A

Seminar report

On

Cluster Computing

Submitted in partial fulfillment of the requirement for the award of degree
Of MCA

SUBMITTED TO:
www.studymafia.org

SUBMITTED BY:
www.studymafia.org

Preface

I have made this report file on the topic **Cluster Computing**; I have tried my best to elucidate all the relevant detail to the topic to be included in the report. While in the beginning I have tried to give a general view about this topic.

My efforts and wholehearted co-corporation of each and everyone has ended on a successful note. I express my sincere gratitude to …………..who assisting me throughout the preparation of this topic. I thank him for providing me the reinforcement, confidence and most importantly the track for the topic whenever I needed it.

Acknowledgement

I would like to thank respected Mr…….. and Mr. ……..for giving me such a wonderful opportunity to expand my knowledge for my own branch and giving me guidelines to present a seminar report. It helped me a lot to realize of what we study for.

Secondly, I would like to thank my parents who patiently helped me as i went through my work and helped to modify and eliminate some of the irrelevant or un-necessary stuffs.

Thirdly, I would like to thank my friends who helped me to make my work more organized and well-stacked till the end.

Next, I would thank Microsoft for developing such a wonderful tool like MS Word. It helped my work a lot to remain error-free.

Last but clearly not the least, I would thank The Almighty for giving me strength to complete my report on time.

Content

- **Introduction**
- **Why cluster?**
- **Types of cluster :-**
  - **High Availability or Failover Clusters.**
  - **Load Balancing Cluster.**
  - **Parallel/Distributed Processing Clusters.**
- **Cluster Component.**
- **Conclusion**

INTRODUCTION

The computing power required by applications is increasing at a tremendous rate. Hence, the search has therefore been towards devising ever faster, ever more powerful computer systems, to help tackle more and more complex problems. In addition, parallel applications have become more and more complex with increasing processing power needs induced essentially by the progress registered in many fields (telecommunication, data-mining, *etc*). In short we can say trend is towards high performance computing systems. Our need for computational resources in all fields of science, engineering and commerce far outstrip our ability to fulfill these needs. The usage of clusters of computers is, perhaps, one of the most promising means by which we can bridge the gap between our needs and the available resources. The usage of a COTS-based cluster system has a number of advantages including:

Price/performance when compared to a dedicated parallel supercomputer.

Incremental growth that often matches yearly funding patterns.

The provision of a multi-purpose system: one that could, for example, be used for secretarial purposes during the day and as a *commodity parallel supercomputer* at night.

WHY CLUTERS?

The question may arise why clusters aredesigned and built when perfectly good commercial supercomputers areavailable on the market. The answer is that the latter is expensive. Clustersare surprisingly powerful. The supercomputer has come to play a larger rolein business applications. In areas from data mining to fault tolerant performance clustering technology has become increasingly important.Commercial products have their place, and there are perfectly good reasonsto buy a commerciallyproduced supercomputer.

If it is within our budget andour applications can keep machines busy all the time, we will also need tohave a data center to keep it in. then there is the budget to keep up with themaintenance and upgrades that will be required to keep our investment up to par. However, many who have a need to harness supercomputing power don’t buy supercomputers because they can’t afford them. Also it isimpossible to upgrade them.

Clusters, on the other hand, are cheap and easyway to take off-the-shelf components and combine them into a singlesupercomputer. In some areas of research clusters are actually faster thancommercial supercomputer. Clusters also have the distinct advantage in thatthey are simple to build using components available from hundreds of sources. We don't even have to use new equipment to build a cluster.

Price/Performance

The most obvious benefit of clusters, and the most compelling reason for thegrowth in their use, is that they have significantly reduced the cost of processing power.

One indication of this phenomenon is the Gordon BellAward for Price/Performance Achievement in Supercomputing, which manyof the last several years has been awarded to Beowulf type clusters. One of the most recent entries, the Avalon cluster at Los Alamos NationalLaboratory, "demonstrates price/performance an order of magnitude superior to commercial machines of equivalent performance.

" This reduction in thecost of entry to high-power computing (HPC) has been due to comodification of both hardware and software over the last 10 years particularly. All the components of computers have dropped dramatically inthat time. The components critical to the development of low cost clustersare:

1.**Processors**- commodity processors are now capable of computational power previously reserved for supercomputers, witness Apple Computer'srecent add campain touting the G4 Macintosh as a supercomputer.
2.**Memory** - the memory used by these processors has dropped in cost rightwith the processors.
3.**Networking Components** - the most recent group of products toexperience co modification and dramatic cost decreases is networkinghardware. High- Speed networks can now be assembled with these productsfor a fraction of the cost necessary only a few years ago.
4.**Motherboards, busses, and other sub-systems-** all of these have become commodity products, allowing the assembly of affordablecomputers from off the shelf components

Overview of Cluster

Cluster is a collection of inter-connected and loosely coupled stand-alone computers working together as a single, integrated computing resource. Clusters are commonly, but not always, connected through fast local area networks. Clusters are usually deployed to improve speed and/or reliability over that provided by a single computer, while typically being much more cost-effective than single computers of comparable speed or reliability.

