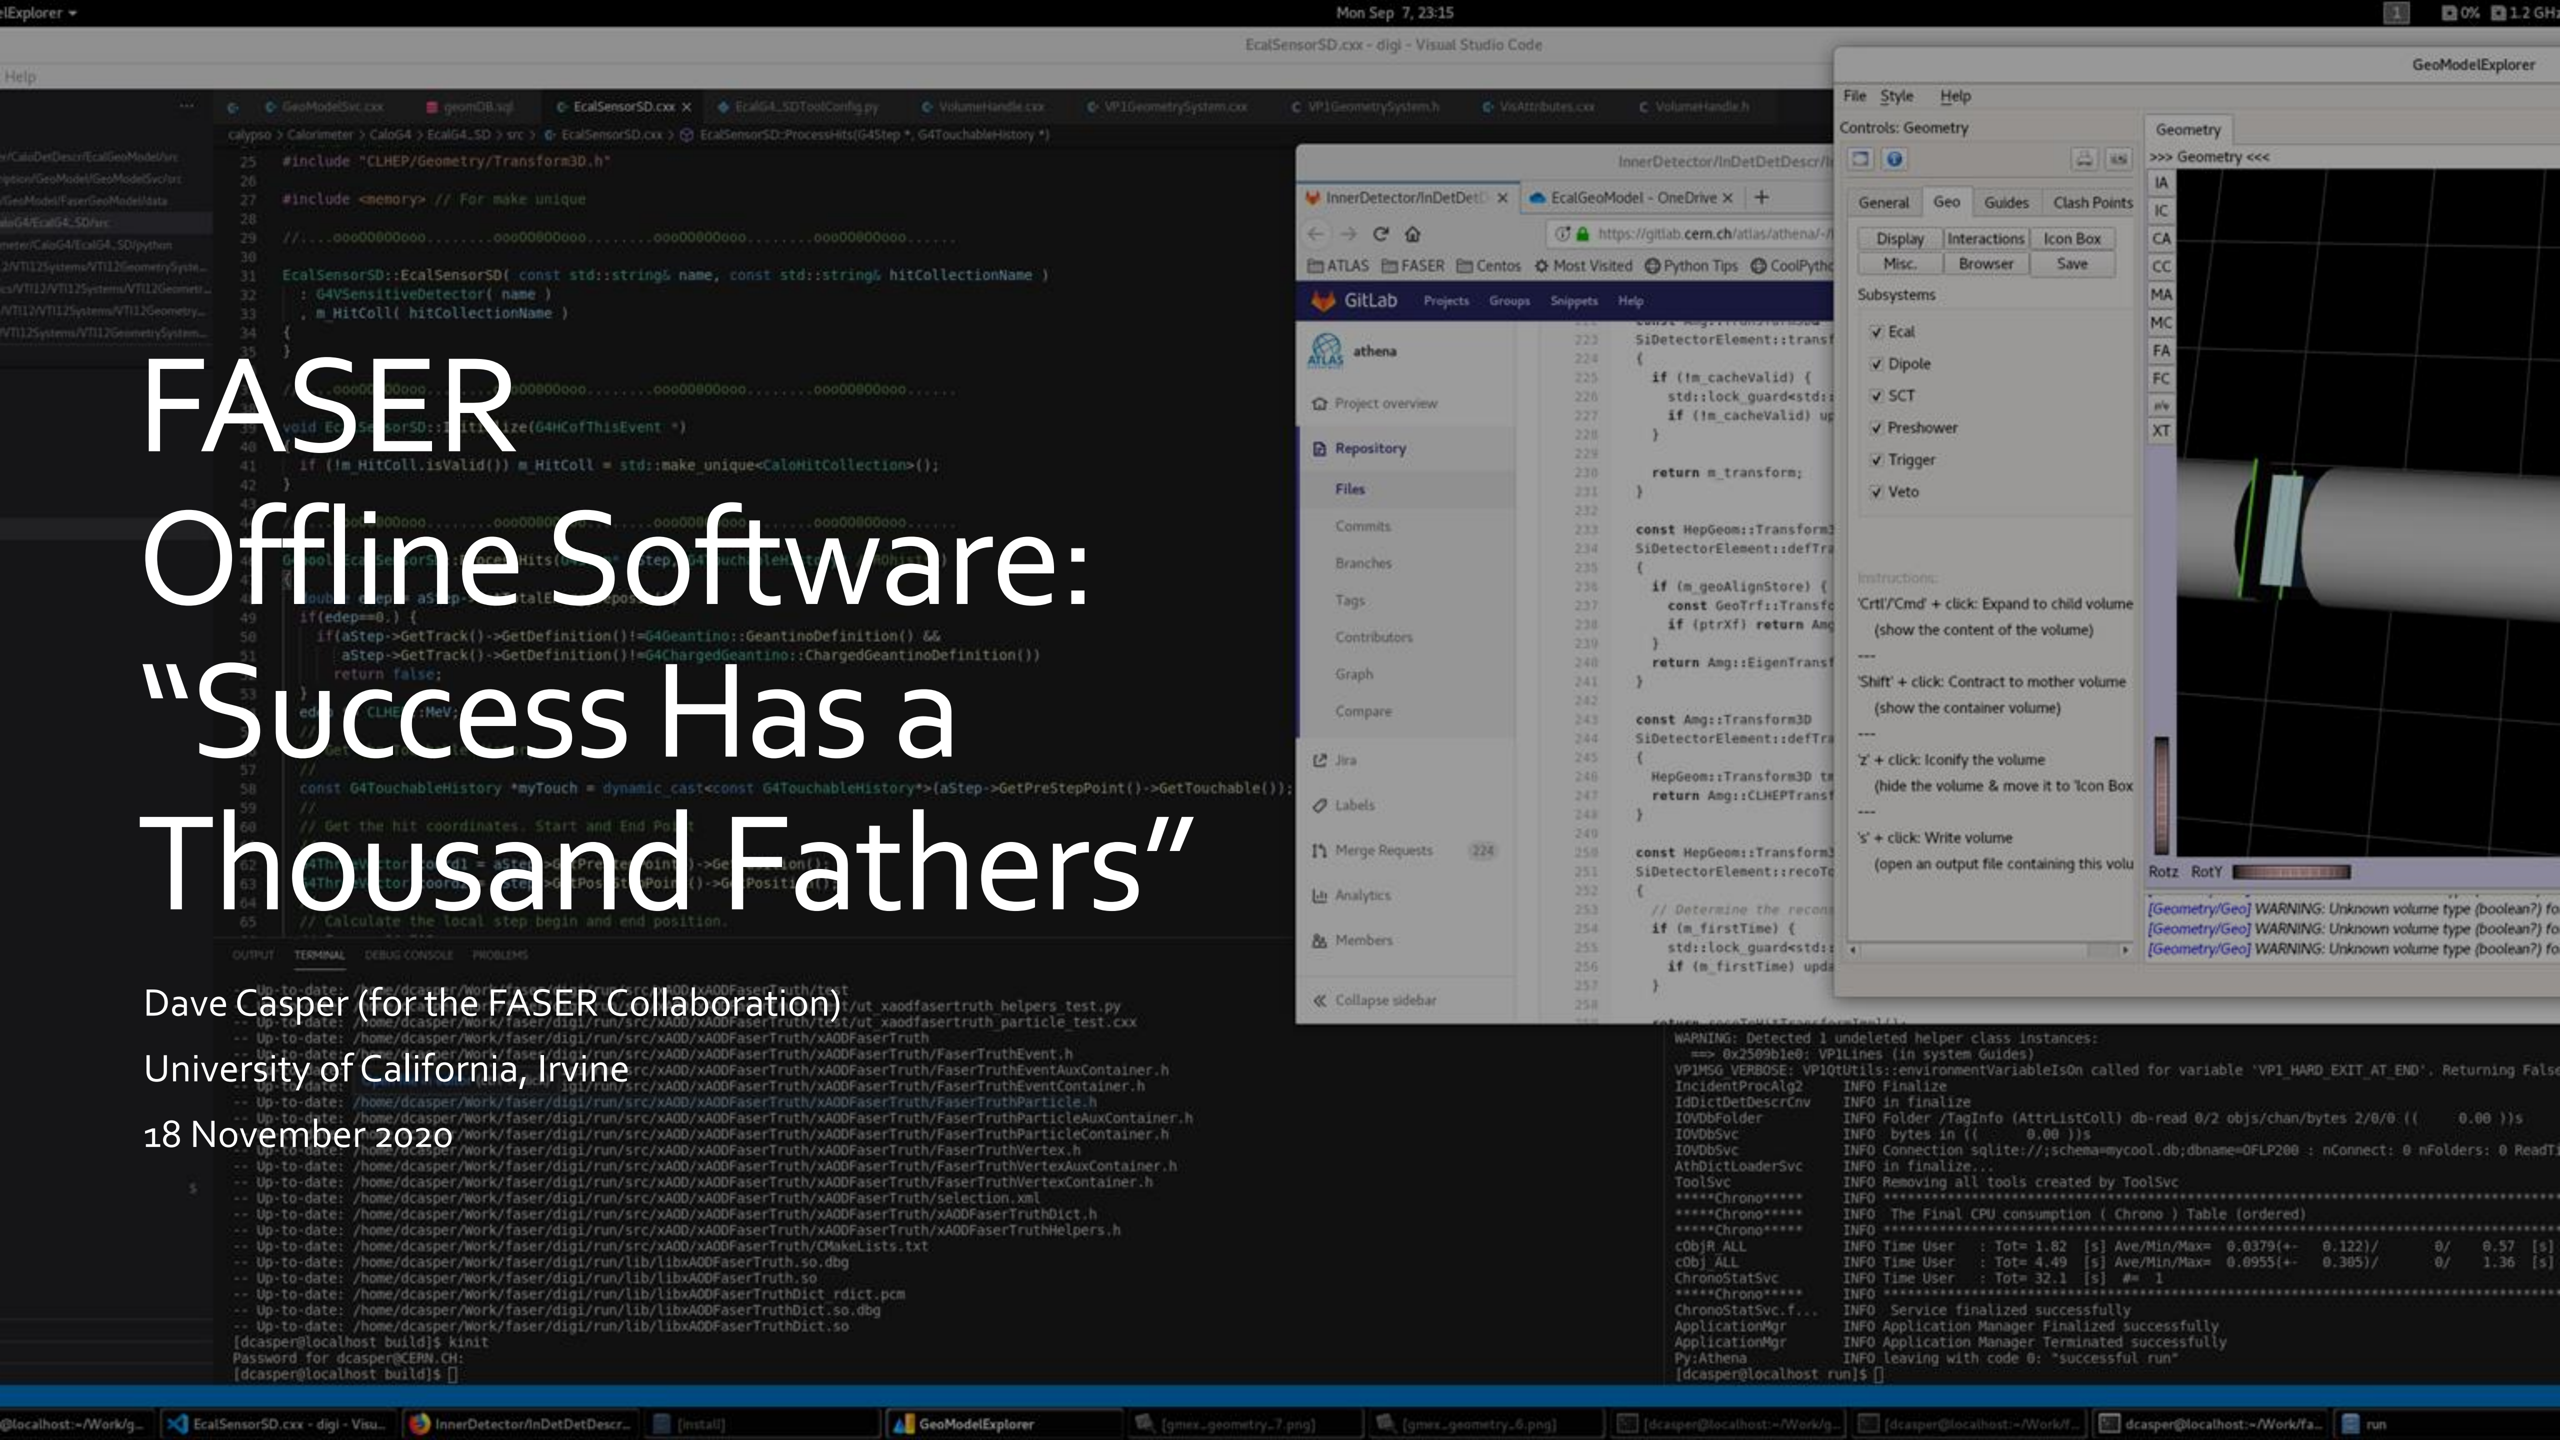

FASER Offline Software: "Success Has a Thousand Fathers"

