

GPU Acceleration of a Monte Carlo Neutral Beam Injection Code using Thrust

Sean A. Dettrick

Tri Alpha Energy, Inc

We examine the use of the Thrust library [1] to accelerate on the GPU a Monte Carlo code used for Neutral Beam Injection in FRCs. The emphasis is on using high-level code to avoid the need for very detailed C-level CUDA code. Each algorithm of the original Fortran code is re-expressed in terms of the generic algorithms available in Thrust, which are similar to those available in the C++ STL and Boost libraries. A particle sort is used and zip iterators are used to encourage memory coalescence. The acceleration of each algorithmic component is analyzed separately, including parallel random number generation, particle push, gather, and reduce-in-bucket. These are compared to the acceleration of the fully assembled code. The algorithm components covered here are also sufficient for some complete PIC models.

[1] Jared Hoberock and Nathan Bell, “Thrust: A Parallel Template Library”, 2010, <http://www.meganeutrons.com>.