

MIMP: Deadline and Interference Aware Scheduling of Hadoop Virtual Machines

Wei Zhang^{*†}, Sundaresan Rajasekaran^{*}, Timothy Wood^{*}, Mingfa Zhu[†]

^{*} *The George Washington University, Washington, D.C., USA*

[†] *Beihang University, Beijing, China*

Abstract—Virtualization promised to dramatically increase server utilization levels, yet many data centers are still only lightly loaded. In some ways, big data applications are an ideal fit for using this residual capacity to perform meaningful work, but the high level of interference between interactive and batch processing workloads currently prevents this from being a practical solution in virtualized environments. Further, the variable nature of spare capacity may make it difficult to meet big data application deadlines.

In this work we propose two schedulers: one in the virtualization layer designed to minimize interference on high priority interactive services, and one in the Hadoop framework that helps batch processing jobs meet their own performance deadlines. Our approach uses performance models to match Hadoop tasks to the servers that will benefit them the most, and deadline-aware scheduling to effectively order incoming jobs. The combination of these schedulers allows data center administrators to safely mix resource intensive Hadoop jobs with latency sensitive web applications, and still achieve predictable performance for both. We have implemented our system using Xen and Hadoop, and our evaluation shows that our schedulers allow a mixed cluster to reduce web response times by more than ten fold, while meeting more Hadoop deadlines and lowering total task execution times by 6.5%.

Keywords—scheduling; virtualization; Map Reduce; interference; deadlines

I. INTRODUCTION

Virtualization has facilitated the growth of infrastructure cloud services by allowing a single server to be shared by multiple customers. Dividing a server into multiple virtual machines (VMs) provides both a convenient management abstraction and resource boundaries between users. However, the performance isolation provided by virtualization software is not perfect, and interference between guest VMs remains a challenge. If the hypervisor does not enforce proper priorities among guests, it is easy for one virtual machine’s performance to suffer due to another guest.

Despite the danger of interference, resource sharing through virtualization has been crucial for lowering the cost of cloud computing services. Multiplexing servers allows for higher average utilization of each machine, giving more profit for a given level of hardware expense. Yet the reality is that many data centers, even those employing virtualization, are still unable to fully utilize each server. This is due in part to fears that if a data center is kept fully utilized there will be no spare capacity if workloads rise, and part due to the risk of VM interference hurting performance even if servers are left underloaded.

In this paper, we first study the causes of interference through virtualization scheduler profiling. We observe that even when set to the lowest possible priority, big data VMs (e.g., Hadoop jobs) interrupt interactive VMs (e.g., web servers), increasing their time spent in the runnable queue, which hurts response times. We control and reduce VM CPU interference by introducing a new scheduling priority for “background” batch processing VMs, allowing them to run only when other VMs are not actively utilizing the CPU.

Our changes in the VM scheduler improve the performance of interactive VMs, but at the cost of unpredictable Hadoop performance. To resolve this challenge, we implement a second scheduler within the Hadoop framework designed for hybrid clusters of dedicated and shared VMs that only use residual resources. We find that when given the same available resources, different tasks will progress at different rates, motivating the need to intelligently match each Hadoop task to the appropriate dedicated or shared server. Our scheduler combines performance models that predict task affinity with knowledge of job deadlines to allow Hadoop to meet SLAs, despite variability in the amount of available resources.

Together, these schedulers form the Minimal Interference Maximal Productivity (MIMP) system, which enhances both the hypervisor’s scheduler and the Hadoop job scheduler to better manage their performance. Our primary contributions include:

- A study of the workloads on 192 production servers that illustrate the vast untapped potential common in today’s data centers.
- A new priority level built into Xen’s Credit Scheduler that prevents batch processing VMs from hurting interactive VM performance.
- Task affinity models that match each Hadoop task to the dedicated or shared VM that will provide it the most benefit.
- A deadline and progress aware Hadoop job scheduler that allocates resources to jobs in order to meet performance goals and maximize the efficiency of a hybrid cluster.

