Contribution ID: **9** Type: **Hands-On Tutorial**

Concurrent Programming in C++

Tuesday, August 27, 2019 1:15 PM (4h 45m)

In this course we will introduce how to program for concurrency in C++, taking advantage of modern CPUs' ability to run multi-threaded programs on different CPU cores. We will briefly review the native C++ concurrency features for
asynchronous execution, thread spawning and locking as well as a few other features useful for concurrent programming. The tutorial will then show you how to use Intel's Threaded Building Block (TBB) library as a much higher level
abstraction onto concurrency that allows concurrent applications to be developed for scientific applications much more quickly. We will examine TBB's basic templates for parallel programming, controlling loops and reductions in 1 and 2 dimensions. Then we will see how TBB's graph execution facilities that allow more sophisticated parallel workflows in a DAG to be run. We will look briefly at the TBB task manager, that allows arbitrary workloads to be executed in parallel when injected into the system by a higher level component. Our goal will be to build up a simple application that exploits multi-level concurrency, executing different tasks, some of which exploit with inner loop optimisations.

Students should be familiar with C++ and the standard template library. Some familiarity with makefiles and/or CMake would be useful.

Presenter: STEWART, Graeme A (CERN)

Session Classification: Tutorials