



Contribution ID: 545

Type: Oral

Real-time HEP analysis with funcX, a high-performance platform for function as a service

Tuesday 5 November 2019 11:00 (15 minutes)

The traditional HEP analysis model uses successive processing steps to reduce the initial dataset to a size that permits real-time analysis. This iterative approach requires significant CPU time and storage of large intermediate datasets and may take weeks or months to complete. Low-latency, query-based analysis strategies are being developed to enable real-time analysis of primary datasets by replacing conventional nested loops over objects with native operations on hierarchically nested, columnar data. Such queries are well-suited to distributed processing using a strategy called function as a service (FaaS).

In this presentation we introduce funcX—a high-performance FaaS platform that enables intuitive, flexible, efficient, and scalable remote function execution on existing infrastructure including clouds, clusters, and supercomputers. A funcX function explicitly defines a function body and dependencies required to execute the function. FuncX allows users, interacting via a REST API, to register and then execute such functions without regard for the physical resource location or scheduler architecture on which the function is executed—an approach we refer to as “serverless supercomputing.” We show how funcX can be used to parallelize a real-world HEP analysis operating on columnar data to aggregate histograms of analysis products of interest in real time. Subtasks representing partial histograms are dispatched as funcX requests with expected runtimes of less than a second. Finally, we demonstrate efficient execution of such analyses on heterogeneous resources, including leadership-class computing facilities.

Consider for promotion

No

Authors: BABUJI, Yadu (University of Chicago); Dr BLAISZIK, Ben (University of Chicago); Dr CHARD, Kyle (University of Chicago); Dr CHARD, Ryan (Argonne National Laboratory); Prof. FOSTER, Ian (University of Chicago); Dr LI, Zhuozhao (University of Chicago); SKLUZACEK, Tyler (University of Chicago); Dr TRISOVIC, Ana (University of Chicago); Dr WOODARD, Anna Elizabeth (University of Chicago); KATZ, Daniel S. (University of Illinois)

Presenter: Dr WOODARD, Anna Elizabeth (University of Chicago)

Session Classification: Track 7 –Facilities, Clouds and Containers

Track Classification: Track 7 –Facilities, Clouds and Containers