

An Optimization Problem in Adaptive Virtual Environments

Ananth I. Sundararaj Manan Sanghi Jack R. Lange Peter A. Dinda
{ais,manan,jarusl,pdinda}@cs.northwestern.edu
Department of Computer Science, Northwestern University

Abstract

A virtual execution environment consisting of virtual machines (VMs) interconnected with virtual networks provides opportunities to dynamically optimize, at run-time, the performance of existing, *unmodified* distributed applications without any user or programmer intervention. Along with resource monitoring and inference and application-independent adaptation mechanisms, efficient adaptation algorithms are key to the success of such an effort. In previous work we have described our measurement and inference framework and also the adaptation mechanisms at hand and proposed simple heuristics as adaptation algorithms. Though we were successful in improving performance as compared to the case with no adaptation, none of them were characterized by theoretically proven bounds. In this paper, we formalize the adaptation problem, show that it is NP-hard and propose research directions for coming up with an efficient solution.

1 Introduction

Virtual machines greatly simplify wide-area distributed computing by lowering the abstraction to benefit both resource users and providers [1, 2, 4]. We have been developing a middleware system, Virtuoso, for virtual machine grid computing¹. Virtuoso, for a user, very closely emulates the existing process of buying, configuring, and using a computer or a collection of computers from a web site. Instead of a physical computer, the user receives a reference to the virtual machine which he can then use to start, stop, reset, and clone the machine.

The nature of the network presence that the virtual machine gets depends solely on the policies of the remote site. To deal with this network problem we developed VNET [10], a simple layer two virtual network tool. VNET is ideally placed to monitor the resource demands of the VMs. The VTTIF (Virtual Topology and Traffic Inference Framework) component of Virtuoso achieves this [3]. In addition, the execution environment can use the naturally occurring traffic of existing, unmodified applications running inside of the VMs to measure the characteristics of the underlying physical network [9].

2 Dynamic adaptation problem in virtual execution environments

Any application running in a distributed environment must adapt to the continuously changing network and computing resources. Despite many efforts adaptation has remained very application specific [7, 8].

A virtual execution environment, such as Virtuoso, provides an opportunity to dynamically optimize, at run-time, the performance of existing, *unmodified* distributed applications running on existing, unmodified operating systems without any user or programmer intervention. However, a number of challenges must first be met. Figure 1 illustrates a simplified adaptation scenario wherein a greedy heuristic drives overlay topology and routing changes leveraging data inferred by VTTIF.

Measurement and inference: This involves (a) measuring the traffic load and topology of applications running inside the virtual machines, (b) monitoring the underlying network and inferring its topology and bandwidth and latency characteristics, and (c) measuring host and VM characteristics such as their size and compute capacities and demands. In previous work [12, 9] we have shown how to successfully accomplish these.

¹<http://virtuoso.cs.northwestern.edu>

Adaptation mechanisms: A wide variety of adaptation mechanisms are made possible in the context of virtual execution environments such as (a) VM migration, (b) overlay topology and routing changes, and (c) network and CPU resource reservation where possible. These have been described previously [11, 12, 6].

Adaptation algorithm: Most importantly, we need an efficient algorithm to drive the adaptation mechanisms while guided by the measured and inferred data. Note that there may be a variety of reasonable adaptation goals in terms of latency, throughput, congestion, workload etc. For this paper we formulate the problem with the goal of maximizing application throughput.

VNET monitors the underlying network and provides a directed VNET topology graph, $G = (H, E)$, where H are VNET nodes (hosts running VNET daemons and capable of supporting one or more VMs) and E are the possible VNET links. Note that this may not be a complete graph as many links may not be possible due to particular network management and security policies at different network sites. VNET also provides estimates for the available bandwidth and latencies over each link in the VNET topology graph. These estimates are described by a bandwidth capacity function, $bw : E \rightarrow \mathbb{R}$, and a latency function, $lat : E \rightarrow \mathbb{R}$.

In addition, VNET also collects information regarding the space capacity (in bytes) and compute capacity made available by each host, described by a host compute capacity function, $compute : H \rightarrow \mathbb{R}$ and a host space capacity function, $size : H \rightarrow \mathbb{R}$. The set of virtual machines participating in the application is denoted by the set VM . The size and compute capacity demand made by every VM is also estimated, denoted by a VM compute demand function, $vm_compute : VM \rightarrow \mathbb{R}$ and a VM space demand function, $vm_size : VM \rightarrow \mathbb{R}$, respectively.

