

Estimating the Reliability of HPC Applications

Daniel Oliveira, Francis Birck Moreira, Paolo Rech, Philippe Navaux
Institute of Informatics, UFRGS
Porto Alegre, Brazil

Abstract—The error rate of current High Performance Computing (HPC) systems is already in the order of one per dozens of hours. Understanding the reliability behavior of HPC applications will be required for the next generation of supercomputers. Using the reliability behavior one can select efficient mitigation techniques for the application and fine-tune parameters such as checkpoint frequency. In this paper, we investigate the application of a machine learning model to predict the reliability behavior of HPC applications. We inject faults in more than 30 HPC

applications executing in the Intel Xeon Phi Knights Landing (KNL) and use profiling information to build a predictive model with Support Vector Machines (SVM). We show that the model can predict the Program Vulnerability Factor (PVF) with an average relative error of 7% for certain classes of algorithm, such as linear algebra and sorting. The average relative error for all algorithm classes is 22%. Such a fast and straightforward prediction model can be effective as a filter to select the most unreliable applications to perform an in-depth analysis.