Virtual Machine Placement in Predictable Computing Clouds

Richard Rauscher
Col. of Medicine and Dept. of Computer Science & Eng.
Pennsylvania State University
Hershey, Pennsylvania, USA
rauscher@psu.edu

Raj Acharya
Department of Computer Science & Engineering
Pennsylvania State University
University Park, Pennsylvania, USA
acharya@cse.psu.edu

***Abstract*—Literature about cloud computing often makes the assumption that the resource demands of computing clouds (and the virtual machines that constitute them) are unpredictable in the short term. There are, however, specific use cases where resource demands can be anticipated. This paper discusses dissertation work-in-progress which shows that, in certain predictable environments, preemptive virtual machine migration can improve both computational resource utilization and the overall user experience. A novel algorithm which reacts to anticipated future resource demands based on past behavior of virtual machines is presented. Simulations are used to quantify performance improvements.**

Keywords—cloud computing, health care information systems, electronic health records, private cloud

I. Introduction

The virtual machine mapping problem (VMMP) is reducible to the bin packing problem which is known to be NP-hard. As such, several heuristically-based solutions have been proposed to introduce a tractable yet useful system to ensure efficient resource utilization and improve user-perceived performance.

Much of the literature about cloud computing assumes that resource demands are not near-term predictable and that cloud computing is best suited for unpredictable resource demands [1]. However, it has been shown that predictability is a characteristic of some types of workloads [2]. We specifically found that this is the case in small-scale systems designed to support private clouds in health care. Hospital-based health care is an around-the-clock operation with well-defined patterns of care. Ambulatory, low-acuity clinics operate with standard office hours. Time-of-day and day-of-year patterns in care provisioning cause reasonably predictable patterns of computer usage.

At its core, the VMMP is a multivariable optimization problem that has historically been attacked using different approaches. There are many algorithms that seek to optimize different aspects of the computing cloud. We have, however, found no algorithms that consider the long-term predictability of the future resource demands as a means of determining virtual-to-physical machine mapping. Consider the trivial example of electronic medical record system that is divided into an application logic server and a database server. Let's assume that we have a virtual machine mapping algorithm (VMMA) that considers inter-VM communications as the primary determinant of virtual machine mapping (VMM). For most of the day, that database server is located proximate to its corresponding application server (see Figure 1, t_0). Each day, at nearly the same time of the day, the content of the database is extracted into a central data warehouse. When the extraction process begins, communications between the database virtual server and the data warehouse virtual server increase significantly. The VMMA then relocates the database virtual server to the physical machine that hosts the data warehouse virtual server (see Figure 1, t_1). Subsequent to the data extraction process, the VMMA relocates the database virtual server back to the physical server that houses its corresponding application virtual server (see Figure 1, t_2). Given foresight, a VMMP that can conceptualize past history may select to migrate the VM that is about to be extracted *before any real-time indicators suggest it should do so.*

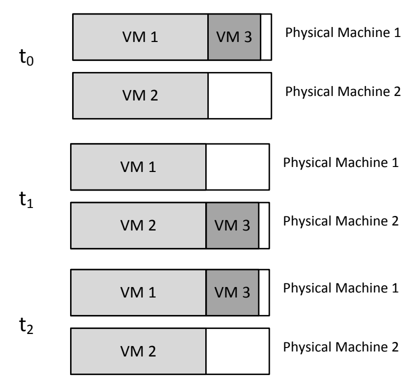

Figure 1: Migration based on behavior

II. Background

A. Load Balancing Approaches

There are several approaches to distributing virtual machines among a set of physical machines. Due to space restrictions, we refer to [3] which provides an excellent survey on the subject of diverse approaches to virtual machine mapping algorithms. To our knowledge, none consider long-term patterns of resource utilization when mapping virtual to physical machines.

III. Novel Algorithm

Most modern VMMAs consider recent activity when making mapping decisions [3]. These algorithms are excellent for unpredictable resource demands but ignore useful information when machine behavior is predictable. Thus, we have created a new algorithm that looks ahead based on past history. Rather than waiting for the hypervisor to react to new resource demands, the algorithm introduces a "look ahead window" which evaluates the current state of the resource demands compared to historical norms and reacts based on the probable resource demands of the entire system. We've called this algorithm "Look Ahead and React Based on History" or "LARBOH".