We have implemented the proposed schedulers by modifying the Xen hypervisor and Hadoop scheduler. Our evaluation shows that MIMP can prevent nearly all interference on a web application, doubling its maximum throughput and

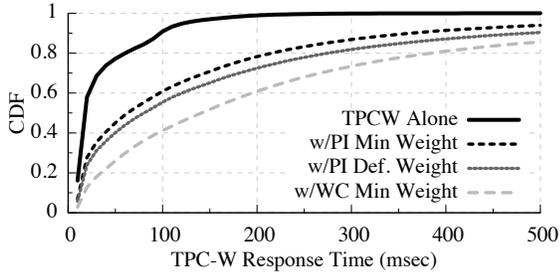


Figure 1. Colocated Hadoop jobs significantly hurt web application performance, even when Xen’s scheduler priority mechanisms are used.

providing nearly identical response times to when it is run alone. For a set of batch jobs, the algorithm can meet more deadlines than EDF, and reduces the total execution time by over four and a half CPU hours, all with minimal impact on interactive VM performance.

II. BACKGROUND AND MOTIVATION

A. Map Reduce in Virtualized Clusters

Map Reduce is a popular framework for distributing data intensive computation [6]. Hadoop is an open source implementation of Map Reduce developed by Yahoo. Users write a program that divides the work that must be performed into two main phases: Map and Reduce. The Map phase processes each piece of input data and generates some kind of intermediate value, which is in turn aggregated by the Reduce phase.

Hybrid Clusters: In this paper we investigate how to run Map Reduce jobs in a hybrid cluster consisting of both dedicated and shared (also known as volunteer) nodes. This problem was first tackled by Clay et al., who described how to pick the appropriate number of shared nodes in order to maximize performance and minimize overall energy costs [5]. Like their work, we focus on scheduling and modeling the Map phase since this is generally the larger portion of the program, and is less prone to performance problems due to slow nodes. Our work extends their ideas both within the virtualization layer to prevent interference, and at the Map Reduce job scheduling level to ensure that *multiple* jobs can make the best use of a hybrid cluster and effectively meet deadlines.

A key issue that has not yet been fully explored is how to prevent batch processing jobs such as Map Reduce from interfering with foreground workloads. Our results suggest that interference can be quite severe if the important performance metric is interactive latency as opposed to coarse grained timing measures (e.g., the time to compile a linux kernel).

As a motivating experiment, we have measured the achieved throughput and response time when running the TPC-W online book store benchmark both alone and alongside a VM running Hadoop jobs. Our results in Figure 1 show that the response time of the web application can be

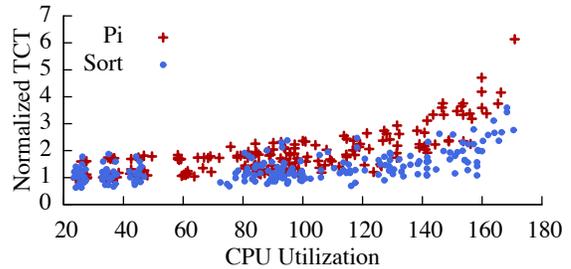


Figure 2. Task completion time varies by job, and increases non linearly as the web service consumes a larger quantity of CPU (out of 2 cores).

dramatically increased when run with a Pi or WordCount (WC) job. This happens even when the Xen scheduler’s parameters are tuned to give Hadoop the lowest possible weight (i.e., the lowest priority). However, the throughput of TPC-W remains similar, as does the amount of CPU that it consumes. Further, we find that if Hadoop is given a separate CPU from TPCW, there is no interference at all. This suggests that the performance interference is due to poor CPU scheduling decisions, not IO interference.

A second major challenge when running in shared environments is that different Hadoop jobs are affected by limitations on available resources in different ways. Figure 2 shows that as the amount of resources consumed by a foreground interactive VM rises, the normalized task completion time (relative to Hadoop running alone) can increase significantly for some jobs. For example, Pi, a very CPU intensive job, suffers more than Sort, which is IO intensive. As a result, the best performance will be achieved by carefully matching a Hadoop job to the servers that will allow it to make the most efficient progress.

