An Opportunity Cost Approach for Job Assignment and Reassignment in a Scalable Computing Cluster

Yair Amir, Baruch Awerbuch, Amnon Barak, R. Sean Borgstrom and Arie Keren[1][2]

Abstract.

A new method is presented for job assignment to and reassignment between machines in a computing cluster. Our method is based on a theoretical framework that has been experimentally tested and shown to be useful in practice. This "opportunity cost" method converts the usage of several heterogeneous resources in a machine to a single homogeneous "cost." Assignment and reassignment is then performed based on that cost. This is in contrast to previous methods for job assignment and reassignment, which treat each resource as an independent entity with its own constraints. These previous methods were intrinsically *ad hoc*, as there was no clean way to balance one resource against another.

1. Introduction

The more powerful a cluster of workstations is, the more important it is to use its resources wisely. A poor job assignment strategy can result in heavily unbalanced loads and thrashing machines, which cripples the cluster's computational power. Resources can be used more efficiently if the cluster can migrate jobs – moving them transparently from one machine to another. However, even systems that can reassign jobs can still benefit from a carefully-chosen assignment strategy.

Job migration is attractive because the arrival rate and resource demands of incoming jobs are unpredictable. In the face of this unpredictability, jobs will sometimes be assigned to a non-optimal machine, and migration gives the system a second (or third, etc.) chance to fix such a mistake. It is intuitively clear that the ability to migrate jobs could lead to better performance – that is, faster completion times for the average job. Unless it is known where a job *should* be at any given time, however, the reassignment strategy could also make mistakes. The Mosix [BGW93, BL97] kernel enhancements to the BSD/OS Unix-like operating system [Bsdi], for example, allow this kind of transparent job migration.

Determining the optimal location for a job is a complicated problem. The most important complication is that the resources available on a cluster of workstations are heterogeneous. In effect, the costs for memory, CPU, process communication and so forth are *incomparable*. They are not even measured in the same units: communication resources are measured in terms of bandwidth, memory in terms of space, and CPU in terms of cycles. The natural greedy strategy, balancing the resources across all of the machines, is not even well defined.

---

[1] Yair Amir, Baruch Awerbuch and Ryan Sean Borgstrom are with the Department of Computer Science, The Johns Hopkins University, Baltimore MD 21218 (yairamir, baruch, rsean@cs.jhu.edu).
Amnon Barak and Arie Keren are with the Institute of Computer Science, The Hebrew University of Jerusalem, 91904 Israel (amnon, ariek@cs.huji.ac.il).

[2] This research is supported, in part, by the Defense Advanced Research Projects Agency (DARPA) under grant F306029610293 to Johns Hopkins University.

In this paper, we present a new job assignment strategy based on "economic" principles and competitive analysis. This strategy enables us to manage heterogeneous resources in a near-optimal fashion. The key idea of this strategy is to convert the total usage of several heterogeneous resources, such as memory and CPU, into a single homogeneous "cost." Jobs are then assigned to the machine where they have the lowest cost.

This economic strategy provides a unified algorithmic framework for allocation of computation, communication, memory and I/O resources. It allows the development of near-optimal online algorithms for allocating and sharing these resources.

Our strategy guarantees near-optimal end-to-end performance for the overall system on each single instance of job generation and resource availability. This is accomplished using online algorithms that know nothing about the future, assume no correlation between past and future, and are only aware of the state. In spite of this, one can rigorously prove that their performance will always be comparable to that of the optimal prescient strategy.

This work shows that the unified opportunity cost approach offers good performance in practice. First, we performed tests using a simulated cluster and a "typical" series of incoming jobs. Our method, with and without reassignments, was compared against the methods of PVM, a dominant static job assignment strategy, and Mosix, one of the more successful system that support transparent process migration. Each method was given an identical stream of jobs. Over 3,000 executions of this Java-based simulation were performed, each representing at least 10,000 simulated seconds. When no reassignments were allowed, our method was shown to be a dramatic improvement over PVM. When reassignments *were* allowed, our method was substantially better than that of the highly tuned, but ad hoc, Mosix strategy.

A second series of tests was performed on a real system, to validate this simulation. The real system was built with BSD/OS machines with a collection of Pentium 133, Pentium Pro 200 and Pentium II machines with different memory capacity, connected by Fast Ethernet. The physical cluster and the simulated cluster were slightly different, but the proportional performance of the various strategies was very close to that given by the Java simulation. This indicates that the simulation appropriately reflects events on a real system.