VTTIF infers the application communication topology to generate the traffic requirements of the application, \mathcal{A} , which is a set of 4-tuples, $A_i = (s_i, d_i, b_i, l_i)$, $i = 1, 2, \dots, m$, where s_i is the source VM, d_i is the destination VM, b_i is the bandwidth demand between the source destination pair and l_i is the latency demand between the source destination pair.

The goal is to find an adaptation algorithm that uses the measured and inferred data and drives the adaptation mechanisms at hand to improve application throughput. In other words we wish to find

- a mapping from VMs to hosts, $vmap : VM \rightarrow H$, meeting the size and compute capacity demands of the VMs within the host constraints. Further, we may also be given a set of constraint mappings from VMs to hosts that have to be maintained at all times, represented by a set of ordered pairs $M_i = (vm_i, h_i)$, $vm_i \in VM$, $h_i \in H$.
- a routing, $R : \mathcal{A} \rightarrow \mathcal{P}$, where \mathcal{P} is the set of all paths in the graph $G = (H, E)$, i.e. for every 4-tuple, $A_i = (s_i, d_i, b_i, l_i)$, allocate a path, $p(vmap(s_i), vmap(d_i))$, over the overlay graph, G , meeting the application demands while satisfying the bandwidth and latency constraints of the network.

Once all the mappings and paths have been decided, each VNET edge will have a residual capacity, rc_e , which is the bandwidth remaining un-utilized on that edge, in that direction.

$$rc_e = bw_e - \sum_{e \in R(A_i)} b_i$$

For each mapped path, $R(A_i)$, we define its bottleneck bandwidth, $bb(R(A_i)) = \min_{e \in R(A_i)} \{rc_e\}$ and its total latency, $tl(R(A_i)) = \sum_{e \in R(A_i)} (lat_e)$

The aim of the adaptation algorithm is to maximize the sum of bottleneck bandwidths over each mapped path. The intuition behind this objective function is to leave the most room for the application to increase performance within the current configuration thereby increasing application throughput.

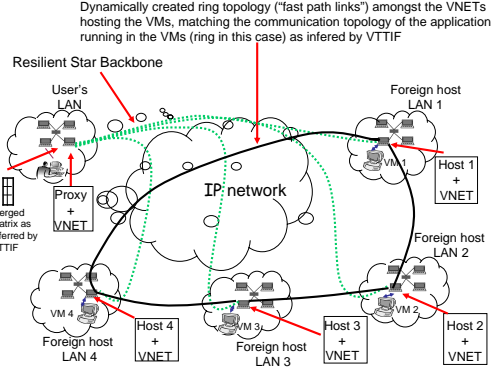


Figure 1: As the application progresses VNET adapts its overlay topology to match that of the application communication as inferred by VTTIF leading to a significant improvement in application performance, without any participation from the user.

Problem 1 (Generic Adaptation Problem In Virtual Execution Environments (GAPVEE))

INPUT:

- A directed graph $G = (H, E)$
- A function $\text{bw} : E \rightarrow \mathbb{R}$
- A function $\text{lat} : E \rightarrow \mathbb{R}$
- A function $\text{compute} : H \rightarrow \mathbb{R}$
- A function $\text{size} : H \rightarrow \mathbb{R}$
- A set, $\text{VM} = (\text{vm}_1, \text{vm}_2 \dots \text{vm}_n)$, $n \in \mathbb{N}$
- A function $\text{vm_compute} : \text{VM} \rightarrow \mathbb{R}$
- A function $\text{vm_size} : \text{VM} \rightarrow \mathbb{R}$
- A set of ordered 4-tuples $\mathcal{A} = \{(s_i, d_i, b_i, l_i) \mid s_i, d_i \in \text{VM}; b_i, l_i \in \mathbb{R}; i = 1, \dots, m\}$
- A set of ordered pairs $\mathcal{M} = \{(\text{vm}_i, h_i) \mid \text{vm}_i \in \text{VM}, h_i \in H; i = 1, 2 \dots r, r \leq n\}$

