G4RD4: Instruction and Data Cache Optimizations

This task aims to identify current bottlenecks via performance profiling with various tools such as Intel VTune, perf, valgrind, etc, and then use the information to propose changes that may improve Geant4 performance.

Status: Some performance measurements have been made and presented. The main bottleneck identified comes from high numbers of cache misses (instruction, data, and TLB) in a few parts of the code. Development is currently ongoing to try to improve performance based on these measurements.

Page: *https://geant4.web.cern.ch/node/1879*

Repository: main Geant4 repository, may move to fork in EP-SFT GitLab group to use GitLab issues

Developers: GA, JA