In Section 2, we will discuss the model we used and our assumptions. In Sections 3 and 4, we will describe our algorithm and the theoretical guarantees that come with it. In Section 5, we will show our experimental evidence that this strategy is useful in practice. Section 6 concludes the paper. For additional information about this research, consult http://www.cnds.jhu.edu/projects/metacomputing.

2. The Model

The goal of this work is to improve performance in a cluster of n machines, where machine i has a CPU resource of speed $r_c(i)$ and a memory resource of size $r_m(i)$. We

will abstract out all other resources associated with a machine, although our framework can be extended to handle additional resources.

There is a sequence of arriving jobs that must be assigned to these machines. Each job is defined by three parameters:

- Its arrival time, *a(j)*,
- The number of CPU seconds it requires, *t(j)*, and
- The amount of memory it requires, *m(j)*.

We assume that *m(j)* is known when a job arrives, but *t(j)* is not. A job must be assigned to a machine immediately upon its arrival, and may or may not be able to move to another machine later.

Let *J(t,i)* be the set of jobs in machine *i* at time *t*. Then the CPU load and the memory load of machine *i* at time t are defined by:

$$l_c(t,i) = |J(t,i)| ,$$

and

$$l_m(t,i) = \sum_{j \varepsilon J(t,i)} m(j) \quad \text{respectively.}$$

We will assume that when a machine runs out of main memory, it is slowed down by a multiplicative factor of τ, due to disk paging. The *effective CPU load* of machine *i* at time *t*, *L(t,i)*, is therefore:

$$l_c(t,i) = |J(t,i)| \qquad \text{if } l_m(t,i) \le r_m(i),$$

and $$l_c(t,i) * \tau \qquad \text{if } l_m(t,i) > r_m(i).$$

For simplicity, we will also assume that all machines schedule jobs *fairly*. That is, at time *t*, each job on machine *i* will receive $1/L(t,i)$ of the CPU resource. A job's completion time, *c(j)*, therefore satisfies the following equation:

$$\int_{a(j)}^{c(j)} \frac{r_c(i)}{L(t,i)} = t(j) \text{, where } i \text{ is the machine the job is on at any given time.}$$

The *slowdown* of a job is equal to $\frac{c(j)-a(j)}{t(j)}$.

Our goal in this paper is to develop a method for job assignment and/or reassignment that will minimize the average slowdown over all jobs.

3. Theoretical Background

We will evaluate the effectiveness of our (online) algorithms by their *competitive ratio*, measured against the performance of an optimal offline algorithm. An online algorithm ALG is *c*-competitive if for any input sequence *I*, $ALG(I) \le c\ OPT(I) + \alpha$, where OPT is the optimal offline algorithm and α is a constant.

3.1 Introduction and Definitions

The theoretical part of this paper will focus on how to minimize the maximum usage of the various resources on a system – in other words, the best way to balance a system's load. One algorithm for doing so, described in [AAF96], proves useful in practice even when our goal is to minimize the average slowdown instead, which corresponds to minimizing the sum of the squares of the loads.

In preparation for a discussion of this algorithm, ASSIGN-U, we will examine this minimization problem with three different machine models and two different kinds of jobs. The three machine models are:

1. **Identical Machines.** All of the machines are identical, and the speed of a job on a given machine is determined only by the machine's load.
2. **Related Machines.** The machines are identical except that some of them have different speeds – in the model above, they have different r_c values, and the memory associated with these machines is ignored.
3. **Unrelated Machines.** Many different factors can influence the effective load of the machine and the completion times of jobs running there. These factors are *known*.

The two possible kinds of jobs are:

1. **Permanent Jobs.** Once a job is on a machine, it will remain there forever.
2. **Temporary Jobs.** Jobs leave the system when they have received a certain amount of CPU time.

We will also examine a related problem, called the *online routing* problem.

3.2 Identical and Related Machines

For now, we will assume that no reassignments are possible, and that the only resource is CPU time. Our goal, therefore, is to minimize the maximum CPU load.

When the machines are identical, and no other resources are relevant, the *greedy algorithm* performs well. This algorithm for job assignment assigns the next job to the machine with the minimum current CPU load. If the machines are identical, and no other resources are relevant, the greedy algorithm has a competitive ratio of 2 – 1/n (see [BFKV92]).

