

Title:

Delivering the OpenCL performance promise: Creating and Optimizing OpenCL Applications with the Intel OpenCL SDK

Details:

OpenCL is an emerging standard heterogeneous programming model which is of growing adoption in the industry. Many researchers are seeking to familiarize or advance themselves with the technology. Performance/optimizations and availability of tools are key factors in the adoption. This tutorial offers one-stop opportunity to learn everything about OpenCL offering from Intel.

After a general overview of OpenCL concepts, we move to the Intel specifics for the traditional client domain (applications for the regular desktop/laptop/ultrabook processors). We discuss performance tips and tools available for CPU and Processor Graphics (which is on-die GPU from Intel) respectively. We also introduce tips on heterogeneous programming, i.e. running OpenCL on CPU and Processor Graphics simultaneously and with load-balancing. We also present details on the Intel OpenCL implementation for MIC (Many-Integrated-Core) architecture. Specifically, we explain important performance implications of the current OpenCL implementation for Intel Xeon Phi. Focus is on the threading scalability and vectorization techniques that are critical to achieve good performance of OpenCL on MIC. We will present tips on Xeon Phi clusters and discuss caveats for multi-coprocessor scenario. Finally, the section on tools will guide people starting from the basics understanding of OpenCL towards non-trivial, tools-assisted optimizations.

About the author:

Maxim Shevtsov is a Software Architect in the OpenCL performance team at Intel located in Russia. He got his Masters degree in Computer Science in 2003, and prior to joining Intel in 2005 was doing various academia studies in computer graphics.