B. Data Center Survey: Residual Capacity

Both private and public data centers are known to have relatively low utilization rates, yet servers are the largest cost in the data center [8], so it is more efficient to keep them fully utilized than shut down. We have studied a cluster of 55 hosts within our university and 137 hosts that run the Wikimedia Foundation’s websites, including wikipedia, to see how much capacity is left idle.

IT Department Virtualized Cluster: We first examine a cluster of servers used by our university’s Division of IT. The cluster uses VMWare’s virtualization platform to run approximately 10 virtual machines on each of the 55 hosts. The VMs run a range of applications such as web servers, domain controllers, databases, and business applications. We have gathered traces of the CPU, memory, network, and disk utilization for each host over a five month period. For each resource we report % utilization out of the total host capacity. We report the percent of memory on each host that VMware considers actively used [20]. All hosts have two Xen E5-2640 CPUs (12 cores total), 96GB of RAM, and are connected to ten gigabit ethernet.

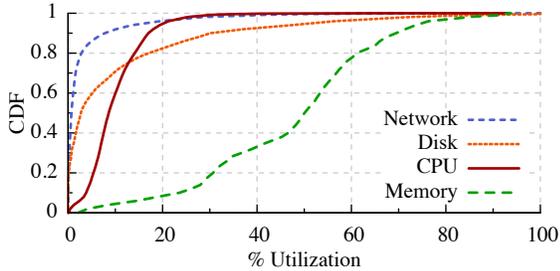


Figure 3. Servers used to host VMs for our IT department illustrate that many resources are often left idle; memory is the only resource with a median utilization level above 50%.

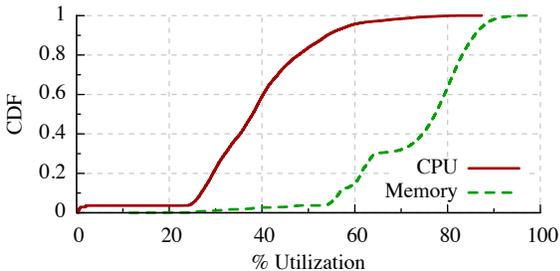


Figure 4. Wikipedia's application servers are moderately loaded for CPU (median utilization of 38%) and highly loaded for memory (median utilization of 77%).

Figure 3 shows a CDF of the utilization of CPU, Memory, Network, and Disk. This data confirms that CPU is lightly utilized, but also shows that network and disk see relatively low loads as well. In contrast, memory has a higher utilization, but still has only about half of the servers using 50% or more of their total RAM.

Wikipedia Application Servers: We next analyze a trace from Wikipedia's Application Servers running in its Equiad data center starting on October 8th and ending on November 3rd, 2013.¹ This cluster is composed of 137 servers, each with either 6 or 12 processing cores and either 12 or 60GB of RAM. While this cluster does not use virtualization, it still gives us insight into the workload seen by a large-scale web application, and illustrates the potential for better resource multiplexing if virtualized servers could be used.

Figure 4 shows a CDF of resource usage on these servers over the full measurement period. For this data set we report only the CPU usage and memory consumption (i.e., the sum of active memory and cached memory as reported by the OS) since disk and network capacities are unavailable. Note that these servers are not virtualized, so it is possible that additional memory could be reclaimed through page sharing or idle memory taxes [20]. This cluster is clearly more highly utilized than our university data center, but it still shows a significant amount of residual capacity—half of the CPU measurements fall below 38%, and half of the 12GB servers

¹We use public data from <http://ganglia.wikimedia.org/>

have at least 2.76GB of memory free.

The relatively low utilizations found in both these data centers is likely due in part to conservative overprovisioning and partly to meet future needs. However, these results suggest that there is a great deal of capacity of *all resource types* that is being wasted in many data centers, even virtualized ones.

III. PROBLEM AND SYSTEM OVERVIEW

This section presents the formal problem MIMP targets, and then gives an overview of the system.

A. Problem Statement

The scenario where we believe MIMP will provide the most benefit is in a hybrid cluster containing a mix of dedicated nodes (virtual or physical) and “volunteer” or “shared” nodes that use virtualization to run both interactive applications and Hadoop tasks. We assume that the interactive applications are *higher priority* than the Hadoop tasks, which is generally the case since users are directly impacted by slowdown of interactive services, but may be willing to wait for long running batch processes. While we focus on web applications, the interactive applications could represent any latency-sensitive service such as a streaming video server or remote desktop application. Although we treat Hadoop jobs as *lower priority*, we still take into account their performance by assuming they arrive with a deadline by which time they must be complete.