When the machines are related, the jobs are permanent, and no other resources are relevant, the *ASSIGN-R* algorithm by Aspnes *et al* [AAFPW93] is effective. Define OPT to be the load of the optimal offline algorithm; an approximation to OPT is given in [AAFPW93]. This algorithm assigns each arriving job to the *slowest* machine with a resulting load below 2 * OPT. If OPT is known, this algorithm has a competitive ratio of 2. A doubling technique can be used to approximate OPT. If this is necessary, the algorithm has a competitive ratio of 8.

For unrelated machines and temporary jobs, without job reassignment, there is no known algorithm with a competitive ratio better than n.

3.3 Unrelated Machines

ASSIGN-U is an algorithm for unrelated machines and permanent job assignments, based on an exponential function for the 'cost' of a machine with a given load. This algorithm then assigns each job to a machine to minimize the total cost of all of the machines in the cluster. More precisely, let:

- a be a constant, $1 < a < 2$,
- $l_i(j)$ be the load of machine i before assigning job j, and
- $p_i(j)$ be the load job j will add to machine i.

The online algorithm will assign j to the machine i that minimizes the marginal cost

$$H_i(j) = a^{l_i(j)+p_i(j)} - a^{l_i(j)}.$$

This algorithm is O(log n) competitive for unrelated machines and permanent jobs. The work presented in [AAPW94] extends this algorithm and competitive ratio to temporary jobs, using up to O(log n) *reassignments* per job. A reassignment moves a job from its previously assigned machine to a new machine. In the presence of reassignments, let

- $h_i(j)$ be the load of machine i just before j was last assigned to i.

When any job is terminated, the algorithm of [AAPW94] checks a 'stability condition' for each job j and each machine M. This stability condition, with i denoting the machine on which j currently resides, is:

$$a^{h_i(j)+p_i(j)} - a^{h_i(j)} \le 2 * (a^{l_M(j)+p_M(j)} - a^{l_M(j)})$$

If this stability condition is not satisfied by some job j, the algorithm reassigns j to machine M that minimizes $H_M(j)$.

3.4 Online routing of virtual circuits

The *ASSIGN-U* algorithm above minimizes the maximum usage of a single resource. In order to extend this algorithm to several resources, we examine the related *online routing of virtual circuits* problem. The reason this problem is applicable will be discussed shortly. In this problem, we are given:

- A graph G(V,E), with a capacity $u(e)$ on each edge e,
- A maximum load mx, and
- A sequence of independent requests (s_j, t_j, $p:E\rightarrow[0,mx]$) arriving at arbitrary times. s_j and t_j are the source and destination nodes, and $p(j)$ is the required bandwidth. A request that is assigned to some path P from a source to a destination increases the load l_e on each edge $e \in P$ by the amount $p_e(j) = p(j)/u(e)$.

Our goal is to minimize the maximum link congestion, which is the ratio between the bandwidth allocated on a link and its capacity.

Minimizing the maximum usage of CPU and memory, where memory usage is measured in the fraction of memory consumed, can be reduced to the online routine

problem. This reduction works as follows: create two nodes, {*s, t*} and *n* non-overlapping two-edge paths from *s* to *t*. Machine *I* is represented by on e of these paths, with a *memory* edge with capacity $r_m(i)$ and a *CPU* edge with capacity $r_c(i)$. Each job *j* is a request with *s* as the source, *t* as the sink, and *p* a function that maps memory edges to the memory requirements of the job and CPU edges to 1. The maximum link congestion is the larger of the maximum CPU load and the maximum memory (over)usage.

ASSIGN-U is extended in [AAFPW93] to address the online routing problem. The algorithm computes the marginal cost for each possible path *P* from s_j to t_j as follows:

$$H_P(j) = \sum a^{l_e + p_e(j)} - a^{l_e},$$

and assigns request *j* to a path *P* that yields a minimum marginal cost.

This algorithm is O(log n) competitive [AAFPW93]. By reduction, it produces an algorithm for managing heterogeneous resources that is O(log n) competitive in its maximum usage of each resource.

4. From Theory to Practice

For each machine in a cluster of *n* machines, with resources $r_1 \ldots r_k$, we define that machine's *cost* to be:

$$\sum_{i=1}^{k} f(n, utilization\ of\ r_i)$$

where *f* is some function. In practice, using ASSIGN-U, we will choose *f* so that this sum is equal to:

$$\sum_{i=1}^{k} n^{\frac{utilized\ r_i}{\max usage\ of\ r_i}}.$$

The *marginal cost* of assigning a job to a given machine is the amount by which this sum increases when the job is assigned there. An "opportunity cost" approach to resource allocation assigns jobs to machines in a way that minimizes this marginal cost. ASSIGN-U uses an opportunity cost approach.

