

Event Processing Frameworks - Animator: Liz

Framework area now in HSF

- Group has resurveyed what has been done in the past two years
- Brand new in terms of coordination work

CMS

- External worker (to hide latency of sending work to GPU and back)
- Task-based

CCE

- What is the proper programming model (CUDA? SYCL, Kokkos)
- ATLAS (Charles L and his work on scheduling)
- Deliverable: present what are the best candidates for the portable parallelization

CERN

Fermilab: near-term focus on portability

RAL

- Atthena

Generic:

- How to integrate external resources and accelerators into our RECO (or, more generically, frameworks)
 - CMS: External Worker
 - ATLAS: Just starting to understand how to put this into Gaudi

LHCb

- Allen - Full framework running on a CPU. Totally new (replaces Gaudi, etc., for trigger).

Gaudi

- Time to schedule work is too long for the trigger
- ATLAS: work on making the scheduling in Gaudi faster
- Spack - config manager for building packages - Does a graph solution to order the build - which is just like ordering algorithm execution.

R&D Ecosystem:

- CMSSW(ART) effort at FNAL
- LBNL effort on Athena
- LBNL effort on Ray (UCB/Rise Lab)
- Potential CCE effort on portability libraries (FNAL, LBNL)
- CERN effort (Attila, unclear if SFT is actively working on this)
- Gaudi - traditionally CERN SFT - not clear if there is the required effort here
- RAL/Edinburgh effort on Athena

Strengths:

- Clear historical groups and R&D tasks
- Reasonable communication and the new HSF group helps

Weaknesses:

- R&D Developers feel overly constrained by weight of existing frameworks and software

Opportunities:

- CCE effort can be kernel, with Framework developers, can drive effort on portability libraries
- We have 4 people on C++ standards committee, our work could drive things more broadly than HEP - get vendors to adopt a single standard for using accelerators, like the GPU, etc.

Threats:

- Inability to evolve the frameworks for new hardware, tasks
- Early choice forced by schedule constrains the field before sufficient R&D is done.

Milestones:

- Decide on programming model that people will use to write an algorithm on an accelerator.
- Frameworks integrate scheduling with this programming model.
 - Generic means of interfacing to accelerators

Simulation - Animator: Jim

R&D Ecosystem:

- Geant:
 - GeantV: (ending)
 - Make sure these developments are preserved/incorporated into future versionGe
 - CERN is resetting its priorities

- Geant4 lives forever?
- Just use it to train GANs?
- Continuous support needed in order to validate “new” Geant
 - (whatever that is?)
 - Development may take 10 years?
 - CERN is currently on the line for maintaining GEANT4
 - However, CERN is complaining there is no one in the USA supporting G4.
- Hope for continued speedup in G4
 - ATLAS: geometry modeling (Joe Boudraux, +)
 - (DD4HEP adoption failed - talk at HSF? Check back)
- Geant X - the future
 - How do we get there
 - If X is when it's ready, what's X?

Fast Simulation

- Not enough R&D to parametrization to make some sort of common project
- GANs:
 - Not performant (yet)
- Projects:
 - CaloGAN - Univ. of Geneva: includes some key US folks
 - Mostly University programs working
 - Some of new 50/50 postdocs
 - No dedicated funding for this yet
 - CERN OpenLab has funded an FTE or so
 - NESAP project is joint CMS and ATLAS effort Ben and JR supervise one PD

Strengths:

- G4 is a solid foundation and shared by experiments (and the outside world).

Weaknesses:

- GV project has ended, future isn't obvious. (And I can't find my shades anyway.)

Opportunities:

- Come up with a plan for GEANTX - and work with funding agencies to see what is possible. The plan is obviously a function of X, where X is the year when it is ready
- Collaboration with ECP (Tom's framework discussion from Liz).
- GEANT workshop at ORNL
- Medium size project focused on simulation using GAN.
- Automated tuning of parametrized FastSim

Threats:

- GEANT4 - A Gathering Storm now that GeantV has been declared dead by CERN

- GANs may never work: do we have a strategy for FastSim?

