Name
Stefan Georgiev
(Bulgaria)

Advisors
A.Univ.-Prof. DI Dr.
Wolfgang Schreiner

Company

Evaluation of Cluster Middleware in a Heterogeneous Computing Environment

Deploying a Beowulf-type high-performance cluster is a challenging task. Many problems, regarding the underlying hardware infrastructure and the used software components need to be solved. The task becomes even more complicated when heterogeneous commodity hardware platforms are utilized. This research is provoked by the idea of using a company's desktop computers for achieving computing performance at a low cost. A systematic analysis of available tools for clustering aims to reveal both their strong and their weak sides. The process of building a fully-functioning parallel environment is described in detail. Usability with regards to different tools like resource managers, schedulers, and MPI, is assessed taking into consideration the underlying hardware platforms.

High-Performance Computing
Clustering
Parallel Computing
ROCKS
CAOS
Cluster Middleware
OSCAR

Introduction

Ever since the invention of the first calculating machines the need for more processing power has been the driving force for further development and innovation. Today's computer hardware evolves rapidly in order to be able to keep up with the extreme demands of the applications and the users. Even though, in the fields of science and research there is still a growing need for more computational power. What is more, professional business applications are becoming more and more advanced and, thus, start requiring more computing resources as well. For years, the solution has been utilization of specialized massive computers that have proprietary hardware and architecture in order to achieve topmost computing performance. However, these machines are still rather costly to be used in everyday business life. An alternative solution, that proves to be also cheaper, is using the joint computing power of collection of more simple computers.

A collection of computers that work together towards achieving a common goal define a parallel environment. The main idea behind parallel environments is to split a heavy computational task into smaller ones and then distribute them to multiple processing units for calculation. Today's multi-core processors aim to achieve parallelism by dividing the workload of a single machine between the separate cores. While using a single machine proves to be rather effective, it also suffers from hardware constraints that limit the reachable performance.

Thus, an alternative solution is creating a parallel environment that consists of separate machines, each of which has its own processor and memory, connected together via means of networking. This architecture is referred to as a cluster of machines. What brings them together to work as a single computer is the middleware software. It is the gluing component that lies between the operating system and the user applications.

This thesis focuses on analyzing the high-performance capabilities of clusters built from commodity hardware machines that are of heterogeneous type. Recent research in the field of HPC shows that there are certain advantages in utilizing different computers in cluster environments.

Cluster middleware presents numerous challenges that need to be studied previous to deploying a production cluster. This thesis presents a solution for creating a fully-functional parallel environment by providing a systematic evaluation and comparison of different techniques and tools. Existing achievements in the field of HPC computing are studied. Appropriate tools were chosen for creating a heterogeneous computing environment based on different criteria like, for instance, ease of utilization and provided functionality. Cluster middleware in the form of the cluster deployment tools ROCKS, OSCAR, CAOS-NSA, is thoroughly analyzed. Detailed description of the experience gained during the installation process aims to provide a comparison between these tools and assess their qualities. Furthermore, usability of the achieved environments is tested by evaluating the tools for job distribution, resource management and resource monitoring included in the software of the clusters.

As a result, two separate clusters were built during the testing process. In accordance to the used hardware, numerous solutions and best practices are described. Test applications prove that the environments are usable for achieving peak performance at low prices.