In this paper, we are interested in only two resources, CPU and memory, and we will ignore other considerations. Hence, the above theory implies that given logarithmically more memory than an optimal offline algorithm, ASSIGN-U will achieve a maximum slowdown within O(log n) of the optimal algorithm's maximum slowdown.

This does not guarantee that an algorithm based on ASSIGN-U will be competitive in its average slowdown over all processes. It also does not guarantee that such an algorithm will improve over existing techniques. Our next step was to verify that such an algorithm does, in fact, improve over existing techniques in practice.

The memory resource easily translates into ASSIGN-U's resource model. The cost for a certain amount of memory usage on a machine is n^u, where *u* is the proportional memory utilization (used memory / total memory.) For the CPU resource, we must know the maximum possible load. Drawing on the theory, we will assume that *L*, the

smallest integer power of two greater than the largest load we have seen at any given time, is the maximum possible load. This assumption, while inaccurate, does not change the competitive ratio of ASSIGN-U.

The cost for a given machine's CPU and memory load, using our method, is:

$$n^{\frac{used\ memory}{total\ memory}} + n^{\frac{CPU\ load}{L}}.$$

In general, we will assign or reassign jobs so as to minimize the sum of the costs of all the machines in the cluster.

To examine the behavior of this "opportunity cost" approach, we evaluated four different methods for job assignment:

1. **PVM**. PVM (for "Parallel Virtual Machine") is a very popular metacomputing system for systems without preemptive process migration. Unless the user of the system specifically intervenes, PVM assigns jobs to machines using a strict Round-Robin strategy. It does not reassign jobs once they begin execution.
2. **Enhanced PVM**. Enhanced PVM is an opportunity cost-based strategy that assigns each job to the machine where the job has the smallest marginal cost. As with PVM, initial assignments are permanent.
3. **Mosix**. The Mosix kernel enhancements to BSD/OS allow the system to migrate processes from one machine to another without interrupting their work. Mosix uses an improved load-balancing strategy that also endeavors to keep some memory free on all machines. Mosix is not omniscient; when the system is exchanging process information in preparation for possible process reassignment, each machine is only in contact with a limited selection of other machines.
4. **Enhanced Mosix**. Enhanced Mosix is an opportunity cost-based strategy intended for use on systems (such as Mosix clusters) that can preemptively migrate processes. It assigns or reassigns jobs to minimize the sum of the costs of all of the machines. Enhanced Mosix has the same limits on its knowledge as unenhanced Mosix.

5. Experimental Results

Our first test of the ASSIGN-U algorithm was a Java simulation of the four job (re)assignment methods above. Our assumptions were as follows:

1. The cluster contains six machines, with the following properties:

Machine Type	# of these Machines	Processing Speed	Installed Memory
Pentium Pro	3	200 MHz.	64 MB of RAM
Pentium	2	133 MHz.	32 MB of RAM
Laptop w/ Ethernet	1	90 MHz.	24 MB of RAM.

Table 1: Machines in the simulated cluster.

If a machine's type is not specified, in the remainder of this paper, it can be assumed to be one of the Pentium Pros. This cluster corresponds to a real-world cluster of

machines at the Center for Networking and Distributed Systems at the Johns Hopkins University.

2. For each incoming job, let r and m be independently-generated random numbers between 0 and 1. A typical process will require $2/r$ seconds of CPU time and $(1/m)\%$ of a Pentium Pro's memory. (The distribution is based on the observations of real-life processes described in [HD96].) Approximately 95% of all jobs are single-process jobs matching this profile; because this is a metacomputing system, 5% of all jobs are assumed to be large parallel jobs utilizing the metacomputer's power. These jobs contain between 1 and 20 identical processes requiring $\mathbf{20}/r$ seconds of CPU time and $(1/m)\%$ of a Pentium Pro's memory. To make the simulation finite, it is assumed that no process requires more than 100% of a Pentium Pro's memory, no process from a single-process job requires more than 1,000 seconds of CPU time, and no process from a large parallel job requires more than 10,000 seconds of CPU time.
3. Jobs arrive at random times during the first 1,000 simulated seconds. Approximately one job arrives every 10 seconds. We observed that for every one of the four methods, there were simulation instances where the system was overloaded, underloaded, and normally loaded, and instances where the system transitioned from any one of these states to any other.
4. When a machine's memory usage is greater than its memory capacity, that machine is assumed to be thrashing. This slows the machine down by a factor τ, which we approximated with a constant factor equal to 10. That is, in the simulation, every job required 10 times as long on a thrashing machine as it would require on a machine with free memory.

