1096-H1-468 **Tim Andersen*** (tim@va.wagner.com), 2 Eaton St., Suite 500, Hampton, VA 23669, and **C Allen Butler** and **Michael Mascagni**. *Memory Efficient Lagged Fibonacci Random Number Generators for GPU Supercomputing.*

Graphics Processing Units (GPUs) bring the promise of supercomputing power for a fraction of the cost of traditional supercomputing, with speed-ups over comparable CPU hardware of one or two orders of magnitude. Rapid development of both proprietary libraries such as NVIDIA's CUDA and an open standard, OpenCL, have opened the doors to their cheap computing power. Unfortunately, random number generators (RNGs) have been slow to catch up with the rapid expansion of GPU computing. The number of types of RNGs available for GPUs is small, and the statistical quality of those provided with standard libraries frequently unknown. Because specific RNGs only have statistical quality for certain applications, new kinds of RNGs must be made available for GPU computing to bring the full power of GPUs to different kinds of research. Lagged Fibonacci Generators, in particular, have been difficult to develop for memory-lacking GPUs because of their large state space, which is unfortunate because they have excellent statistical properties for many applications. In this talk, we discuss our implementation of a memory efficient additive LFG for both CUDA and OpenCL. We also discuss reproducibility and portability. (Received September 04, 2013)