Ignacio Laguna

Protecting Scientific Applications From Silent Data Corruption via Static Analysis and Machine Learning

Lawrence Livermore National Laboratory
P O Box 808
L-561
Livermore CA 94551-0808
ilaguna@llnl.gov
Martin Schulz
David F. Richards
Jon Calhoun
Luke Olson

A possible increase of soft error rates is one of the major concerns for the HPC community as larger and more complex HPC systems are built. In exascale systems parallel scientific applications may be subjected to higher degrees of silent data corruption (SDC) errors, which may arise from unprotected components of the machine. As a result, programmers of scientific codes would have to incorporate protection mechanisms in computation code that could be vulnerable to these errors. In this talk we present a framework to automatically protect scientific applications from SDC in their output. The framework uses fault injection and machine learning to learn code instructions that may yield SDC in applications output. Using the LLVM compiler, these vulnerable instructions are duplicated to detect errors. We present case studies of the technique in HPC kernels and mini applications, such as AMG and HPCCG, and show that the overhead of the technique is small in comparison to existing instruction duplication approaches.