In each execution of the simulation, all four methods were provided with an identical scenario, where the same jobs arrived at the same rate.

5.1 Java Simulation Results

Each execution returned the average slowdown over all jobs in that execution, as well as some information about the scenario itself. These results were evaluated in two different ways:

- An important concern is the overall slowdown experienced using each of the four methods. The *average slowdown by execution* is an unweighted average of all of the simulation results, regardless of the number of jobs in each execution. The *average slowdown by job* is the average slowdown over all of the jobs in all of the executions of the simulation. These results, incorporating 3000 executions, are given in Table 2.
- The behavior of Enhanced PVM and Enhanced Mosix is different in lightly-loaded and heavily-loaded scenarios. This behavior is illustrated in Figures 1 to 4, detailing the first 1000 executions of the simulation.

Slowdown for …	PVM	Enhanced PVM	Mosix	Enhanced Mosix
(average by execution)	14.3338	9.79463	8.55676	7.47886
(average by job)	15.4044	10.7007	9.4208	8.20262

Table 2: Average slowdown in the Java simulation for the different methods.

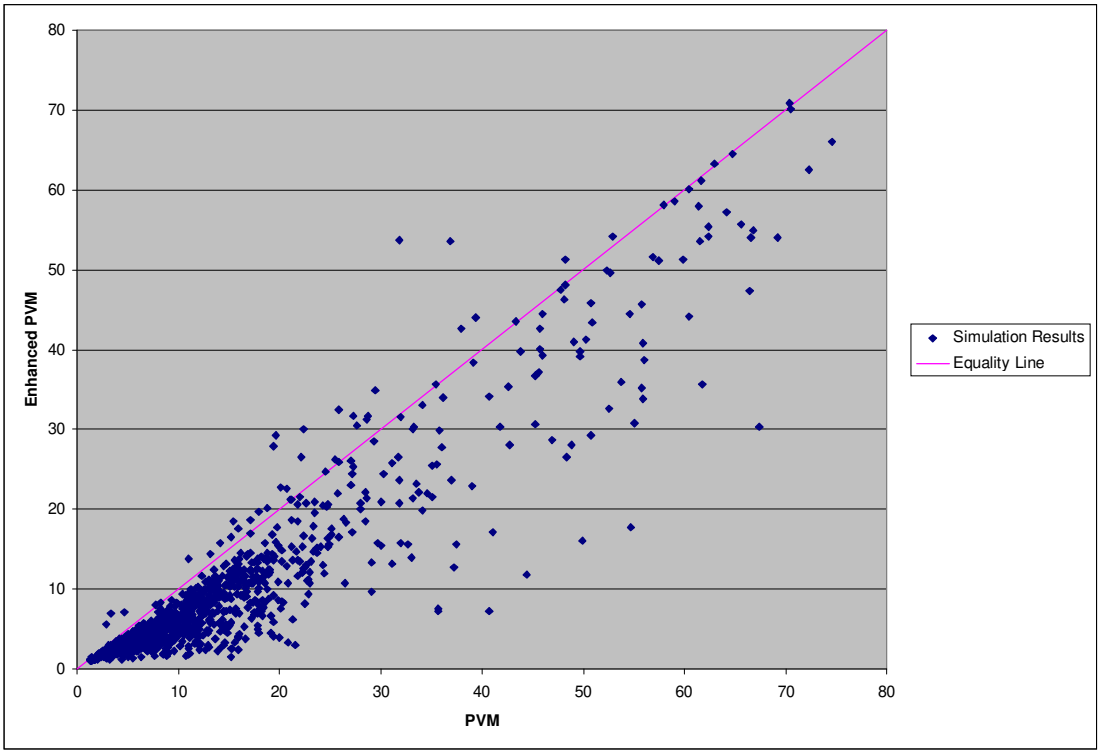

Figure 1: PVM vs. Enhanced PVM

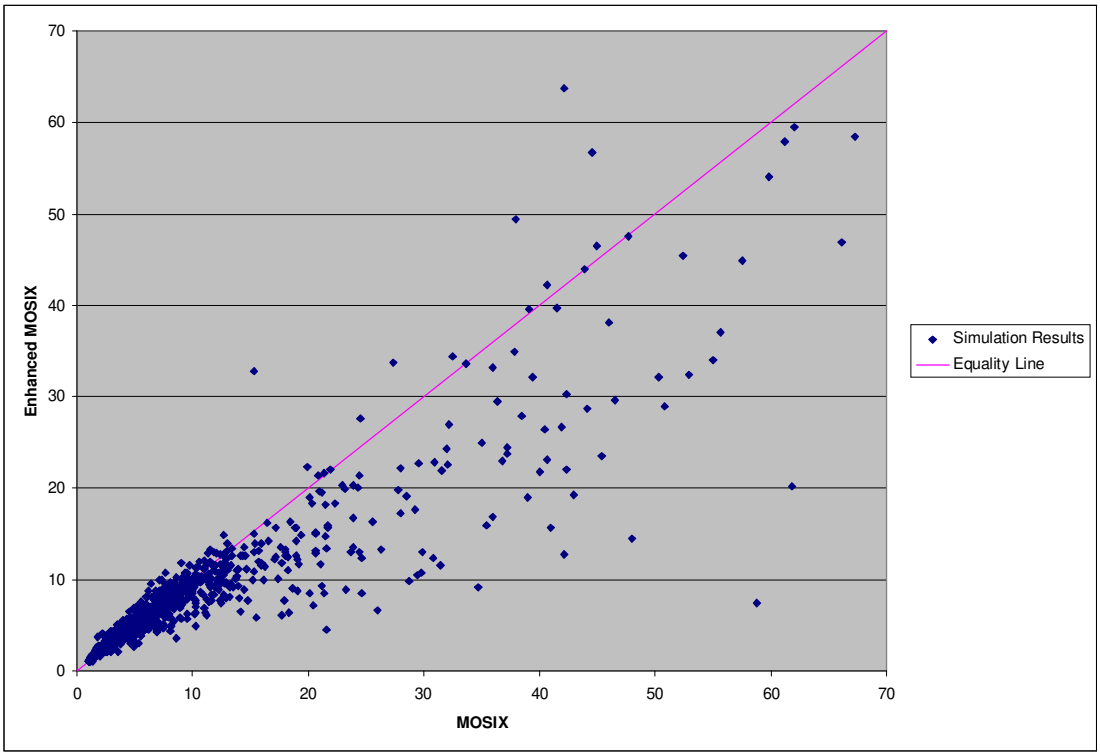

Figure 2: Mosix vs. Enhanced Mosix

Each point in the figures represents a single execution of the simulation for the two methods named. In Figure 1, the X axis is the average slowdown for PVM, and the Y axis is the average slowdown for enhanced PVM. Similarly, in Figure 2, the X axis is the average slowdown for