

MAR13-2012-020299

Abstract for an Invited Paper  
for the MAR13 Meeting of  
the American Physical Society

**Programming for 1.6 Million cores: Early experiences with IBM's BG/Q SMP architecture<sup>1</sup>**

JAMES GLOSLI, Lawrence Livermore National Laboratory

With the stall in clock cycle improvements a decade ago, the drive for computational performance has continued along a path of increasing core counts on a processor. The multi-core evolution has been expressed in both a symmetric multi processor (SMP) architecture and cpu/GPU architecture. Debates rage in the high performance computing (HPC) community which architecture best serves HPC. In this talk I will not attempt to resolve that debate but perhaps fuel it. I will discuss the experience of exploiting Sequoia, a 98304 node IBM Blue Gene/Q SMP at Lawrence Livermore National Laboratory. The advantages and challenges of leveraging the computational power BG/Q will be detailed through the discussion of two applications. The first application is a Molecular Dynamics code called ddcMD. This is a code developed over the last decade at LLNL and ported to BG/Q. The second application is a cardiac modeling code called Cardioid. This is a code that was recently designed and developed at LLNL to exploit the fine scale parallelism of BG/Q's SMP architecture. Through the lenses of these efforts I'll illustrate the need to rethink how we express and implement our computational approaches.

<sup>1</sup>This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.