

Recent Trends in HPC Architecture and Random Number Generation in the Exascale Era

Tuesday, 25 August 2020 09:50 (40 minutes)

The number of nodes/cores of the HPC systems has been increasing rapidly in recent years; some currently available systems are already equipped with more than one million cores. As the degree of parallelism in the system reaches this level, the performance of some of the traditional numerical algorithms need to be re-examined, due to a widening imbalance between the computational capability vs available memory size per core and /or inter-processor communication bandwidth. In this regard, the Monte Carlo Methods (MCMs) possess the following advantageous characteristics:

- (1) Very high degree of parallelism exists inherently, where very few information is exchanged across nodes/cores,
 - (2) Required memory size per node/core can be smaller than some deterministic mesh-based algorithms
 - (3) Due to statistical nature of MCM algorithms, fault resilience is expected to some extent.
- Therefore, we should give more considerations in expanding the application areas of MCMs such as the solution of the elliptic partial differential equations, for example.

In order to conduct the Monte Carlo calculations correctly and efficiently, the capability to generate high quality pseudo-random number sequence is crucial. The desirable features are:

- (1) Long period of the sequence,
- (2) High statistical qualities both with serial and across the multiple sequences,
- (3) Fast computation in generating the random number sequences,
- (4) Efficient initialization of multiple sequences for millions of cores.

We have been developing the "Multiple Recursive Generators (MRG)" based on the 8-th order full polynomials with a large modulus ($2^{31}-1$), which can generate statistically high-quality random number sequences with very long periods ($\sim 10^{74}$), and easily implementable jump-ahead scheme for effective parallelization. We have recently reformulated the MRG8 for Intel's KNL and NVIDIA's P100 GPU – named MRG8-AVX512 and MRG8-GPU, respectively.

This research has been conducted with Dr. Yusuke Nagasaka of Tokyo Institute of Technology (currently with Fujitsu Laboratories) and Dr. John Shalf of Lawrence Berkeley National Laboratory.

Presenter: Dr MIURA, Kenichi (NII)

Session Classification: Opening Ceremony & Keynote Session