2009 REU in Telematics and Cyber Physical Systems

Performance Variations in Cloud Computing

Faculty Mentor: Dr. Cheng-Zhong Xu
Grad Student Mentor: Jia Rao

Claire Maki, University of Michigan, Ann Arbor, Michigan

Abstract:
Cloud computing, the long-held dream of computing as a utility, has the potential to transform a large part of the IT industry [1]. It is a cost-effective, scalable solution to traditional computing. Amazon and Google have already created their own clouds, and many more companies are predicted to follow. The following sections introduce the concept of cloud computing along with its advantages and challenges.

Introduction:
Traditional computing has relied on the computing power of the user's local machines to perform tasks. This is a major disadvantage, especially for smaller companies. Buying and maintaining servers can be costly. It can be difficult for companies who are just starting out to know the amount of resources they will need, and also be able to adjust their resources when demand increases or decreases.

Cloud computing refers to both the applications delivered as services over the Internet and the hardware and systems software in the datacenters that provide those services [1]. It has the potential to replace traditional computing and provide a flexible, on-demand computing infrastructure for numerous applications [2].

Background of Cloud Computing:
Cloud computing uses virtual machines to make use of the hardware on the cloud. A virtual machine is a software implementation of a machine. It can support individual processes or a complete system depending on the level of abstraction used [3].

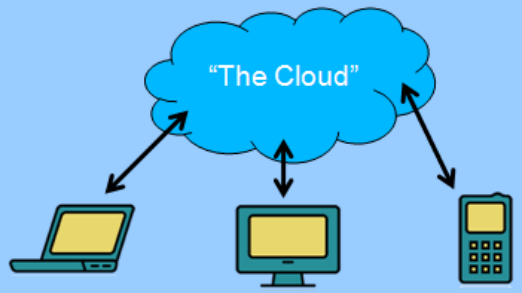

In the diagram above, a laptop computer, desktop computer, and mobile device are all using the resources in the cloud.

Advantages and Challenges of Cloud Computing:
There are many advantages to cloud computing:
- Use resources in a pay-as-you-go manner
- Easily expand resources if needed
- Give users access to higher computing power than on their local machine

While cloud computing may seem like the answer to all of our computing problems, it also has some challenges to overcome:
- Widespread Internet availability
- Guarantee data confidentiality
- Providing predictable performance in virtualized environments

Identifying Performance Variations:

Performance tends to be unpredictable in a cloud environment when applications from different users are competing for the same resources. Instead of performance being the same for both Virtual Machines, one gets better performance than the other. I have been looking into the performance unpredictability of cloud computing by running benchmark tests and comparing their performance in the cloud.

Benchmark Test Results:

The first benchmark I ran was called IOzone, which focuses on I/O performance. I ran the benchmark alone, and then simultaneously on two Virtual Machines. My results are in Figure 1 below:

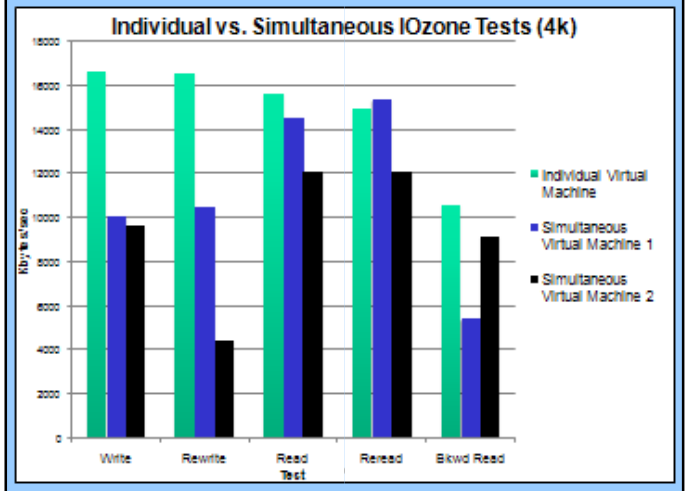

The figure above shows that the simultaneous Virtual Machines always have different performance. The Rewrite and Backward Read tests showed the biggest performance difference.

The next benchmark I ran was the NAS Parallel Benchmark (NPB). It is a collection of CPU-intensive tests designed to compare parallel computational performance. The three tests with the best results are shown in Figure 2 above.

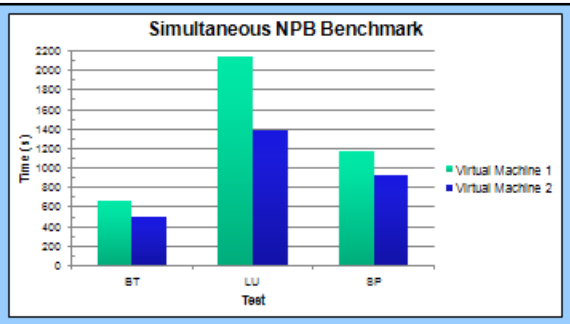

Figure 2. Performance Differences during Simultaneous NPB Benchmark

The percent difference in run time for the BT, LU, and SP tests were 25%, 35%, and 22% respectively. This difference is significant and needs to be addressed when usage charges are being calculated.

Discussion and Possible Solutions:
One possible reason for the performance problem is hard disk memory. Flash memory used in place of mechanical hard disks could decrease I/O interference [1]. Flash memory is much faster and uses less energy. This could decrease the amount of interference between Virtual Machines and reduce performance degradation.
L2-Cache sharing could also be a factor in performance differences. Figure 3 shows the L2-Cache miss rate during the NPB benchmark.

	BT Time	BT L2-Cache Miss Rate	LU Time	LU L2-Cache Miss Rate	SP Time	SP L2-Cache Miss Rate
CPU-Arrangement 1	680.6	**0.214**	2167.9	**0.283**	1143.7	**0.451**
CPU-Arrangement 2	568.0	0.013	1815.7	0.008	968.6	0.028

Figure 3. L2-Cache Miss Rate compared to NPB run time

I found that when the miss rate was higher, the performance was decreased. If we can monitor the cache-miss rate, we could use this information to charge users according to their performance.

Conclusions:

Overall, cloud computing has the potential to completely replace traditional computing. Various applications are being explored, including offloading data-intensive applications to the cloud from mobile devices to increase performance. Performance issues need to be more deeply explored for this technology to grow. In the future, however, I predict cloud computing will become the computing standard.

References:

[1] Michael Armbrust, Armando Fox, Rean Griffith, Anthony D. Joseph, Randy Katz, Andy Konwinski, Gunho Lee, David Patterson, Ariel Rabkin, Ion Stoica, and Matei Zaharia. Above the Clouds: A Berkeley View of Cloud Computing. Technical report, University of California at Berkeley, 2009.

[2] Christina Hoffa, Gaurang Mehta, Timothy Freeman, Ewa Deelman, Kate Keahey, Bruce Berriman, John Good. On the Use of Cloud Computing for Scientific Workflows. Indiana University, University of Southern California, Argonne National Laboratory, Caltech, 2008.

[3] Smith, Daniel E.; Nair, Ravi (2005). "The Architecture of Virtual Machines". Computer (IEEE Computer Society) 38 (5): 32–38.