Milestones/Action Items:

- Figure out how much Geant X will cost (aiming for Run 4?)

Visualization

ROOT

- For CMS it is happening in the context of the ROOT project
- Rewrite of EVE
-

Effort to abstract geometry

ATLAS & CMS both support Viz with OPS programs

Reco /Trigger: Animator: David Lange

Which US R&D projects are working in this area?

Are there international efforts?

What is the FTE effort levels for each US project?

What are the goals of each R&D project?

Is the R&D connected to the Experiments and/or the Operations programs?

Projects have their own goals and milestones

Tracking R&D projects:

- mkFit
- Exa.TrkX
- ACTS
- Vertexing w/ ML
- (CMS) CERN Patatrack
- FNAL Tracking/Reco SciDAC (collaborating with mkFit, neutrinos)
- (ATLAS) RAL group doing GPU tracking for HLT (ATLAS Mgmt organizing dedicated forum for this and other projects)
- IRIS-HEP: UIUC-Princeton collaboration on GNN tracking on TrackML dataset

Are these coherent? Should they be?

Workshop looking at level of commonality? We think encouraging this would be useful.

- Leverage “Connecting the Dots”?

CCE: using the seeding codes from these projects can be used as demonstrator/test bed for code portability testing/validation

Accelerator demonstrators R&D projects

FPGA demonstrator for HLT

- Lower barriers for technology test bed by providing examples
- Formalize documentation and code bases - build on FastML workshop but need the facilities generally available, is this a role for SLATE?

Calo/Jets R&D projects

Graph networks

- FASTJET: CCE Target demonstrator for portability
- LDRD: exa.Trk - using graph networks for clustering 1.2 FTE (two shared postdocs)
 - Some targeting of CMS HGCal
- IRIS-HEP: benchmark datasets for tool evaluation
-

Strengths:

Weaknesses:

Reco a big driver for HL-LHC

Can Atlas/CMS capitalize on real-time analysis developments in LHCb - no real effort here in US projects

ML vs traditional approaches

Multiple approaches

Computational vs physics evolutions - rewrite vs improve/evolve scientific outcome

Current evolution in algorithms towards GPUs vs CCE R&D timescales on programming models

Evolve single experiment (or dominated by single experiment efforts)

Incoherence

Mkfit and ACTS have not estimated effort to move to accelerators - Paolo: “by run4 for sure”

Opportunities:

- Parts of the CWP program are uncovered, eg. QA/QC tools, real-time analysis

- DOE has initiated a Basic Research Needs (BRN) study (report due in Feb 2020, townhall, community meetings in Dec 2019 at CPAD and elsewhere). The aim is to identify priority research directions and key challenges covering detector R&D and motivate seeking funding for them. TDAQ is one subgroup covering accelerators at HLT, exploration of heterogeneous architectures, fast machine learning for triggering etc.

Threats:

Milestones:

Opportunities for Machine Learning

Text from the AI for Science Townhall here in DC:

- **Usable** tools for **large-scale distributed training and optimization** of ML models to enable scaling up the complexity of models to orders of magnitude above the current state-of-the-art
- Training methodologies that are able to detect **rare features** in high-dimensional spaces while being robust against systematic effects
- Tools to quantify the impact of **systematic effects** of the accuracy and stability of complex ML models
- High-quality generative models satisfying physical constraints and symmetries
- Fast methods for solving high-dimensional statistical inverse problems

\$70M is dedicated to this work in ASCR (Paolo: “Not that much money”)

Can HEP get some of its interests in as a use case (e.g. rare feature extraction)?

Potential Action items:

- Catalogue or group of white papers describing the HL-LHC problems
 - Prepared to collaborate
- Form HEP Computational Data group to respond/interact with the ECP groups
 - “Math groups” are hungry for collaborators

Weakness:

- Strong connections to the foundational groups in the ML community

Opportunity:

- Summarize the killer-apps for ML for HEP that can be shown around to garner interest
 - Before Dec 15th

Killer Apps:

- Generator -> Reco/analysis in one step (a Gigaflop, not a Teraflop)
- FastML