



Contribution ID: 19

Type: **Parallel**

GOoDA: The Generic Optimization Data Analyzer

Monday, May 21, 2012 1:55 PM (25 minutes)

Modern superscalar, out-of-order microprocessors dominate large scale server computing. Monitoring their activity, during program execution, has become complicated due to the complexity of the microarchitectures and their IO interactions. Recent processors have thousands of performance monitoring events. These are required to actually provide coverage for all of the complex interactions and performance issues that can occur. Knowing which data to collect and how to interpret the results has become an unreasonable burden for code developers whose tasks are already hard enough. It becomes the task of the analysis tool developer to bridge this gap.

To address this issue, a generic decomposition of how a microprocessor is using the consumed cycles allows code developers to quickly understand which of the myriad of microarchitectural complexities they are battling, without requiring a detailed knowledge of the microarchitecture. When this approach is intrinsically integrated into a performance data analysis tool, it enables software developers to take advantage of the microarchitectural methodology that has only been available to experts.

The Generic Optimization Data Analyzer (GOoDA) project integrates this expertise into a profiling tool in order to lower the required expertise of the user and, being designed from the ground up with large-scale object-oriented applications in mind, it will be particularly useful for large HENP codebases

Primary authors: LEVINTHAL, David (Google); CALAFIURA, Paolo (Lawrence Berkeley National Lab. (US)); VITILLO, Roberto Agostino (LBNL); ERANIAN, Stephane (Google)

Presenter: VITILLO, Roberto Agostino (LBNL)

Session Classification: Software Engineering, Data Stores and Databases

Track Classification: Software Engineering, Data Stores and Databases (track 5)