Dave Casper (for the FASER Collaboration)
University of California, Irvine

18 November 2020

Outline

- About FASER
  - Collaboration
  - Goals and design
  - Status and plans
  - Offline software requirements
- Software Framework
- Detector Description
- Event Generation
- Track Reconstruction
- Conclusion

FASER Collaboration

FASER Collaboration Members

Henso Abreu (Technion), Yoav Afik (Technion), Claire Antel (Geneva), Akitaka Ariga (Bern), Tomoko Ariga (Kyushu/Bern), Florian Bernlochner (Bonn), Tobias Boeckh (Bonn), Jamie Boyd (CERN), Lydia Brenner (CERN), Dave Casper (UC Irvine), Franck Cadoux (Geneva), Xin Chen (Tsinghua), Andrea Coccaro (INFN), Monica D'Onofrio (Liverpool), Candan Dozen (Tsinghua), Yannick Favre (Geneva), Deion Fellers (Oregon), Jonathan Feng (UC Irvine), Didier Ferrere (Geneva), Stephen Gibson (Royal Holloway), Sergio Gonzalez-Sevilla (Geneva), Carl Gwilliam (Liverpool), Shih-Chieh Hsu (Washington), Zhen Hu (Tsinghua), Peppe Iacobucci (Geneva), Sune Jakobsen (CERN), Enrique Kajomovitz (Technion), Felix Kling (SLAC), Umut Kose (CERN), Susanne Kuehn (CERN), Helena Lefebvre (Royal Holloway), Lorne Levinson (Weizmann), Ke Li (Washington), Jinfeng Liu (Tsinghua), Chiara Magliocca (Geneva), Josh McFayden (CERN), Sam Meehan (CERN), Dimitar Mladenov (CERN), Mitsuhiro Nakamura (Nagoya), Toshiyuki Nakano (Nagoya), Marzio Nessi (CERN), Friedemann Neuhaus (Mainz), Hidetoshi Otono (Kyushu), Carlo Pandini (Geneva), Hao Pang (Tsinghua), Brian Petersen (CERN), Francesco Pietropaolo (CERN), Markus Prim (Bonn), Michaela Queitsch-Maitland (CERN), Filippo Resnati (CERN), Jakob Salfeld-Nebgen (CERN), Osamu Sato (Nagoya), Paola Scampoli (Bern), Kristof Schmieden (Mainz), Matthias Schott (Mainz), Anna Sfyrla (Geneva), Savannah Shively (UC Irvine), John Spencer (Washington), Yosuke Takubo (KEK), Ondrej Theiner (Geneva), Eric Torrence (Oregon), Serhan Tufanli (CERN), Benedikt Vormwald (CERN), Gang Zhang (Tsinghua)

- 64 members from 18 institutions and 8 countries

- Major financial support provided by:

Physics Goals and Design

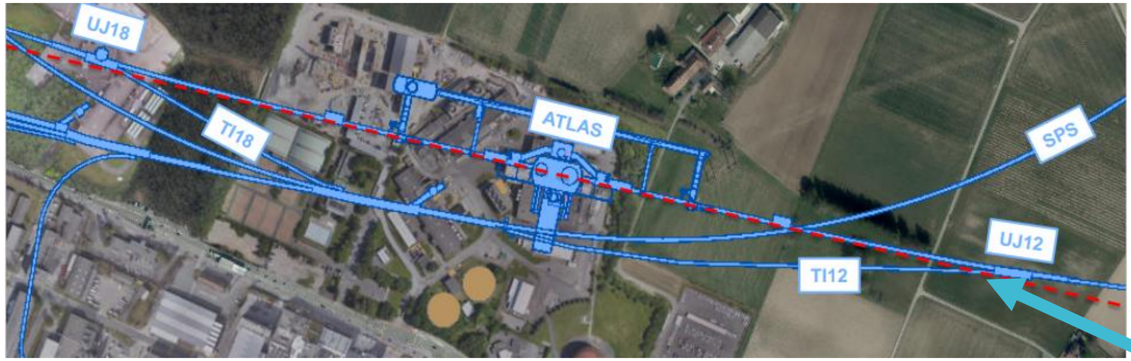

- Primary physics goals:
  - Long-lived dark photon search (e.g. $A' \to e^+e^-$)
  - TeV-scale measurements of all three neutrino flavors
- Detector design:
  - Magnetic spectrometer to measure energetic charged particles
    - 96 ATLAS Silicon Tracker modules
    - 3 permanent dipole magnets (0.57 Tesla)
  - Electromagnetic calorimeter for independent energy measurement and particle ID
    - Using 4 LHCb Ecal modules
  - Emulsion neutrino detector
  - Plastic scintillators for veto, trigger and preshower

