

CEA-EDF-INRIA - Computer Science Summer School

Course 1: How to build a visualization application for very large data.

Speaker: Hank Childs (Lawrence Berkley National Lab)

Abstract :

Supercomputing simulations generate massive data sets and the software to visualize these data sets is complex. In this course, we will discuss the algorithmic and design issues that go into such software.

The course will be designed around two questions: "how is parallel visualization software designed today?" and "how will parallel visualization software be designed in the future?" For the first question, we will discuss how to construct a library of interoperable modules (data flow networks), how to make such a design work in a parallel setting, parallel rendering, and the details of non-embarrassingly parallel algorithms, such as particle advection and volume rendering. To better understand the second question, we will discuss problems introduced by petascale- and exascale-class machines, namely I/O limitations, massive concurrency, data movement, power, and architectural limitations. We will then describe solutions: in situ processing, multi-resolution processing, and data subsetting, as well as practical issues like what it takes to work on accelerators like a GPU. Finally, we will discuss the process of deploying production visualization software to a community.