A. Computational Overhead of LARBOH

LARBOH requires additional computation by the hypervisor as well as historic state information. The hypervisor must be able to access a record of previous performance characteristics of each virtual machine running within the cloud. Depending upon the level of temporal resolution and the cycle duration, this will result in varying sizes of datasets. In our implementation, the computational overhead of looking ahead was of the same order of complexity as current VMMAs. The storage required to manage past history was also not significant as we focused on a period of one day with a sample frequency of one minute. Furthermore, we implemented LARBOH in our simulations such that it only had seven days of history.

IV. CHARACTERIZATION OF PREDICTABILITY

A basic assumption of this research is that there exist systems that have predictable resource utilization behaviors. Stokely et al examined the predictability of resource usage in a cloud environment [4]. They found that aggregate behavior among many users was predictable with reasonable accuracy (+/- 12%). We have had access to resource utilization data at diverse academic health centers throughout the eastern United States (Tampa, Detroit, Boston and Hershey, PA). At these centers, past resource utilization was consistently a good predictor of future utilization. The following data are anecdotal but we contend that it is characteristic of health care resource utilization. During December, 2013, the same-day (week-to-week) differences in the number of users on the electronic medical record system varied by less than 1% (0.463%) for non-holiday weeks for a population of providers of more than 2000.

V. EXPERIMENT

A. Trace Creation

We constructed a simulation that used traces of actual performance of different component virtual machines of an electronic medical record (EMR) system. We used OpenEMR as the sample EMR. The traces were generated using the same configuration as described in [2]. We varied the number of simulated client machines using Apache Jmeter. We used SNMP to query each virtual machine regarding load average, RAM usage, RAM allocated to file caching (so that file caching was not a confounder of memory utilization), network bandwidth utilization and disk I/O operations per second. The results from the SNMP queries were written to text files and were used as the source of trace data that was fed into the subsequent simulation. RAM was ultimately the most contentious resource during our simulations. The demand for RAM scaled linearly with the number of simultaneous users.

B. Simulation

We constructed a simulation written in C to examine how the timing of the migration of virtual machines between different physical machines impacted the utilization of the physical machines and the responsiveness of the virtual machines. We also examined how the timing of the migration affected user-perceived performance as well as the complexity of the migration. As we indicated previously, the use of RAM in our model systems scaled linearly with more users. The amount of RAM allocated is also a prime determinant of the computational cost of migration. To simplify the experiment, we relied on a single threshold variable: RAM utilization. Although others have successfully integrated multiple variables into the migration decision [5], it was sufficient to use one key resource as a trigger for migration to demonstrate that there is value in preemptive migration of virtual machines. We measured the impact on users as the product of the number of users and the amount of time required to migrate the virtual machine.

VI. RESULTS

Our simulation demonstrated that preemptive migration of virtual machines in anticipation of changes in resource requirements can improve the operation of the cloud system while introducing only a slight increase in computational overhead for the hypervisor. Depending upon the specific states of the machines and in highly predictable systems, LARBOH was able to reduce the user impact of virtual machine migration by up to 99.7%.

VII. FUTURE WORK

There were several arbitrary constants used in the initial algorithm that may be opportunities for optimization. Specifically, the look-ahead time range and the duration of history considered should be examined. Furthermore, it would be useful to quantify the definition of "predictable" and specify the level of predictability that makes LARBOH useful.

VIII. CONCLUSION

Preemptive positioning of virtual machines within a cloud, given a reasonably predictable environment, results in improved user-perceived performance.

REFERENCES

[1] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. H. Katz, A. Konwinski, G. Lee, D. A. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, "Above the Clouds: A Berkeley View of Cloud Computing," EECS Department, University of California, Berkeley UCB/EECS-2009-28, February 10 2009.

[2] R. Rauscher and R. Acharya, "Performance of Private Clouds in Health Care Organizations," in *CloudCom 2013* Bristol, UK, 2013.

[3] K. A. Nuaimi, N. Mohamed, M. A. Nuaimi, and J. Al-Jaroodi, "A Survey of Load Balancing in Cloud Computing: Challenges and Algorithms," Los Alamitos, CA, USA, 2012, pp. 137-42.

[4] M. Stokely, A. Mehrabian, C. Albrecht, F. Labelle, and A. Merchant, "Projecting disk usage based on historical trends in a cloud environment," in *Proceedings of the 3rd workshop on Scientific Cloud Computing Date* Delft, The Netherlands: ACM.

[5] A. Ankit, J. Lakshmi, and S. K. Nandy, "Virtual Machine Placement Optimization Supporting Performance SLAs," in *CloudCom 2013*, Bristol, UK, 2013.