Towards HPC in the Cloud

Liqiang Wang
Department of Computer Science
University of Wyoming
wang@cs.uwyo.edu

1. Recent Advances

Parallel scientific applications require massive computing resources. Traditionally, such applications run on dedicated high performance computing (HPC) clusters located at research institutes or government agencies. The recent advances of Cloud Computing have made computing resources widely available as a utility. Infrastructure-as-a-service (IaaS) Cloud such as Amazon EC2 makes it possible to run HPC applications in a pay-as-you-go fashion. There are two major problems for current traditional HPC clusters.

- Many HPC cluster users constantly feel the pain of waiting in a queue for dedicated computing resources to become available.
- It is difficult to have a customized software environment because users have no administrator privilege.

Virtual clusters in Cloud solve both problems. A user can easily construct a virtual cluster at any time by renting virtual machines (VMs) from the Cloud provider. The user has full control over the virtual cluster and can configure a customized software environment for his/her application.

However, the performance of a virtual cluster in Cloud is currently still inferior to a dedicated HPC cluster, due to low network bandwidth and high virtualization overhead. Nevertheless, virtual clusters are sufficient and ideal for many loosely coupled HPC applications. For a university or enterprise department, a virtual cluster is a good alternative to a small physical cluster. For an HPC developer, the ease of use in obtaining a virtual cluster greatly improves productivity. Even for production HPC workloads, a virtual cluster can provide additional capacity when the local dedicated cluster runs out resources.

2. Research Challenges

To apply cloud techniques to traditional HPC clusters, there are still many open problems to be solved in order for cloud to be more amenable to HPC applications.

2.1 Performance

People usually believe that virtual clusters have inferior performance, which might be the major obstacle to apply cloud techniques to traditional dedicated HPC clusters. However, this is not always true. We have demonstrated that virtualization technique in cloud computing can improve the performance and reduce design effort for HPC programs [1]. We proposed an approach using adaptive resource adjustment based on virtualization techniques to automatically tune performance of MPI programs through a dynamic load balancing in cloud computing platforms. Specifically, our approach periodically adjusts CPU resource allocations based on the measured information, i.e., execution time in VMs and the capacity in hosts. The adjustment is based on VM virtualization, and happens at two levels, the host level and the cluster level. At the host level, our approach moves CPU resource allocation from under-utilized VMs to overloaded VMs based on their execution time. At the cluster level, our approach further monitors resource utilization of physical hosts, then migrates VM from busy ones to less-utilized ones. This further reduces load imbalance across the entire cluster by effectively reassigning overloaded jobs. This approach runs in a pure automatic way. It avoids manual code modification, does not require users to understand the particular MPI programs or any expertise on MPI programming, and does not need to transform MPI programs into specific programming models for automatic load balancing. We have implemented and

evaluated our approach using real-world scientific computation benchmarks and applications on Amazon EC2, which demonstrates that the performance can be improved up to 28.36% (with an average of 14.65%) when comparing with the baseline of application runtime.

2.2. Resource Management Using Virtualization Techniques

Today, there is an emerging trend in utilizing a resource management middleware to coordinate the upper level applications and the lower level resources. To ease deployment and improve portability and resource utilization, computing resources are often virtualized to support flexible allocation and management. Traditional HPC clusters use dedicated compute resource and lack of flexible resource management and multiplexing. To improve the utilization, compute resources could be shared by multiple HPC programs based on virtualization. The allocation ratio can be adjusted according to the current execution status. In recent years, as an emerging lightweight virtualization technique, container is being adopted by many systems to efficiently manage computing resource. Many big data computing platforms, such as Hadoop and Spark, have already utilized containers. Traditional HPC clusters could significantly benefit from such cutting-edge virtualization techniques.

2.3 Running HPC Programs on Public Cloud

Public cloud providers usually offer various types of VM instances at different prices, with varying compute, network, and storage capabilities. Cloud providers also offer different pricing models. For example, Amazon EC2 provides on-demand, reserved, and spot instances. It is a challenging issue for how to run HPC applications in Cloud with an efficient and cost-effective manner [2]. It is often hard for customers to determine the proper number of VMs that can meet the job deadline while minimizing the cost. On the one hand, using a smaller cluster may not finish the job in time. On the other hand, using an unnecessarily large cluster may incur a high cost without proportionally reducing the execution time, because an HPC application usually does not scale beyond a certain cluster size.

3. Conclusions

With advances of High-Performance Computing (HPC) in scientific computing, the idea of cloud as a platform for HPC draws large attention since it has high potential to provide great benefits for researchers and engineers in implementing and running large-scale computations. Users can simply rent computing resources in the form of a cluster of instances from cloud providers such as Amazon EC2 on demand with reasonable cost instead of owning and maintaining physical clusters. The major benefits of HPC in cloud include the following.

- *Better resource utilization*: compute resource can be managed in a fine-grained way and shared by multiple HPC programs.
- *Utility computing and scalability*: we can scale out HPC resources much easily by renting computing resource.
- *Shorter waiting time*: users do not need to wait for a long time in a queue for dedicated computing resources to become available. Based on virtualization, users could start all processes with a limited resource, then conduct migration and scale-out when more resources are available.
- *Customized execution environment*: each user may set his/her own configuration to run HPC programs in a customized way.

[1] H. Ma, LQ Wang, L. Wang, BC Tak, CQ Tang, and H. Huang. Auto-tuning Performance of MPI Programs in Cloud Using Dynamic Resource Management. Technical Report. 2014, University of Wyoming.

[2] H. Huang, LQ. Wang, BC Tak, L. Wang, and C. Tang. CAP3: A Cloud Auto-Provisioning Framework for Parallel Processing Using On-demand and Spot Instances. In the IEEE 6th International Conference on Cloud Computing (IEEE Cloud). 2013, IEEE Press.