

# Identifying memory allocation patterns in HEP software

*Tuesday 11 October 2016 11:00 (15 minutes)*

HEP applications perform an excessive amount of allocations/deallocations within short time intervals which results in memory churn, poor locality and performance degradation. These issues are already known for a decade, but due to the complexity of software frameworks and the large amount of allocations (which are in the order of billions for a single job), up until recently no efficient mechanism has been available to correlate these issues with source code lines. However, with the advent of the Big Data era, many tools and platforms are available nowadays in order to do memory profiling at large scale. Therefore, a prototype program has been developed to track and identify each single de-/allocation. The CERN IT Hadoop cluster is used to compute memory key metrics, like locality, variation, lifetime and density of allocations. The prototype further provides a web based visualization backend that allows the user to explore the results generated on the Hadoop cluster. Plotting these metrics for each single allocation over time gives new insight into application's memory handling. For instance, it shows which algorithms cause which kind of memory allocation patterns, which function flow causes how many shortlived objects, what are the most commonly allocated sizes etc. The paper will give an insight into the prototype and will show profiling examples for LHC reconstruction, digitization and simulation jobs.

## Primary Keyword (Mandatory)

Software development process and tools

## Secondary Keyword (Optional)

Analysis tools and techniques

## Tertiary Keyword (Optional)

**Authors:** RAUSCHMAYR, Nathalie (CERN); KAMA, Sami (Southern Methodist University (US))

**Presenter:** KAMA, Sami (Southern Methodist University (US))

**Session Classification:** Track 5: Software Development

**Track Classification:** Track 5: Software Development