

Thrust: A Parallel Template Library for CUDA

Jared Hoberock

NVIDIA

jhoberock@nvidia.com

Thrust is a parallel template library for developing CUDA applications. Modeled after the C++ Standard Template Library (STL), Thrust brings a familiar abstraction layer to the realm of GPU computing. Thrust provides host and device variants of the STL vector container to simplify memory management and facilitate data transfers. These containers are complemented with a large collection of generic data-parallel algorithms and a suite of useful iterator adaptors. Together, these features form a flexible high-level interface for GPU programming that greatly enhances developer productivity.

In this talk we'll walk through Thrust's features and demonstrate how Thrust is used to develop larger-scale libraries and applications. We'll also discuss how Thrust enables "kernel fusion" and circumvents the "array of structs" vs. "structure of arrays" problem. Lastly, we'll present evidence that Thrust implementations are fast, while remaining concise and readable.