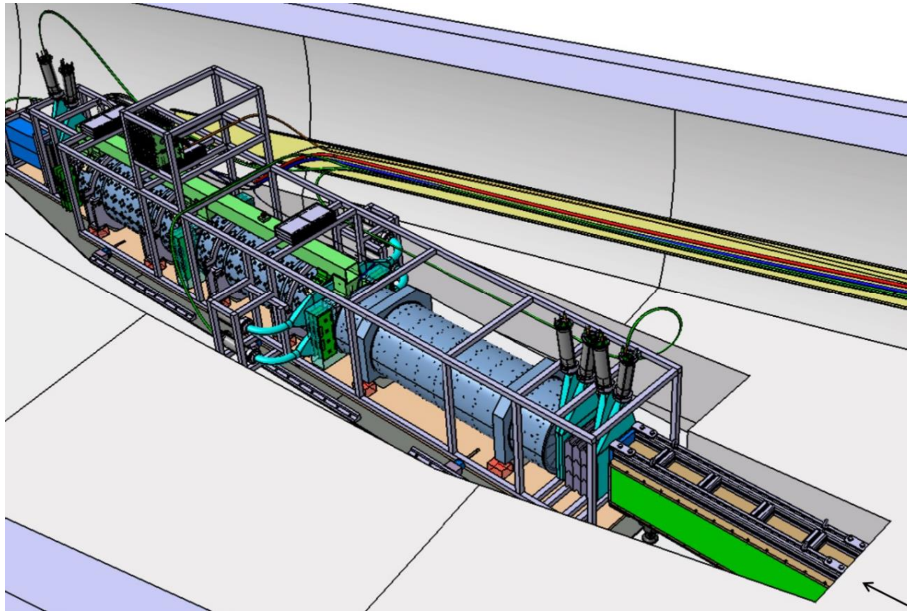

Expected raw trigger rate < 1 kHz (single muons)
Average raw event size: ~25 kB (uncompressed)
Raw data rate: $\sim 1\ \mathrm{TB/fb^{-1}}$ (uncompressed)
(does not include emulsion detector scans)

Detector Installation Status and Plans

- FASER conceived: August 2017
- Experimental collaboration forms: ~January 2018
- Letter of intent: July 2018
- Technical proposal: September 2018
- Tentative approval & initial funding: December 2018
- Final approval by CERN Research Board: March 2019
- Cosmic ray tests on surface: July – November 2020
- Cabling, installation of magnets and other hardware underground began: ~3 weeks ago!
  - Outstanding support from CERN technical staff under difficult conditions
- Second phase (installation of tracker stations, scintillators and calorimeter) to begin early January
- FASER's *hardware* should be ready when collisions begin!

Offline Software: Scaling Violations

- Our detector is physically (and logically) small
  - Tracker has only about 2% of ATLAS SCT channels
  - Calorimeter is 0.07% of LHCb Ecal channels
  - Designed and constructed rapidly and inexpensively
    - Thanks to hardware donated by ATLAS and LHCb
- Collaboration is also small, and most have other commitments
  - Fewer than 10 developers actively working (most part-time) on FASER offline software!
- Unfortunately, the offline software system does not scale to the size, cost or construction time of the experiment, or the size of the collaboration
  - Our offline software must do most of the things a much larger experiment's would
  - We have fewer subdetectors but the same requirements for them

Offline Software Requirements

- Core framework
- Detector description
- Alignment/calibration/conditions
- Event data model and persistency
- Data preparation
- Data quality validation and monitoring
- Track reconstruction
- Event generation
- Detector simulation
- Electronics simulation (“Digitization”)
- Event display

In Search of a Framework

- In addition to the daunting scope of our software requirements, we also have much less time than a typical large experiment (with a life-cycle measured in decades) would
  - Obvious conclusion: we can't do it by ourselves
- In late 2018, as FASER was nearing final approval, ATLAS released their offline software framework ("Athena") under the Apache 2.0 open-source license
  - Perfect timing!
- Athena is derived from the LHCb Gaudi framework
- Both rely heavily on the LHC Computing Grid (LCG) software stack maintained at CERN
  - ROOT
  - Geant4
  - many other packages

From Athena to Calypso

- “Pros”
  - Most actual or potential FASER developers are ATLAS members
  - Athena will be maintained and improved for the lifetime of FASER
  - We could not hope to create something as functional and robust within FASER
  - Good for students to learn from a well-designed, state-of-the-art framework
- “Cons”
  - Complexity: contains many features we don’t want or need
  - Athena was not designed with use outside ATLAS in mind
    - Many parts are experiment-agnostic, but others are not
    - Because Gaudi became a joint LHCb/ATLAS effort relatively early in its development, it is better in this regard
  - Parallel development will require care
- FASER’s extension of Athena and Gaudi is named “Calypso” (a daughter of the Greek titan Atlas).

Calypso: The Art of Subtraction

- Michelangelo is said to have remarked that in creating his masterpiece, "all I did was chip away everything that didn't look like David."
  - In some sense, Calypso is created by chipping away everything in Athena that doesn't look like FASER.

- Requirements we get (almost) for free from Athena:
  - Core software framework
  - Event data model and persistency
  - SCT data preparation (clustering, spacepoints)
  - Detector description and event display (more below)
  - Alignment/conditions infrastructure (via CORAL/COOL)
  - Geant4 detector simulation interface
  - SCT electronics simulation
- Features we do not currently expect to need or support
  - Multi-threading
  - GRID jobs
  - Distributed databases (Oracle)

Managing Parallel Development

- Athena is under very active development within ATLAS
  - Multi-threading
  - New job configuration
  - Python3 migration
  - Detector description for Run-4
  - ACTS
- Our initial development system linked directly against ATLAS Athena releases or nightlies on cvmfs
  - Extremely convenient, but also fragile
- New paradigm is to maintain our own (forked) branch of Athena, and build only the parts we need
  - Currently using 354/2088 Athena packages
- Install "our" Athena + Calypso on cvmfs
  - No more reference to ATLAS binaries