OUTPUT: $\text{vmap} : \text{VM} \rightarrow H$ and $R : \mathcal{A} \rightarrow \mathcal{P}$ such that

- $\sum_{\text{vmap}(\text{vm})=h} (\text{vm_compute}(\text{vm})) \leq \text{compute}(h), \forall h \in H$
- $\sum_{\text{vmap}(\text{vm})=h} (\text{vm_size}(\text{vm})) \leq \text{size}(h), \forall h \in H$
- $h_i = \text{vmap}(\text{vm}_i), \forall M_i = (\text{vm}_i, h_i) \in \mathcal{M}$
- $(\text{bw}_e - \sum_{e \in R(A_i)} b_i) \geq 0, \forall e \in E$
- $(\sum_{e \in R(A_i)} \text{lat}_e) \leq l_i, \forall e \in E$
- $\sum_{i=1}^m (\min_{e \in R(A_i)} \{rc_e\})$, where $rc_e = (\text{bw}_e - \sum_{e \in R(A_i)} b_i)$, is maximized

3 A special case of the adaptation problem

The generic adaptation problem seeks a mapping, vmap from VMs to hosts and routing, R of VM traffic over the overlay network, G . To establish the hardness of the problem, we consider a special case of the problem wherein all the VM to host mappings are constrained by the ordered pairs \mathcal{M} and latency demands are dropped, leaving us only with the routing problem.

Since the mappings are pre-defined, we can formulate the problem in terms of only the hosts and exclude all VMs. Also, as the latency demands have been dropped, the application 4-tuple reduces to 3-tuple, $A_i = (s_i, d_i, b_i)$, $s_i, d_i \in H$, $b_i \in \mathbb{R}$, $i = 1, 2 \dots m$. Notice that now $s_i, d_i \in H$ as VM to host mappings are fixed and VMs are synonymous with the hosts that they are mapped to.

Problem 2 (Routing Problem In Virtual Execution Environments (RPVEE))

INPUT:

- A directed graph $G = (H, E)$
- A function $\text{bw} : E \rightarrow \mathbb{R}$
- A set of ordered 3-tuples $\mathcal{A} = \{(s_i, d_i, b_i) \mid s_i, d_i \in H; b_i \in \mathbb{R}; i = 1, \dots, m\}$

OUTPUT: $R : \mathcal{A} \rightarrow \mathcal{P}$ such that

- $(\text{bw}_e) - (\sum_{e \in R(A_i)} b_i) \geq 0, \forall e \in E$,
- $\sum_{i=1}^m (\min_{e \in R(A_i)} \{rc_e\})$, where $rc_e = (\text{bw}_e - \sum_{e \in R(A_i)} b_i)$, is maximized

4 Analysis

Theorem 1 *RPVEE is NP-hard.*

The NP-hardness for the problem is established by reduction from the Edge Disjoint Path Problem (EDPP) which is shown to be NP-complete in [5]. In the interest of space we provide only a brief sketch of the reduction here.²

The input to EDPP consists of a graph G and a set of source destination pairs $S = \{(s_1, t_1), \dots, (s_n, t_n)\}$ and the problem is to determine whether there exist edge disjoint paths in $G = (V, E)$ for each source destination pair in S .

For reducing EDPP to an instance of the decision version of RPVEE, construct a complete graph $G' = (V, E')$ where $\text{bw}((u, v)) = 1 + \epsilon$ if $(u, v) \in E$ and $\text{bw}((u, v)) = 1$ if $(u, v) \notin E$. Further for all $(s_i, t_i) \in S$, let $(s_i, t_i, 1) \in \mathcal{A}$. Note

²The complete proofs for Theorem 1 and Theorem 2 can be found on our website: <http://virtuoso.cs.northwestern.edu/mama-proofs-rev1.pdf>

that there exists edge disjoint paths for the EDPP if and only if the sum of bottleneck bandwidths in the instance of RPVEE is greater than $k \cdot \epsilon$.

Since GAPVEE is a special case of RPVEE, the following theorem immediately follows.