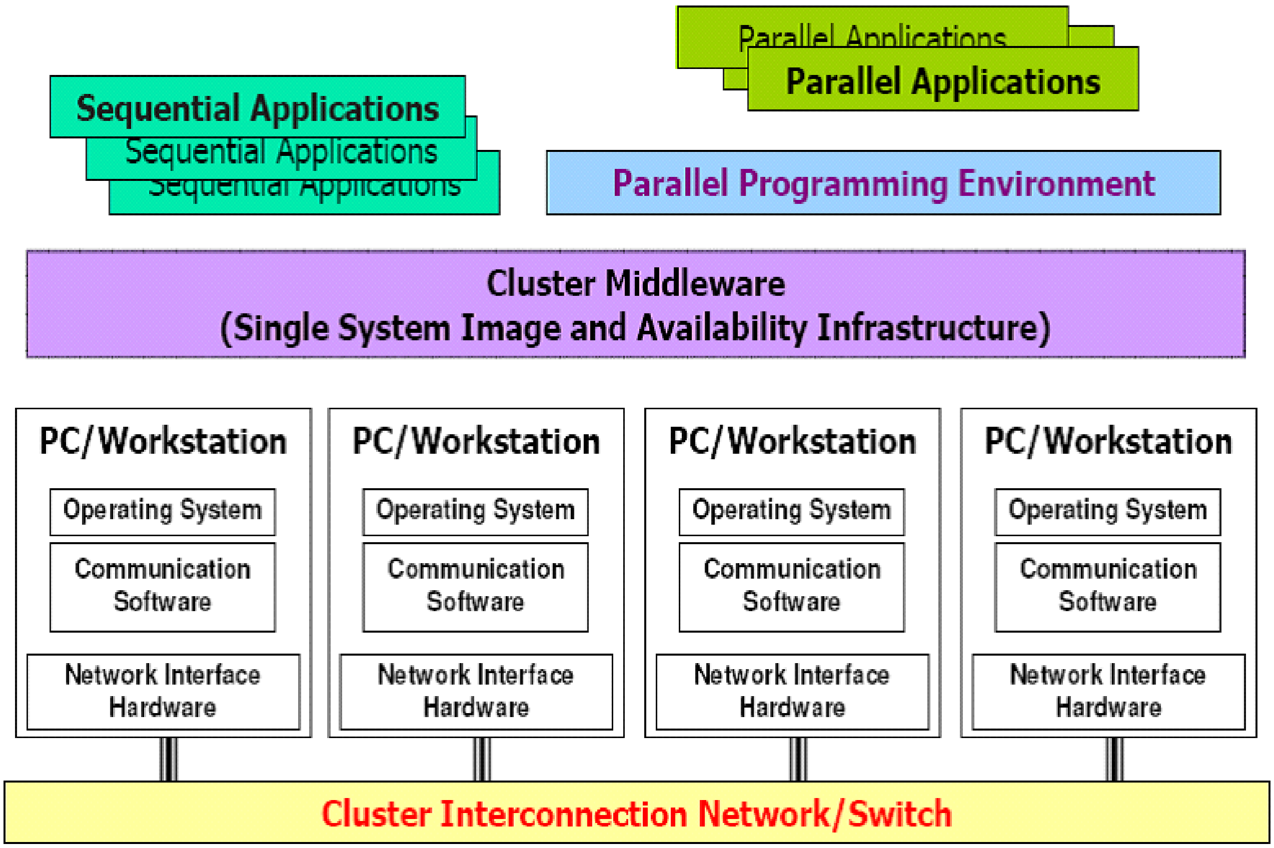

Figure 1: Cluster Architecture

Cluster Technology permits organizations to boost their processing power using standard technology (commodity hardware and software components) that can be acquired/purchased at a relatively low cost. This provides expandability- an affordable upgrade path that lets organizations increase their computing power- while preserving their existing investment and without incurring a lot of extra expenses. The performance of applications also improves with the support of scalable software . environment. In addition, failover capability allows a backup computer to take over the tasks of a failed computer located in its cluster.

COMPONENTS OF CLUSTER COMPUTER

1. **Multiple High Performance Computers**
a. PCs
b. Workstations
c. SMPs (CLUMPS)

2. **State of the art Operating Systems**
a. Linux (Beowulf)
b. Microsoft NT (Illinois HPVM)
c. SUN Solaris (Berkeley NOW)
d. HP UX (Illinois - PANDA)
e. OS gluing layers(Berkeley Glunix)

3. **High Performance Networks/Switches**
a. Ethernet (10Mbps),
b. Fast Ethernet (100Mbps),
c. Gigabit Ethernet (1Gbps)
d. Myrinet (1.2Gbps)
e. Digital Memory Channelf. FDDI

4. **Network Interface Card**
a. Myrinet has NIC
b. User-level access support

5. **Fast Communication Protocols and Services**
a. Active Messages (Berkeley)
b. Fast Messages (Illinois)
c. U-net (Cornell)d. XTP (Virginia)