As discussed previously, we focus on the Map phase of Map Reduce, as this is generally more parallelizable and is less prone to straggler performance problems (i.e., a single slow reduce task can substantially hurt the total completion time). As in [5], we use dedicated servers to run both the shared Hadoop file system and all reduce tasks.

We assume that the interactive applications running in the high priority VMs have relatively low disk workloads, meaning that sharing the IO path with Hadoop tasks does not cause a resource bottleneck. While this is not true for some disk intensive applications such as databases, for others it can be acceptable, particularly due to the increasing use of networked storage (e.g., Amazon's Elastic Block Store) rather than local disks.

Given this type of cluster, a key question is how best to allocate the available capacity in order to maximize Hadoop job performance (i.e., minimize the number of deadline misses and the total job completion times) while minimizing the interference on the interactive services (i.e., minimizing the change in response time compared to running the web VMs alone).

B. MIMP Overview

We have developed MIMP to tackle this pair of challenges. The system is composed of two scheduling components, as illustrated in Figure 5.

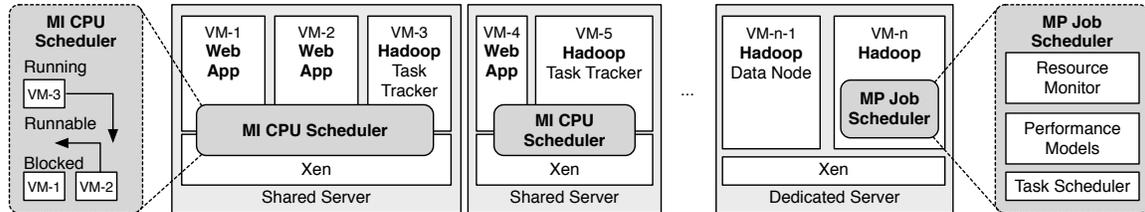


Figure 5. The MI CPU scheduler only runs Hadoop VMs when others are blocked, so VM-2 immediately preempts VM-3 once it becomes ready. The MP Job Scheduler gathers resource availability information from all nodes and schedules jobs based on performance model results.

Minimal Interference CPU Scheduler: The MI CPU Scheduler tries to prevent lower priority virtual machines from taking CPU time away from interactive VMs. We do this by modifying the Xen CPU scheduler to define a new priority level that will always be preempted if an interactive VM becomes runnable.

Maximal Productivity Job Scheduler: Next we modify the Hadoop Job scheduler to be aware of how available resources affects task completion time. The MP Scheduling system is composed of a training module that builds performance models, a monitoring system that measures residual capacity throughout the data center, and a scheduling algorithm. Our MP Scheduler combines this information to decide which available resources to assign to each incoming Hadoop Job to ensure it meets its deadline while making the most productive use of all available capacity.

IV. INTERACTIVE VM SCHEDULING IN XEN

This section first diagnoses the performance issues of Xen’s current Credit scheduler when mixing latency sensitive and computationally intensive virtual machines. We then describe how we have enhanced Xen’s scheduler to help minimize this interference.

A. Interactive Performance with Xen’s Credit Scheduler

Xen’s Credit scheduler is a non-preemptive weighted fair-share scheduler. As a VM runs, its VCPUs are dynamically assigned one of three priorities - over, under, or boost, ordered from lowest to highest. Each physical CPU has a local run queue for runnable VCPUs, and VMs are selected by their priority class. Every 30ms, a system-wide accounting thread updates the credits for each VCPU according to its weight share and resorts the queue if needed. If the credits for a VCPU are negative, Xen assigns “over” priority to this VCPU since it has consumed more than its share. If the credits are positive, it is assigned “under” priority. Every 10ms, Xen updates the currently running VCPU’s credits based on the VCPU running time. In order to improve a virtual machine’s I/O performance, if a VCPU is woken up (e.g., because an IO request completes) and it has credits left, it will be given “boost” priority and immediately scheduled. After the boosted VCPU consumes a non-negligible amount of CPU resources, then Xen resets the priority to “under”.

