GaudiHive in ATLAS

Charles Leggett

What We've Done

- Simulation of ATLAS reconstruction w/ CPUCrunchers
  - shows lots of potential parallelism from just the data dependency graph
  - only cloning the most expensive Algorithms still gives us excellent performance
- Small sub-detector testbeds
  - Calorimeter and Inner Detector
  - real reco code, real data
  - identified lots of issues with framework and user code
    - much better understanding of what we need to address
  - shown performance benefits and memory savings
- Extended Hive for a hybrid MP/MT framework
  - when concurrency is limited, get better performance with more MP workers
  - balance between performance and memory usage

Where We Are Now

- build current Atlas dev release against Hive every night
  - builds about as cleanly as our regular dev builds
  - merging Hive specific changes into code trunk
    - not many #ifdef ATHENAHIVES or branch tags left!
  - running/developing several Hive enabled packages

- Produced a Future Frameworks Requirements Document, identifying requirements for what ATLAS needs in Run3 and beyond
  - https://cernbox.cern.ch/public.php?service=files&t=69ccb3f0a2cbfca389c7a469c319bc69

- preparing for Hive workshop/sprint at CERN in early July
  - identify certain areas/weaknesses that need to be worked on
  - developers from LHCb and CMS will participate so that we get other points of views

What We're Working On

- G4Hive
  - Atlas Geant4 Simulation in multi-threaded and hybrid environment
    - 1 data reader alg, 1 G4 Sim alg (5 SD enabled), 1 data output alg
      - I/O is serial and mutually exclusive. only small fraction of CPU time
      - no mag field
    - leverages G4 v10 with MT enabled
    - scales well up to total number of cores on platform
  - processed 9M events with 450 ranks w/ 24 threads each on edison.nersc.gov
    - 2% failure rate due to threading issues – may be solved
    - output data file sizes scale inversely with #threads.... may have solved this too
- Calorimeter
  - jet finding with 16 parallel JetFinders (after clusters made)
  - scales well with #threads
  - some possible memory issues

Short Term Goals

- This year

- Start working on making important Services thread safe
  - Magnetic Field
- Implement a short term solution for the Data Dependency registration problem
- Transition away from public Tools
- Address IncidentSvc
- Start working on making I/O more concurrency friendly, or at least less of a hater

- get something real to run in production
  - G4Hive

Medium Term Goals

- Next year

- Full merging of GaudiHive with Gaudi trunk
- Transition ATLAS code to VarHandles
  - will fix Data Dependency problem
- Implement HLT EventViews
- Test/convert many more Algorithms with Hive
- Concurrent I/O

- big production jobs on Cori

- Educate developers in multi-threaded programming techniques, or at least how not to shoot themselves (and us) in the foot. or actually head.