DUNE Workflow Development and Analysis

Elizabeth Paton, University of Alaska Anchorage
Kenneth Herner, Fermilab

Introduction

The Deep Underground Neutrino Experiment (DUNE) is planned to be the largest liquid argon neutrino detector. DUNE will take place between Fermilab and The Long Baseline Neutrino Facility (LBNF) in South Dakota, which is currently under construction. To prepare for DUNE, the DUNE prototype, ProtoDUNE is currently used for testing components of DUNE.

DUNE will be very important in the role of understanding neutrinos by allowing us to their interactions with other particles. Observing neutrinos will be a key step to understanding why there is matter and other big questions in physics.

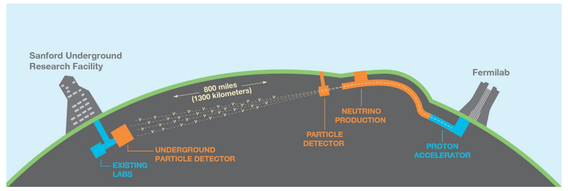

Purpose

Neutrino interactions are rare, so a lot of data must be stored to allow for meaningful interpretation. Recent changes in the ProtoDUNE workflow have led to an increased memory footprint, which results in a reduced overall throughput.

To lower this increase in memory consumption and increasing throughput, steps in the workflow are further split into stages. By making these steps small, less work is done in each stage. These modifications reduce how many subdetectors are in memory at the same time, which will save computing resources for DUNE in the future.

Methods and Materials

Using DUNE Lar TPC, we ran the single phase ProtoDUNE simulation. After editing the config files to test a change in the workflow, we run the jobs again to compare where the job count vs Resident Set Size peaks.

The workflow steps are the following:

- Neutrino interaction generation (Gen)
- Geant 4 simulation of interaction with detector (G4)
- Electronics Simulation (Detsim)
- Event Reconstruction (Reco)

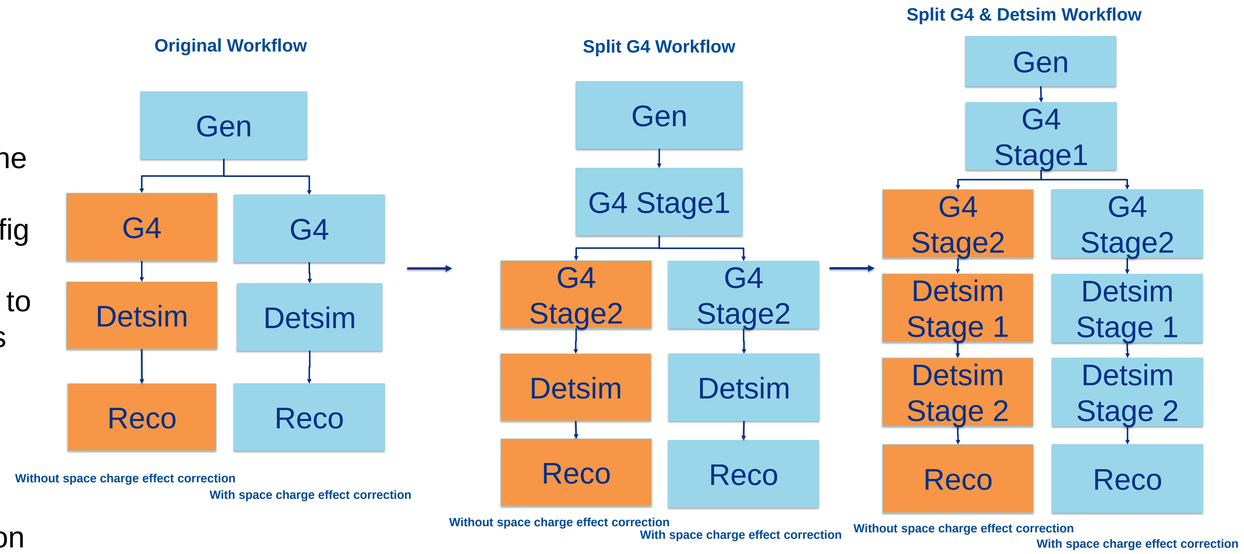

The charts above shows how the ProtoDUNE Single Phase simulation chain was reiterated into more steps to test each iteration's memory consumption.

In each iteration, one step in the chain is split into two stages, which is further split into two sides. There are two sides since one notes the space charge of the particles, while the other side does not.

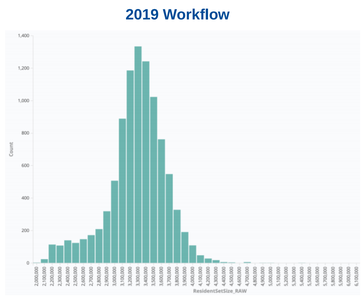

Job Count vs Resident Set Size (KB)

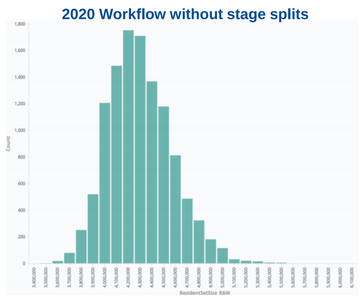

Job Count vs Resident Set Size (KB)

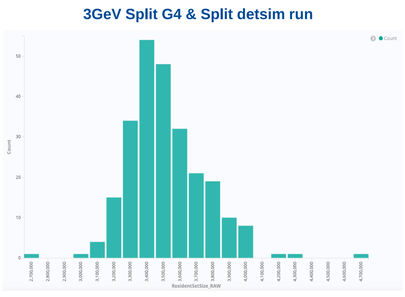

Job Count vs Resident Set Size (KB)

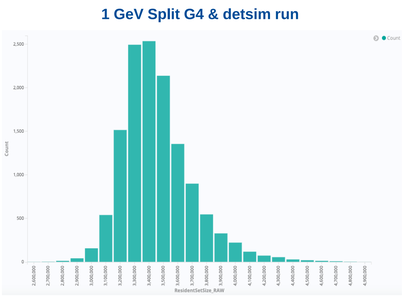

Job Count vs Resident Set Size (KB)

The graphs to the upper right illustrate the differences in memory consumption between several workflow chains.

Notice that the graphs don't peak in the same place.

Conclusion

We found that the methods of reducing memory consumption were successful but did not improve as much as previously expected.

We also found that memory consumption has a dependence on the momentum of the beam of particles supplied to the detector since a simulation run with less momentum consumes less memory.

Work supported by the Fermi National Accelerator Laboratory, managed and operated by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for U.S. Government purposes.