

Abstract

Manal Helal

Performance Evaluation of Checkpoint/Restart Techniques for MPI Applications on Amazon Cloud,

Distributed applications running on a large cluster environment, such as the cloud instances will have shorter execution time. However, the application might suffer from sudden termination due to unpredicted computing node failures, thus losing the whole computation. Checkpoint/restart is a fault tolerance technique used to solve this problem. In this work we evaluated the performance of two of the most commonly used checkpoint/restart techniques (Distributed Multithreaded Checkpointing (DMTCP) and Berkeley Lab Checkpoint/Restart library (BLCR) integrated into the OpenMPI framework). We aimed to test their validity and evaluate their performance in both local and Amazon Elastic Compute Cloud (EC2) environments. The experiments were conducted on Amazon EC2 as a well-known proprietary cloud computing service provider. Results obtained were reported and compared to evaluate checkpoint and restart time values, data scalability and compute processes scalability. The findings proved that DMTCP performs better than BLCR for checkpoint and restart speed, data scalability and compute processes scalability experiments.