Mosix, and the Y axis is the average slowdown for enhanced Mosix. The light line is defined by 'x = y'. Above this line, the un-enhanced algorithm does better than the enhanced algorithm. Below this line, the enhanced algorithm does better than the un-enhanced algorithm.

Enhanced PVM, as Table 2 has already shown, does significantly better than straight PVM in almost every circumstance. More interesting, however, is the behavior of enhanced Mosix when compared to Mosix. The larger Mosix's average slowdown was on a given execution, the more improvement our enhancement gave. Intuitively, when an execution was hard for all four models, Enhanced Mosix did much better than unenhanced Mosix. If a given execution was relatively easy, and the system was not heavily loaded, the enhancement had less of a positive effect.

This can be explained as follows. When a machine becomes heavily loaded or starts thrashing, it does not just affect the completion time for jobs already submitted to the system. If the machine does not become unloaded before the next set of large jobs is submitted to the system, it is effectively unavailable to them, increasing the load on all other machines. If many machines start thrashing or become heavily loaded, this effect will build on itself. Every incoming job will take up system resources for a much longer span of time, increasing the slowdown experienced by jobs that arrive while it computes. Because of this pyramid effect, a 'wise' initial assignment of jobs and careful re-balancing can result (in the extreme cases) in a significant improvement over standard Mosix, as shown in some of the executions in Figure 2.