TPC-W 600 clients	Alone	+WC (min weight)
Avg. Resp. Time	25 msec	175.5 msec
Avg. CPU Utilization	84.8%	91.1%
Running (sec)	605.9	656.4
Runnable (sec)	1.9	524.4
Blocked (sec)	1092.4	520.2

Table I
XEN CREDIT SCHEDULER STATISTICS WHEN RUNNING A WEB APPLICATION ALONE OR WITH A WORD COUNT VM.

As this is a weight-based scheduler, it primarily focuses on allocating coarse grained shares of CPU to each virtual machine. The Boost mechanism is relied upon to improve performance of interactive applications, but as shown previously, it has limited effect.

Table I shows how much time was spent in each scheduler state when a TPCW VM is run either alone or with a VM running the Word Count Hadoop job that has been given the lowest possible scheduler weight. As was shown in Figure 1, this significantly affects TPCW performance, raising average response time by seven times. We find that the Credit Scheduler’s weight system does do a good job of ensuring that TPCW gets the overall CPU time that it needs—the CPU utilization (out of 200% since it is a 2-core machine) and time spent in the Running state are similar whether TPC-W is run alone or with word count. In fact, TPC-W actually gets more CPU time when run with word count, although the performance is substantially worse. While the overall CPU share is similar, the timeliness with which TPC-W is given the CPU becomes very poor when word count is also running. The time spent in the Runnable state (i.e., TPC-W could be servicing requests) rises substantially, causing the delays that increase response time.

This happens because Credit uses coarse grain time accounting, which means that 1) at times TPC-W may be woken up to handle IO, but it is not able to interrupt Hadoop; and 2) at times Hadoop may obtain boost priority and interrupt TPC-W if it is at the beginning of an accounting interval and has not yet used up its quantum.

B. Minimal Interference CPU Scheduler

Our goal is to run processor or data intensive virtual machines in the background, without affecting the more important interactive services. Therefore, we have modified Xen’s scheduler to define a new extra low priority class.

Virtual machines of this class are always placed at the end of the Runnable queue, after any higher priority VMs. We also adjust the Boost priority mechanism so that “background” VMs can never be boosted, and so that if a regular VM is woken up due to an I/O interrupt, it will always be able to preempt a background VM, regardless of its current priority (i.e., under or over).

This scheduling algorithm minimizes the potential CPU interference between interactive and Hadoop virtual machines, but it can cause starvation for background VMs. To prevent this, we allow a period, p , and execution time, e , to be specified. If over p seconds the VM has not been in the Running state for e milliseconds, then its priority is raised from background to over. After it is scheduled for the specified time slice, it reverts back to background mode. We use this to ensure that Hadoop VMs do not become completely inaccessible via SSH, and so they can continue to send heartbeat messages to the Hadoop job scheduler. While this mechanism is not necessary when running interactive VMs that typically leave the CPU idle part of the time, it can be important if MIMP is run either with CPU intensive foreground tasks, or with a very large number of interactive VMs.

V. PROGRESS AWARE DEADLINE SCHEDULING IN HADOOP

A Hadoop Job is broken down into multiple tasks, which each perform processing on a small part of the total data set. When run on dedicated servers, the total job completion time can be reliably predicted based on the input data size and previously trained models [26], [18]. The challenge in MIMP is to understand how job completion times will change when map tasks are run on servers with variable amounts of spare capacity. Using this information, MIMP then instructs the Hadoop Job Tracker on how to allocate “slots” (i.e., available shared or dedicated workers) to each job.

Monitoring Cluster Resource Availability: MIMP monitors resource usage information on each node to help guide task placement and prevent overload. MIMP runs a monitoring agent on each dedicated and shared node, and sends periodic resource measurements to the centralized MP Job Scheduler component. MIMP tracks the CPU utilization and disk read and write rates of each virtual machine on each host. These resource measurements are then passed on to the modeling and task scheduling components as described in the following sections.