6. **Cluster Middleware**
a. Single System Image (SSI)
b. System Availability (SA) Infrastructure

7. **Hardware**
a. DEC Memory Channel, DSM (Alewife, DASH), SMP Techniques

8. **Operating System Kernel/Gluing Layers**
a. Solaris MC, Unixware, GLUnix

9. **Applications and Subsystems**
a. Applications (system management and electronic forms)
b. Runtime systems (software DSM, PFS etc.)
c. Resource management and scheduling software (RMS)

10. **Parallel Programming Environments and Tools**
a. Threads (PCs, SMPs, NOW..)
b. MPI
c. PVMd. Software DSMs (Shmem)
e. Compilers
f. RAD (rapid application development tools)
g. Debuggersh.
Performance Analysis Tools
i. Visualization Tools

11. **Applications**
a. Sequential
b. Parallel / Distributed (Cluster-aware app.)

CLUSTER CLASSIFICATIONS

Clusters are classified in to several sections based on thefacts such as 1)Application target 2) Node owner ship 3)Node Hardware 4) Node operating System 5) Nodeconfiguration.Clusters based on Application Target are again classified intotwo:

•High Performance (HP) Clusters
•High Availability (HA) Clusters

Clusters based on Node Ownership are again classified intotwo:
•Dedicated clusters
Non-dedicated clusters

Clusters based on Node Hardware are again classified into three:
•Clusters of PCs (CoPs)
•Clusters of Workstations (COWs)

Clusters of SMPs (CLUMPs)Clusters based on Node Operating System are again classified into:
•Linux Clusters (e.g., Beowulf)
•Solaris Clusters (e.g., Berkeley NOW)
•Digital VMS Clusters
•HP-UX clusters

•Microsoft Wolfpack clustersClusters based on Node Configuration are again classifiedinto:

•Homogeneous Clusters -All nodes will have similararchitectures and run the same OSs
•Heterogeneous Clusters- All nodes will have differentarchitectures and run different Oss

ISSUES TO BE CONSIDERED

Cluster Networking

If you are mixing hardware that has different networking technologies, therewill be large differences in the speed with which data will be accessed andhow individual nodes can communicate. If it is in your budget make surethat all of the machines you want to include in your cluster have similar networking capabilities, and if at all possible, have network adapters fromthe same manufacturer.

Cluster Software

You will have to build versions of clustering software for each kind of system you include in your cluster.

Programming

Our code will have to be written to support the lowest common denominator for data types supported by the least powerful node in our cluster. Withmixed machines, the more powerful machines will have attributes thatcannot be attained in the powerful machine.

Timing

This is the most problematic aspect of heterogeneous cluster. Since thesemachines have different performance profile our code will execute at different rates on the different kinds of nodes. This can cause serious bottlenecks if a process on one node is waiting for results of a calculation ona slower node. The second kind of heterogeneous clusters is made fromdifferent machines in the same architectural family: e.g. a collection of Intel boxes where the machines are different generations or machines of samegeneration from different manufacturers.

Network Selection

There are a number of different kinds of network topologies, including buses, cubes of various degrees, and grids/meshes. These network topologieswill be implemented by use of one or more network interface cards, or NICs,installed into the head-node and compute nodes of our cluster.

Speed Selection

No matter what topology you choose for your cluster, you will want to getfastest network that your budget allows. Fortunately, the availability of highspeed computers has also forced the development of high speed networkingsystems. Examples are 10Mbit Ethernet, 100Mbit Ethernet, gigabitnetworking, channel bonding etc

Conclusion

As cluster sizes scale to satisfy growing computing needs in various industries as well as in academia, advanced schedulers can help maximize resource utilization and QoS. The profile of jobs, the nature of computation performed by the jobs, and the number of jobs submitted can help determine the benefits of using advanced schedulers. An important problem with traditional parallel job-scheduling algorithms is their specialization for specific types of workloads, which results in poor performance when the workload characteristics do not fit the model for which they were designed.

For example, batch and gang scheduling perform poorly under dynamic or load-imbalanced workloads, whereas implicit co-scheduling suffers from performance penalties for fine-grained synchronous jobs. Most job schedulers offer little adaptation to externally and internally fragmented workloads. The result is reduced machine utilization and response times. We still don't have a cluster scheduler that works equally well for all kinds of workloads. Hence, there is still a lot of scope of improvement to achieve high level of resource utilization while maintaining low job turnaround time.

A major area of ongoing and future research is locality based scheduling. That is, **scheduling based upon the topology of the interconnect**, which might include interconnects with a tree structure and will certainly include SMP building blocks. This type of scheduling will become even more important in the near future since it becomes increasingly difficult and expensive to build a flat interconnect as the cluster size grows.

In addition, new interconnect technologies are appearing which use loop, mesh and torus topologies. Improving Resource Matching through Estimation of Actual Job Requirements is another area which will command a lot of attention in the near future. Effective runtime estimation directly effects effectiveness of backfill algorithm and may even lead to improved system utilization by about 10-15% while maintaining the same average turnaround time.

REFERENCES

- www.google.com
- www.wikipedia.com
- www.studymafia.org