

Towards a cross-platform performance portability math kernel library in SYCL

Wednesday, 19 May 2021 18:32 (13 minutes)

The increasing number of high-performance computing centers around the globe is providing physicists and other researchers access to heterogeneous systems – comprising multiple central processing units and graphics processing units per node – with various platforms. However, it is more often than not the case that domain scientists have limited resources such that writing multiple implementations of their codes to target the different platforms is unfeasible. To help address this, a number of portability layers are being developed that aim to allow programmers to achieve performant, portable codes; for example, Intel(R) oneAPI, which is based on the SYCL programming model. Nevertheless, portable application programming interfaces often lack some features and tools that are manifest in a platform-specific API. High-energy physicists in particular rely heavily on large sets of random numbers in nearly their entire workflow, from event generation to analysis. In this paper, we detail the implementation of a cuRAND backend into Intel's oneMKL, permitting random number generation within oneAPI applications on NVIDIA hardware using libraries optimised for these devices. By utilizing existing optimisations, we demonstrate the ability to achieve nearly native performance in cross-platform applications.

Primary authors: PASCUZZI, Vincent (Lawrence Berkeley National Lab. (US)); GOLI, Mehdi (Codeplay Software Ltd.)

Presenter: PASCUZZI, Vincent (Lawrence Berkeley National Lab. (US))

Session Classification: Accelerators

Track Classification: Offline Computing