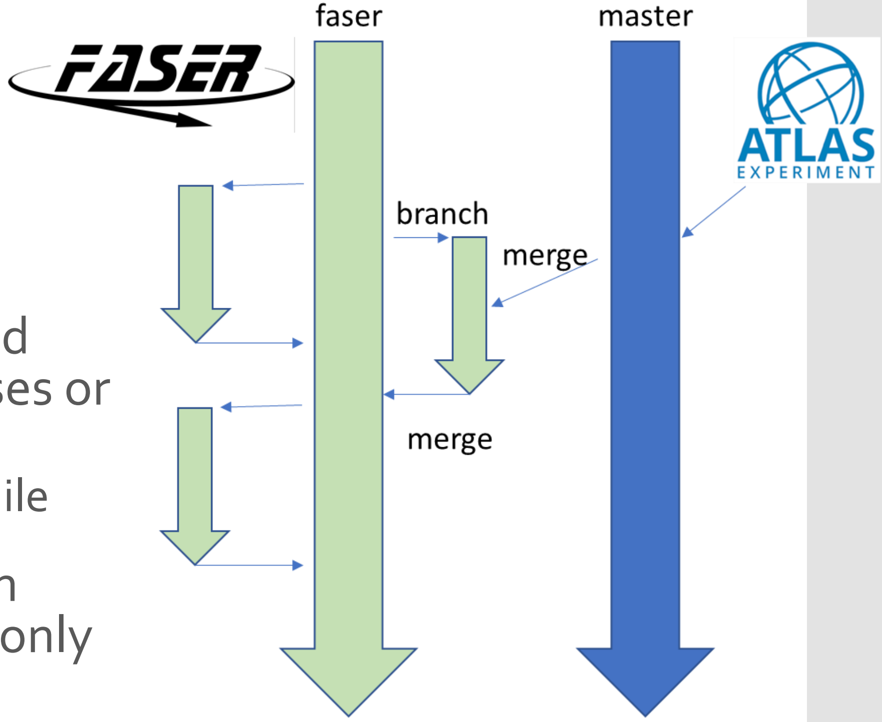

Detector Description: GeoModel

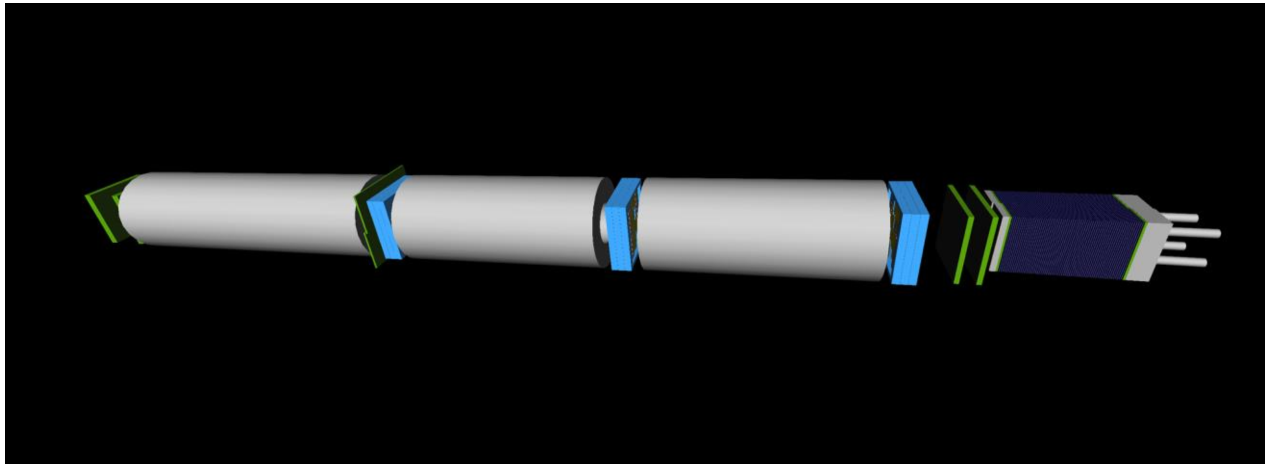

- The ATLAS detector description framework, GeoModel, has been spun-off into a standalone package
- Old paradigm:
  - With versioned primary numbers from database, use “detector factories” to construct detector volume tree “on the fly” at run-time
- New paradigm (under development for Run-4)
  - Use “plugins” to build subdetector volumes ahead of time; store “as-build” volume tree in SQLite database; load directly into memory at run-time. A given SQLite database represents a single version of the geometry.
  - FASER will likely serve as one “proof-of-concept” test
- Another FASER-friendly improvement to GeoModel is the ability to load a volume tree from GDML
  - This allows us to use the native LHCb Ecal detector description file without modification.

