Partnership for International Research and Education
A Global Living Laboratory for Cyberinfrastructure Application Enablement

Development of a Simulator for Scheduling on Virtual Machines

Student: Javier Delgado, PhD student, Florida International University (FIU)
FIU Advisors: S. Masoud Sadjadi, PhD, FIU and Dr. Malek Adjouadi, PhD, FIU.
PIRE Local Industry Partner Advisor: Liana Fong, IBM T.J. Watson Research Lab
PIRE International Partner Advisor: Vijay Mann, IBM India Research Lab

National Science Foundation

I. Research Overview and Outcome

Abstract

Traditionally, job schedulers have been designed to assign jobs to physical machines. Recent advances in virtual machine technology have made execution of most applications on virtual machines a considerable option. Virtual machine technologies have several advantages that may help with running jobs remotely, but they introduce challenges. In this work, we experiment with an algorithm for scheduling jobs on virtual machines. A simulation environment to test this IBM technology was implemented. The results thus far indicate that this methodology can improve resource utilization and job turn-around time.

Benefits and Problems

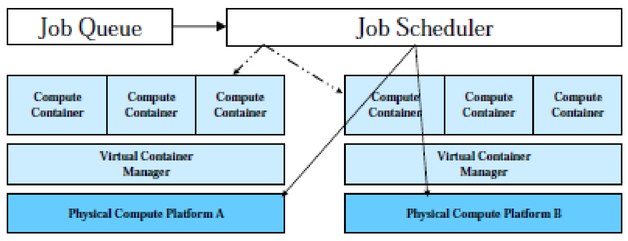

There are several benefits to be gained by using virtual machines (VMs):
•Reliable partitioning of resources in a system (e.g. to different jobs).
•Novel scheduling methods are possible (e.g. [1])
•Built-in checkpoint/restart/migration capability
•Ease deployment of software on remote systems
On the other hand, it is necessary to
•Minimize Contention
•Address resource competition issue
•Give different jobs different priorities so that certain jobs do not cause delays to one or more other jobs
•Mitigate VM overhead (although its not really too high [2])

Problem: Picking the best resource container (e.g. VM) to run a job. The example below shows how allocating a new job (D) can affect the run time of existing jobs (A-C), since their available share will be decreased.

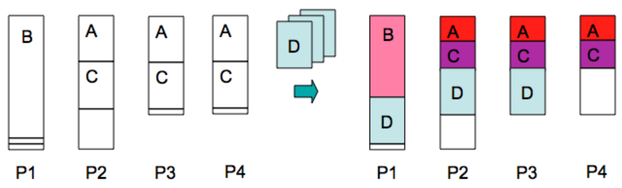

Solution: Send jobs to VMs based on their "potential free capacity," which takes into account how much physical capacity is available in the underlying physical machine (and bound by the VM's constraints).

Implementation and Evaluation

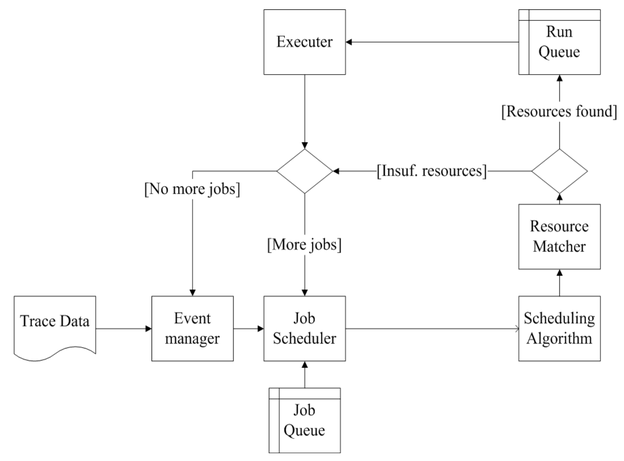

- Event-based simulation environment for scheduling jobs on VMs was created using MATLAB.
- Grid'5000 workload archive trace file, containing job data spanning approximately 1.5 years, was used to test simulator and algorithms
- Virtual data center of 30-90 physical machines, with 4 VMs per machine, was used.

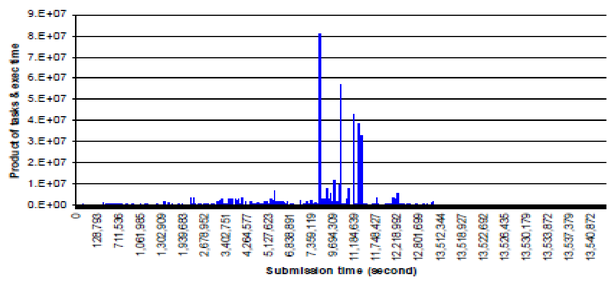

3 **node selection algorithms** were implemented.
•**Potential free capacity** (**PFC**) – Select the node(s) with the highest "potential free capacity"
•**PM then VM (PMVM)** – Select the physical machine with the highest free capacity, then select the VM with highest weight
•**Baseline/Fixed VM** – Provide a static, equal share to all VMs
Three **performance metrics**:
•**Average utilization** – The overall CPU utilization of the data center
•**Average queue time** – The average amount of time that a job spends waiting in the queue
•**Average expansion factor (*xfactor*)** –

