
Steven Hamilton
**GPU Implementation of a Monte Carlo Linear Solver
Algorithm**

P O Box 2008
Bldg 5700
Rm H327
`hamiltonsp@ornl.gov`
Thomas Evans
Stuart Slattery
Massimiliano Lupo Pasini
Michele Benzi

The use of Monte Carlo methods for solving linear systems of equations has garnered attention in recent years due to the natural resiliency and scalability of Monte Carlo methods. With most of the largest supercomputers turning towards vectorized architectures such as GPUs or the Intel Xeon Phi, we must ask whether Monte Carlo methods are capable of efficiently utilizing these architectures. At face value, the stochastic nature of the Monte Carlo approach would appear to lead to limited vectorization and poor memory access patterns. In this talk, we discuss an implementation of a Monte Carlo linear solver algorithm on Nvidia GPUs using the CUDA programming language. We show that careful organization of data in memory, combined with appropriate use of GPU-specific capabilities (such as texture fetching), can have a significant impact on performance.