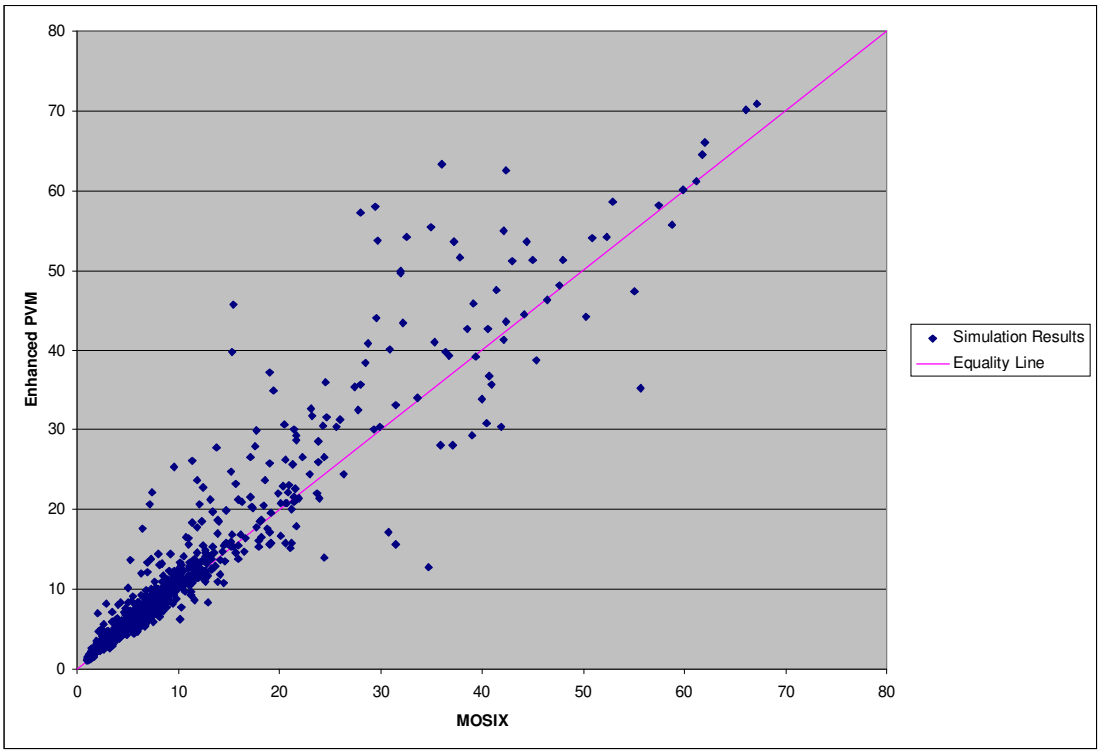

Figure 3: Mosix vs. Enhanced PVM.

It is particularly interesting to note that, as seen in Table 2 and Figure 3, the enhanced PVM method, which makes no reassignments at all, manages to achieve respectable (though inferior) performance compared to Mosix. This emphasizes the power of the opportunity cost approach: its performance on a normal system is not overwhelmed by the performance of a much superior system that can correct initial assignment mistakes.

The importance of migration is demonstrated by Figure 4. Even when using the opportunity cost algorithm, it is still very useful to have the migration ability in the system. In fact, Enhanced Mosix outperform Enhanced PVM in all of the cases, sometimes considerably.