Simulated 1 TeV Electron in FASER Event Display

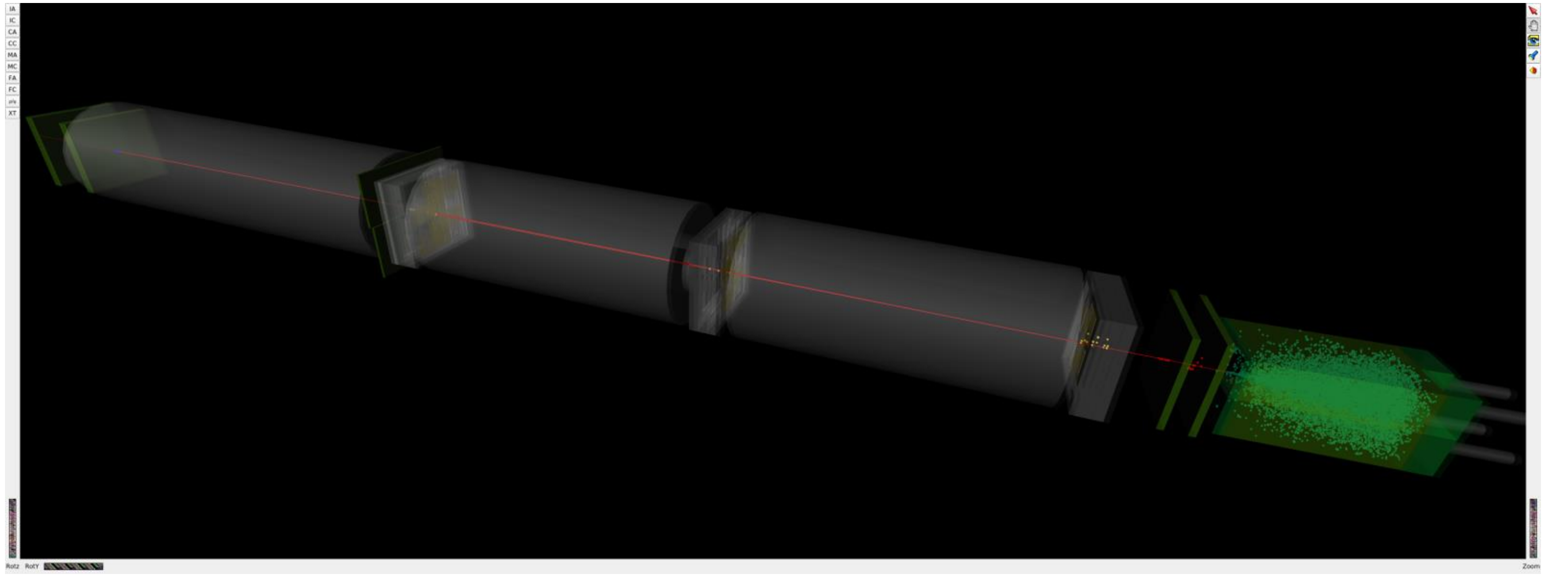

Tracker modules use native ATLAS SCT detector description (GeoModel).
Calorimeter uses native LHCb Ecal detector description (GDML).

Event Generation Overview

- Event generation in FASER is arguably more complex than in ATLAS
  - The most commonly-used pp generators are tuned to describe particles emerging with high p_T. The very forward region relevant for FASER, where traditional collider experiments are blind, is less well understood.
  - Any pp reaction product detected by FASER (muon, neutrino or exotic) is the result of decay, scattering and/or bending into a tiny region of solid angle nearly collinear with the beams.
  - The physics of neutrino scattering will need to be simulated for the first time in a collider experiment, and in a previously unexplored energy regime.

Propagation and Decays: FLUKA and BDSIM

- BDSIM is a Geant4-based beam transport simulation used by FASER to predict muon and neutrino fluxes at TI-12
  - Includes detailed geometry of tunnel and beamline elements
  - Uses primary pp generator (e.g. CRMC) events as input
  - BDSIM developers are members of FASER
- FLUKA simulations of comparable detail performed by the CERN STI group for comparison

