Run-time Visualization with the Cactus framework and VisIt's Simulation Library

Eloisa Bentivegna (Center for Computation and Technology at LSU, U.S.A.)
Oleg Korobkin (Center for Computation and Technology at LSU, U.S.A.)
Thomas Radke (Albert Einstein Institute, Germany)
Erik Schnetter (Center for Computation and Technology at LSU, U.S.A.)

This poster will illustrate the motivation, techniques and progress in run-time visualization from large-scale simulations instrumented with the Cactus framework, using the capabilities provided by VisIt and its simulation library `libsim`. In particular, we will show how Cactus variables (actual data like meshes and arrays or metadata representing the simulation's status) can be exposed to VisIt on-the-fly, circumventing the need for disk I/O and, especially, data file transit over the network.

The tool takes advantage of VisIt's modular architecture, and in particular the logical separation into viewer and engine components and the remote communication infrastructure between them, for processing the simulation data locally and concurrently, and only communicating the rendered image to the remote viewer. The user can then manipulate the image using VisIt's standard toolbox, while the raw data never leaves the compute node's memory, dramatically downsizing the latency between data production and state-of-the-art visualization that large simulation code users have to face today.

The tool also allows the user to control the data production by halting, resuming or advancing it in individual steps of adjustable granularity. Generating and interacting with data thus becomes one single, highly interactive process that has the potential to spotlight simulation faults faster than traditional methods.

The poster will illustrate the visualization and communication model for this tool, the network security solutions necessary for running it on large supercomputers, and several snapshots of a typical session.

This functionality has been developed within the Cactus Tools for Application-Level Performance And Correctness Analysis (Alpaca) project, a NSF-funded scheme for endowing Cactus with built-in, high-level HPC debugging and profiling capabilities.