G-CSC Kolloquium

Time - An Often Neglected Dimension of Parallel Performance Analysis

Felix Wolf, German Research School for Simulation Sciences, RWTH Aachen

Abstract

Advanced modeling techniques are pushing the complexity of high-performance computing applications to new limits; a development that finds its obvious expression in the dramatically rising numbers of cores used for individual jobs. At the same time, the deployment of adaptive algorithms and sophisticated load-balancing schemes make the execution behavior of simulation programs increasingly dynamic. Understanding parallel performance therefore requires capturing and analyzing performance data along the dimensions of both space and time, each presenting its unique scalability challenges. In this talk, we will discuss how recent research in the Scalasca project, a performance-analysis tool specifically designed for large-scale applications, will improve the understanding of time-dependent performance behavior. We will consider three approaches suitable for different granularity levels together with a general methodology of how to combine them most effectively. The first approach establishes an overview of how the performance behavior of a parallel program evolves over time, keeping both memory requirements and runtime dilation as low as possible. The two remaining approaches examine cause-effect chains of wait state formation within smaller sections of the overall execution, allowing the scalable attribution of waiting times to their origins and the simplified characterization of dynamic load imbalance. The three new techniques will be illustrated using examples of real-world applications.