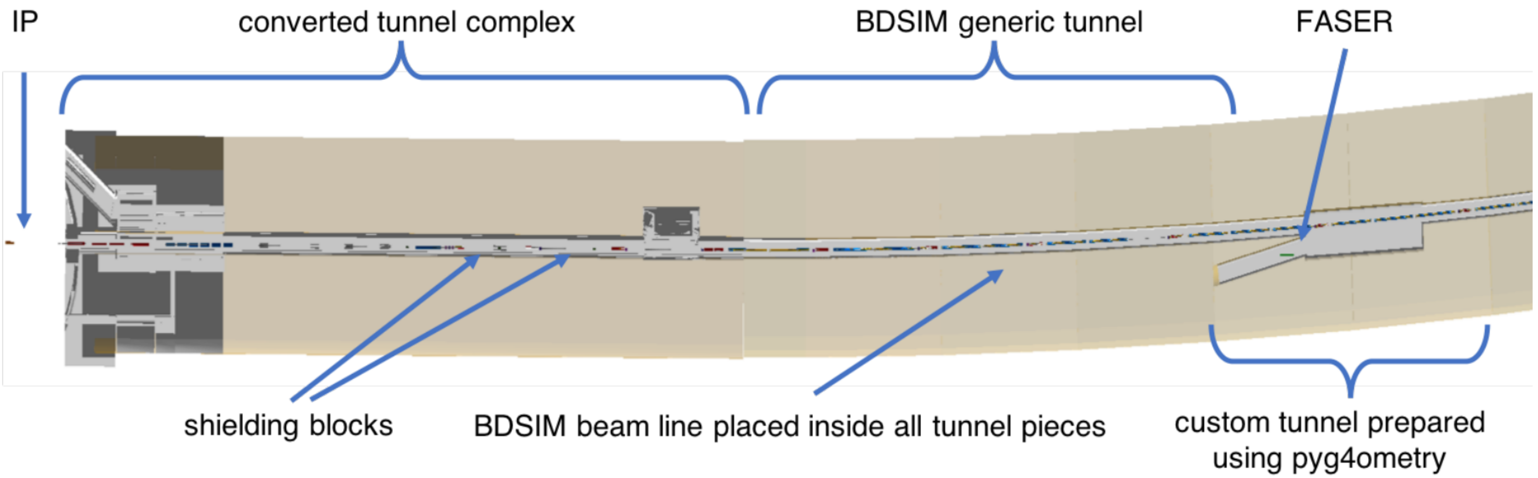

Neutrino Scattering: GENIE

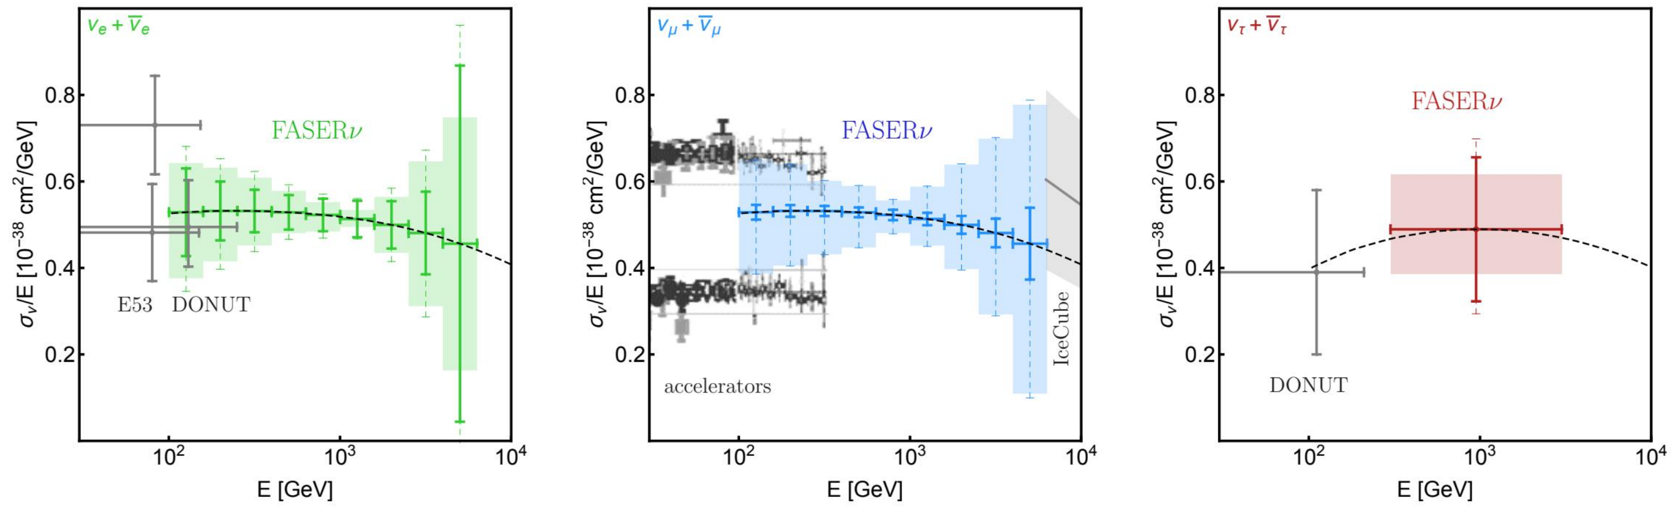

- GENIE is a general-purpose neutrino event generator
  - Simulates exclusive (quasi-elastic, resonant, coherent) and inclusive (deep-inelastic) neutrino reactions
- Requires integration with detector description to properly distribute interactions according to target mass and composition

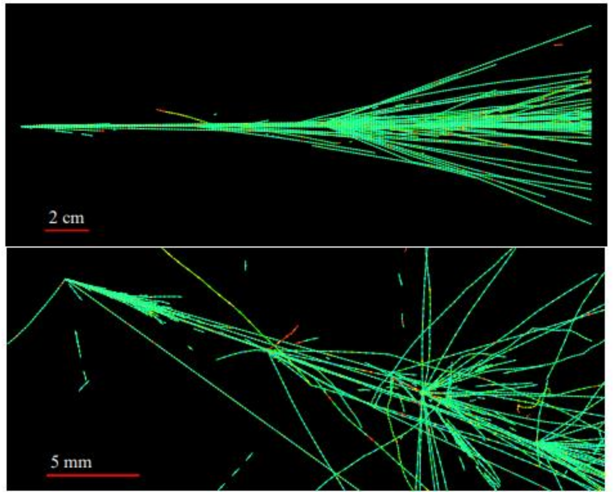

1 TeV ν_μ charged-current interaction generated by GENIE

Track Reconstruction with ACTS

- ACTS (“A Common Tracking Software”) is a modern tracking toolkit based on lessons learned from ATLAS track reconstruction in Runs 1 and 2.
  - FASER has always planned to adopt ACTS, and has been in contact with the developers (and attending each others’ meetings) since Summer 2018
  - Work on ACTS in Calypso is proceeding in parallel with, but independent of, ACTS work in Athena
- ATLAS plans to use ACTS for track reconstruction in the future, but has mature legacy code available until then.
  - FASER has no such “insurance policy,” and adapting the ATLAS legacy code to FASER is likely not feasible.
  - Establishing baseline track reconstruction functionality with ACTS is therefore urgent for us

ACTS in Calypso: Status

- We are close to accomplishing end-to-end track finding and reconstruction with ACTS.
- ACTS elements successfully integrated with Calypso:
  - Tracking geometry
  - Magnetic field map
  - Propagator
  - Combinatorial Kalman Filter