$$\frac{wall\_clock\_runtime + queue\_time}{runtime\_at\_full\_cpu}$$

Results

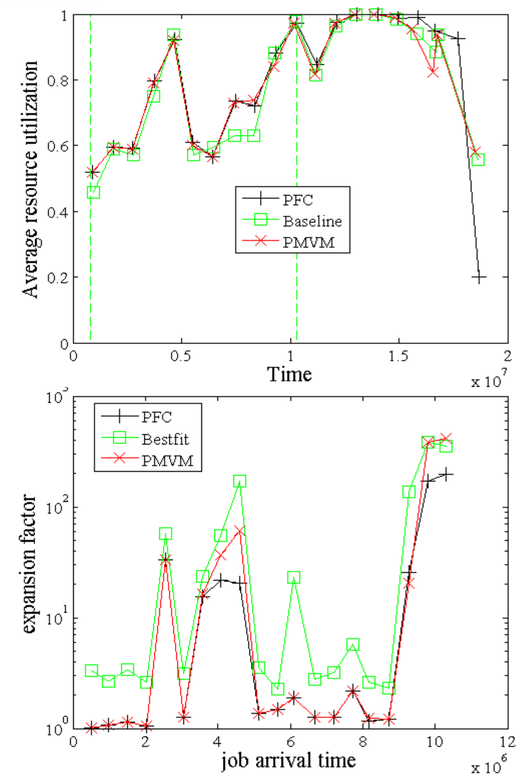

Conclusion

•Our algorithm schedules jobs on *resource containers* using an IBM-developed, non-trivial placement algorithm
•Current results show promise for improving system utilization and job turn-around time

Future Work

•Test with other trace data.
•Add support for multicore scheduling.

References

[1] M. Stillwell, F. Vivien, H. Casanova, *Dynamic Fractional Resource Scheduling for HPC Workloads*. in Proceedings of the 24th IEEE International Parallel and Distributed Processing Symposium (IPDPS'10), April 2010.

[2] Lamia Youseff, Rich Wolski, Brent Gorda, Chandra Krintz, *Paravirtualization for HPC Systems*, Workshop on XEN in HPC Cluster and Grid Computing Environments (XHPC), held in conjunction with The International Symposium on Parallel and Distributed Processing and Application (ISPA 2006), December 2006.

II. International/Industry Experience

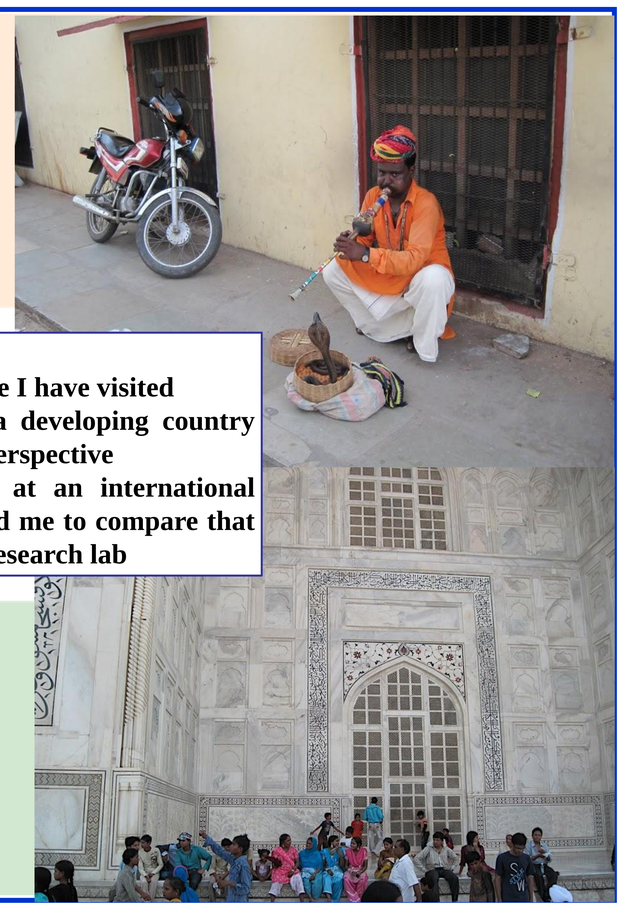

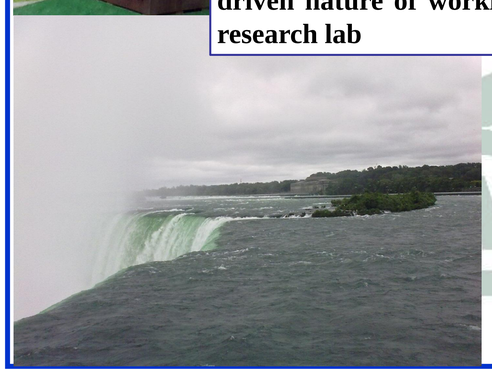

Personal Impact

•This was my first time implementing technology devised by IBM R&D
•Learned about the competitiveness and deadline-driven nature of working on a project in an industry research lab

Personal Impact

•India is the most interesting place I have visited
•Experiencing the problems of a developing country first hand really put things into perspective
•It is also my first experience at an international industry research lab and allowed me to compare that lifestyle to that of the New York research lab

III. Acknowledgement

The material presented in this poster is based upon the work supported by the National Science Foundation under Grant No. OISE-0730065. The author appreciates NSF grant HRD-0833093 for supporting me throughout my research. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.