Theorem 2 *GAPVEE is NP-hard.*³

5 Status

We have previously developed a variety of heuristics to drive the adaptation mechanisms [6, 12]. Though they were successful in improving performance relative to the naive approach (with no adaptation), we believe there is a significant scope for improvement. Therefore, as a first step, we have formalized the adaptation problem and given preliminary results of its hardness.

Even though the problem in this most generic incarnation is computationally hard, special cases of the problem amenable to efficient solutions will be of significant interest as we have a working system wherein various adaptation algorithms can be deployed and studied very conveniently.

There are a number of well studied variants of our problem such as routing un-splittable flows. We believe that the rich collection of approximation algorithms already available for them can be adapted to our specific problem. Accordingly, we have begun work in that direction.

6 Conclusion

We formalize the adaptation problem that arises in virtual execution environments consisting of virtual machines inter-connected by virtual networks. Such a platform provides opportunities to dynamically optimize, at run-time, the performance of existing, *unmodified* distributed applications running on existing, unmodified operating systems without any user or programmer intervention. We show that the adaptation problem is NP-hard and propose potential research directions for coming up with efficient solutions to certain interesting special cases.

References

- [1] FIGUEIREDO, R., DINDA, P. A., AND FORTES, J. A case for grid computing on virtual machines. In *Proceedings of the 23rd International Conference on Distributed Computing Systems (ICDCS)* (May 2003).
- [2] GARFINKEL, T., PFAFF, B., CHOW, J., ROSENBLUM, M., AND BONEH, D. Terra: A virtual machine-based platform for trusted computing. In *Proceedings of the 19th ACM symposium on Operating systems principles SOSP* (2003), pp. 193–206.
- [3] GUPTA, A., AND DINDA, P. A. Inferring the topology and traffic load of parallel programs running in a virtual machine environment. In *Proceedings of the 10th Workshop on Job Scheduling Policies for Parallel Program Processing(JSPPP)* (June 2004).
- [4] JIANG, X., AND XU, D. Soda: A service-on-demand architecture for application service hosting platforms. In *Proceedings of the 12th IEEE International Symposium on High Performance Distributed Computing (HPDC)* (June 2003), pp. 174–183.
- [5] KARP, R. *Complexity of Computer Computations*. Miller, R.E. and Thatcher, J.W. (Eds.). Plenum Press, New York, 1972, ch. Reducibility among combinatorial problems, pp. 85–103.
- [6] LANGE, J. R., SUNDARARAJ, A. I., AND DINDA, P. A. Automatic dynamic run-time optical network reservations. In *Proceedings of the Fourteenth International Symposium on High Performance Distributed Computing (HPDC)* (July 2005). To appear.
- [7] NOBLE, B. D., SATYANARAYANAN, M., NARAYANAN, D., TILTON, J. E., FLINN, J., AND WALKER, K. R. Agile application-aware adaptation for mobility. In *Proceedings of the 16th ACM Symposium on Operating Systems Principles (SOSP)* (1997).
- [8] SIEGELL, B., AND STEENKISTE, P. Automatic generation of parallel programs with dynamic load balancing. In *Proceedings of the Third International Symposium on High-Performance Distributed Computing (HPDC)* (August 1994), pp. 166–175.
- [9] SUNADARARAJ, A. I., ZANGRILLI, M., GUPTA, A., DINDA, P., AND LOWEKAMP, B. Free network measurement for virtual machine distributed computing. In submission.
- [10] SUNDARARAJ, A. I., AND DINDA, P. A. Towards virtual networks for virtual machine grid computing. In *Proceedings of the 3rd USENIX Virtual Machine Research and Technology Symposium (VM 04)* (May 2004).
- [11] SUNDARARAJ, A. I., GUPTA, A., AND DINDA, P. A. Dynamic topology adaptation of virtual networks of virtual machines. In *Proceedings of the Seventh Workshop on Languages, Compilers and Run-time Support for Scalable Systems (LCR)* (October 2004).
- [12] SUNDARARAJ, A. I., GUPTA, A., AND DINDA, P. A. Increasing application performance in virtual environments through run-time inference and adaptation. In *Proceedings of the Fourteenth International Symposium on High Performance Distributed Computing (HPDC)* (July 2005). To appear.

³The complete proofs for Theorem 1 and Theorem 2 can be found on our website: <http://virtuoso.cs.northwestern.edu/mama-proofs-rev1.pdf>