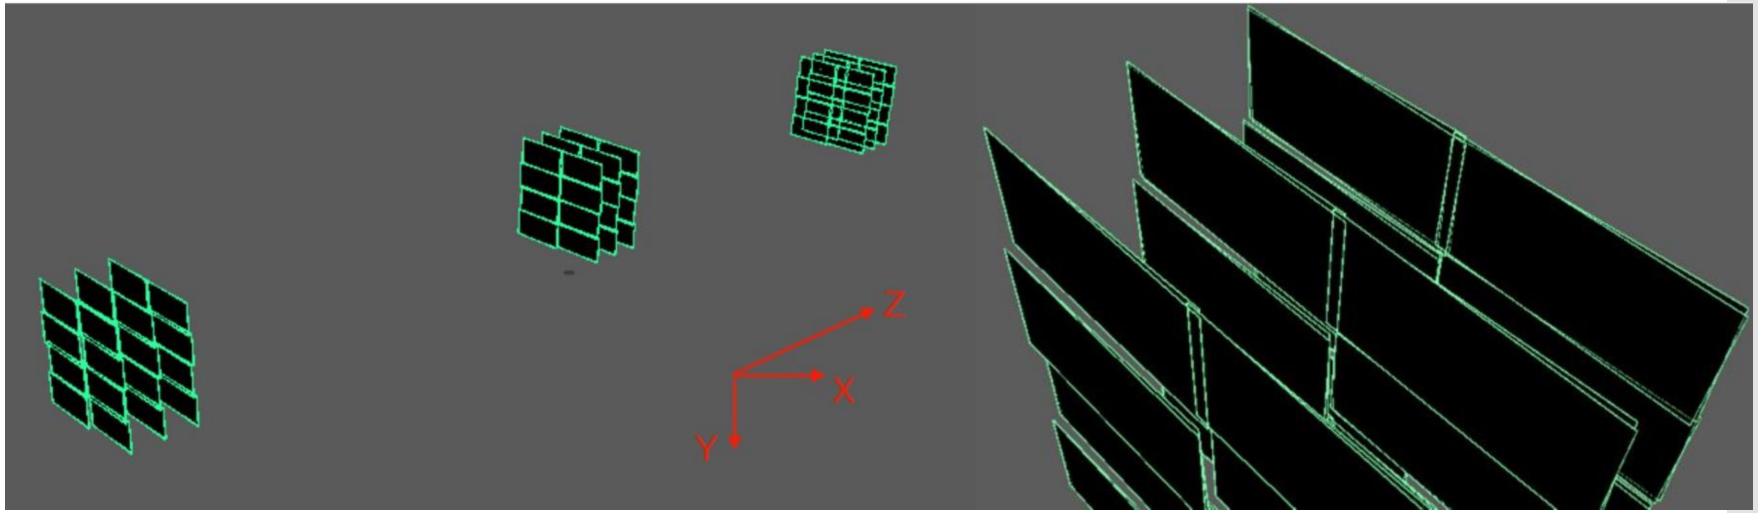

ACTS in Calypso: Future Plans

- ACTS work required to complete end-to-end tracking:
  - FASER-specific track seed finder
  - ACTS-based tracking event data model
- Less urgent but necessary work to follow:
  - ACTS material-mapping
  - ACTS alignment
  - Visualization of ACTS data objects
  - Vertex fitting

Summary

- FASER went from an idea on a theorist's whiteboard to a fully-approved and funded LHC experiment in just 18 months
  - First installation period just completed successfully, on schedule
  - Installation will be completed after the new year
- Despite FASER's small physical (and human) size, its software requirements are comparable to those of much larger experiments
- "Standing on the shoulders of giants"
  - Aggressive re-use and repurposing of software originally written for ATLAS and LHCb, and for the wider HEP community, is allowing us to successfully meet the experiment's requirements with very modest resources and available time

Collider neutrino candidate (2018)

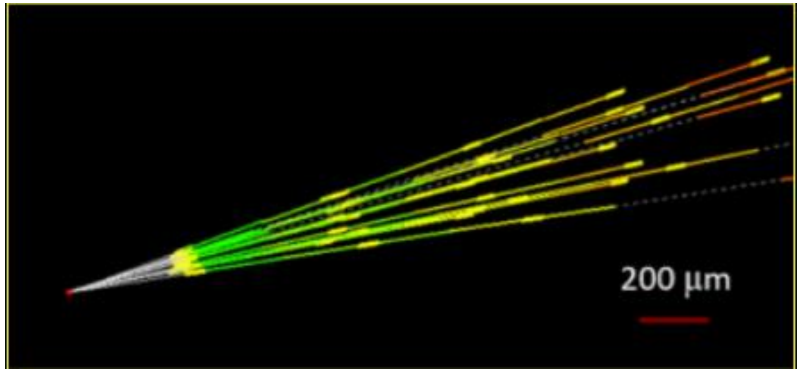

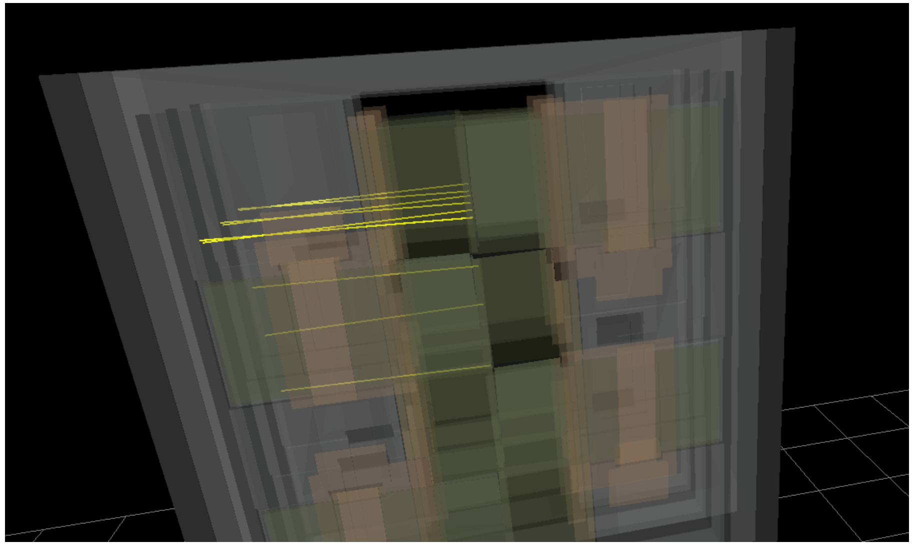

Cosmic-ray muon test-stand data (August 2020)