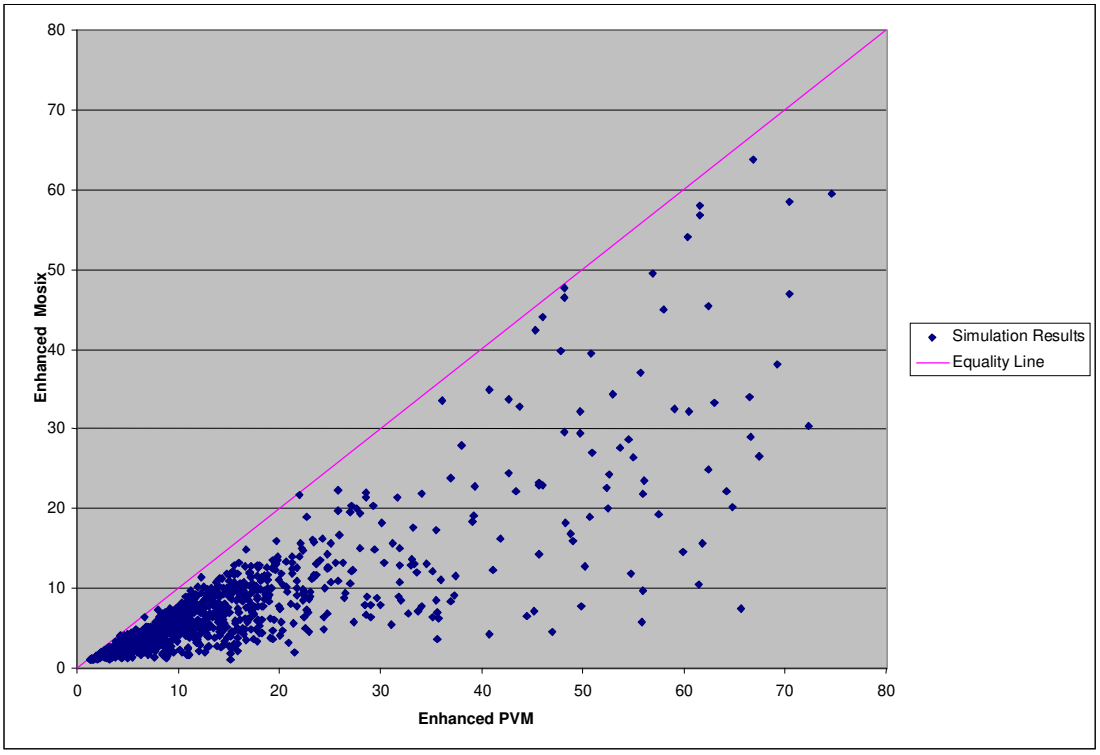

Figure 4: Enhanced PVM vs. Enhanced Mosix

5.2 Real System Executions

Our algorithms were also tested on a real cluster. The same model for incoming jobs was used, and jobs were assigned using the PVM, Enhanced PVM, and Mosix strategies. Enhanced Mosix has not yet been implemented on a real system. The results are as follows:

Slowdown for …	PVM	Enhanced PVM	Mosix
(average by execution)	29.98788	16.29643	13.67707
(average by job)	33.31620	16.76646	14.00990

Table 3: Average slowdown in the real cluster for 3 (re)assignment methods.

These initial results imply that the real-life thrashing constant, the smaller cluster, and various miscellaneous factors increased the average slowdown, which indicates that we were too conservative in picking the parameters for the simulation. Nevertheless, the

results do not substantially change the relative values. In fact, the Enhanced methods proved to be even better in real-life, compared to their original versions, than what our Java simulation has predicted. We consider that as a validation of our initial Java simulations and as a strong confirmation for the merit of our opportunity cost approach.

Slowdown on Real System for …	**PVM / Mosix**	**Enh. PVM / Mosix**	**PVM / Enh. PVM**
(average by execution)	2.19257	1.20416	1.84015
(average by job)	2.37804	1.20946	1.98707
Slowdown in Simulation for …	**PVM / Mosix**	**Enh. PVM / Mosix**	**PVM / Enh. PVM**
(average by execution)	1.67514	1.14467	1.46346
(average by job)	1.63515	1.13585	1.43957

Table 4: Average *relative* slowdowns for 3 job (re)assignment methods.

6. Conclusions

The opportunity cost approach is a universal framework for efficient allocation of heterogeneous resources. The theoretical guarantees are weak: one can only prove a logarithmic bound on the gap between the algorithm and the optimum offline schedule. However, the optimum offline schedule is not really an option; in reality, our algorithm competes with naive online heuristics.

In practice, this approach yields simple algorithms that significantly outperform widely used and carefully optimized methods. We conclude that the theoretical guarantees of logarithmic optimality is a good indication that the algorithm will work well in practice.

References

[AAF96] B. Awerbuch, Y. Azar and A. Fiat. Packet Routing via Min-Cost Circuit Routing. In *Proceedings of the Israeli Symposium on Theory of Computing and Systems,* 1996.

[AAFPW93] J. Aspnes, Y. Azar, A. Fiat, S. Plotkin and O. Waarts. On-Line Machine Scheduling with Applications to Load Balancing and Virtual Circuit Routing. In *Proceedings of the ACM Symposium on Theory Of Computing (STOC)*, May 1993.

[AAPW94] B. Awerbuch, Y. Azar, S. Plotkin and O. Waarts. Competitive Routing of Virtual Circuits with Unknown Duration. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, January 1994.

[BFKV92] Y. Bartal, A. Fiat, H. Karloff and R. Vohra. New algorithms for an ancient scheduling problem. In *Proceedings of the ACM Symposium on Theory of Algorithms*, 1992.

[BGW93] A. Barak, S. Guday and R. Wheeler. The Mosix distributed operating system, load balancing for Unix, Volume 672, May 1993.

[BL97] A. Barak and O. La'adan, The MOSIX Multicomputer Operating System for High Performance Cluster Computing, Journal of Future Generation Computer Systems, April 1998, to appear.

[Bsdi] Berkeley Software Design, Inc. http://www.bsdi.com.

[HD96] M. Harchol-Balter and A. Downey. Exploiting Process Lifetime Distributions for Dynamic Load Balancing. In *Proceedings of ACM Sigmetrics Conference on Measurement and Modeling of Computer Systems*, May 1996.