

Generating HPC Job Profiles and Expectations with Time-Series Data

In the field of High Performance Computing it can be challenging to quantify the expected performance of job, especially if that job utilizes multiple nodes executing in parallel. Current techniques for job analysis allow users to track their code via function calls and to determine what resource their code is bound by. However, these techniques often demand a significant amount of time and effort on the part of the developer to fully understand the limitations of their code and to investigate issues. We are demonstrating a system for automatically collecting a job's relevant time series data, and using that data to produce complete job profiles for every job on a machine. In our system, we are using LDMS data such as CPU utilization, memory, and bandwidth. These job profiles not only encompass raw time series data, but also preprocessed statistical series, and can be readily visualized by simply navigating to a web page. In addition to visualizing a profile on its own, profiles from past successful runs can be used to generate a statistical expectation for a job's behavior, which can be visualized as a cloud path. When a job deviates significantly from this path, this indicates anomalous behavior. It's also worth noting that we have developed a machine learning technique to classify jobs into workload types when that information isn't reliably supplied by the user. Providing a readily available live visualization of a job's performance allows users and administrators to take immediate action when a job is idling or in a failed state, and as a result save valuable compute hours. The utility of job profiles extends beyond providing a window into an individual job's performance. For example, job profiles taken together across a whole system can help us to visualize and automatically detect system-wide anomalies. Also, job expectations can be applied to system benchmark jobs to monitor for system changes. Finally, detection of anomalous job behavior can be used to label log data and improve log analytics. In other words, encapsulating this type of data into job profiles enables a variety of new avenues for data analysis and investigation. To summarize, we have built a system that automatically generates statistical expectations for a job's performance based on metric time series data, and makes that information readily available through web visualization for all jobs on a system.