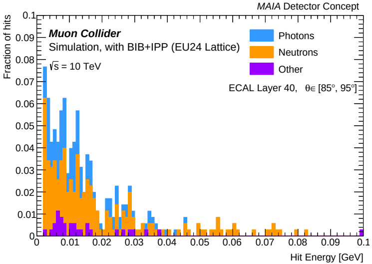

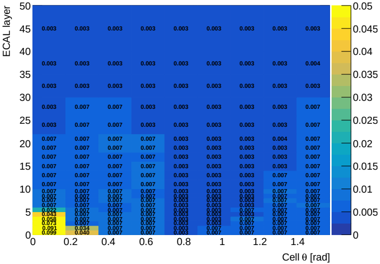

- BIB in the ECAL:
  - Lower layers dominated by **photons**.
  - **Very low energy**, soft, diffuse: hard to reconstruct.
- High cell thresholds needed for BIB:
  - Derive **cell-dependent** thresholds from BIB.
  - Higher thresholds to reject photons.
  - Strong handle at reducing fakes; at cost of worse resolution.

FPGA Tracking: (Semi) Classical

- EF Tracking process selecting **technology**, not final trigger system.
  - For FPGAs: three maximal pipelines built on top of minimal "data preparation" (F-100).
  - Classical (no-ML); semi-classical; and full ML.
- Classical (F-600): use custom conformal transform to find **5-layer pixel seeds**.
  - Extend tracks outward using estimated track parameters, simple $\Delta\phi/\Delta z$ search windows.
  - Linear fit with precomputed constants on FPGA; final Kalman Filter refit on CPU.
- Semi-classical (F-610): very similar, except **neural networks** for extrapolation/fitting.

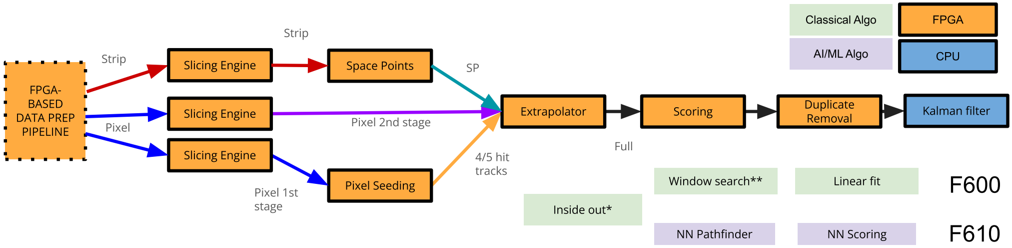

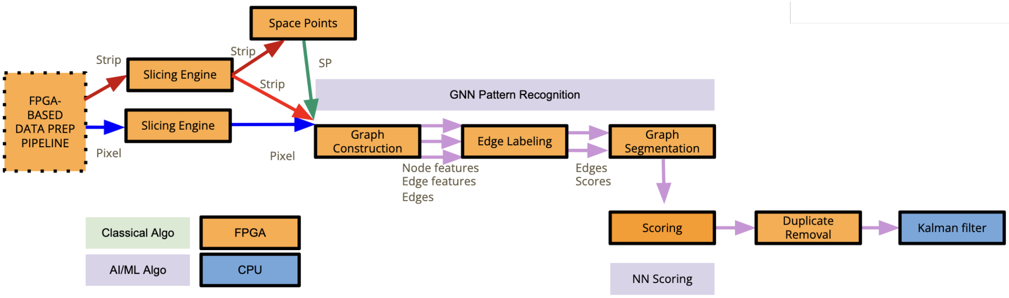

- Use machine learning for **full** track reconstruction (F-410):
  - GNN finds full-detector tracks, no need for separate extrapolation step.
  - Use same NNs as semi-classical pipeline (F-610) to score and predict track parameters.
  - Track scores used to run on-FPGA duplicate removal.