A. Modeling Background Hadoop Jobs

MIMP uses task completion time models to predict the *progress rate* of different job types on a shared node with a given level of resources. As shown previously in Figure 2, each job needs its own task completion time model. The model is trained by running map tasks on shared nodes with different available CPU capacities. This can either be done

offline in advance, or the first set of tasks for a new job can be distributed to different nodes for measurement, and then a model can be generated and updated as tasks complete. Our current implementation assumes that all jobs have been trained in advance on nodes with a range of utilization levels. Once a job has been trained for one data input size, it can generally be easily scaled to accurately predict other data sizes [18].

Job Progress: The progress model for a job of type j is a function that predicts the *task completion time* on a shared node with residual capacity r . From Figure 2 we see that this relationship is highly non-linear, so we use a double exponential formula, *exp2*, provided by MATLABs Non-linear Least Squares functionality :

$$TCT_j(r) = a * e^{b*r} + c * e^{d*r} \quad (1)$$

where a , b , c , and d are the coefficients of the regression model trained for each job. The coefficients b and d represent the rate at which $TCT_j(r)$ exponentially grows. In order to compare the progress that will be made by a job on an available slot, we use the normalized TCT:

$$NormTCT_j(r) = \frac{TCT_j(r)}{TCT_j(r_{dedicated})} \quad (2)$$

where the denominator represents the task completion time when running on a dedicated node. This allows MIMP to compare the relative speeds of different jobs.

Checking Deadlines: The task completion time model can then be used to determine if a job will be able to meet its deadline given its current slot allocation. MIMP tracks a resource vector, R , for each active job. The entry R_i represent the amount of resources available on worker slot i that this job has been allocated for use: 100% for an available dedicated slot, 0% for a slot assigned to a different job, or something in between for a shared slot allocated to this job. If there is currently $t_{remaining}$ seconds until the job’s deadline, then MIMP can check if it will meet its deadline using:

$$CompletableTasks(j, R) = \sum_{slot\ i=1}^n \frac{t_{remaining}}{TCT_j(R_i)} \quad (3)$$

If $CompletableTasks(R)$ is greater than n_{tasks} , the number of remaining tasks for the job, then it is on track to meet its deadline.

Map Phase Completion Time: We can also obtain a direct prediction of the map phase completion time using:

$$CompletionTime(j, R) = n_{tasks} \sum_{slot\ i=1}^n \frac{TCT_j(R_i)}{n} \quad (4)$$

which estimates the total map phase completion time based on the average TCT of each slot and the number of remaining tasks.

Data Node I/O: Hybrid clusters like the ones considered in MIMP are particularly prone to disk I/O bottlenecks since there may be a relatively small number of dedicated nodes acting as the data store. If too many I/O intensive tasks are run simultaneously, task completion times may begin to rise [5]. To prevent this, we use a model to predict the I/O load incurred by starting a new map task. During MIMP’s model training phase, we measure the read request rate sent to the data nodes by a dedicated worker. Since I/O accesses can be erratic during map tasks, we use the 90th percentile of the measured read rates to represent the I/O required by a single worker, per data node available. In order to calculate the read I/O load incurred by a new task on a shared worker, we use the normalized TCT from Equation 2 as a scaling factor:

$$IO_j(r) = \frac{read_j^{90^{th}}}{NormTCT_j(r)} \quad (5)$$

to predict its I/O requirement. This can then be used to determine whether running the task will cause the data nodes to become overloaded, as described in the following section.

B. Progress Aware Earliest Deadline First

We now present two standard Hadoop job schedulers, and then discuss how we enhance these in MIMP so that it accounts for both deadlines and the relative benefit of assigning a worker to each job.

FIFO Scheduler: The simplest approach to scheduling Hadoop jobs is to service them in the order they arrive—all tasks for the first job are run until it finishes, then all tasks of the second job, and so on. Not surprisingly, this can lead to many missed deadlines since it has no concept of more or less urgent tasks to perform.

EDF Scheduler: Earliest Deadline First (EDF) is a well known scheduling algorithm that always picks the job with the earliest deadline when a worker becomes available; that job will continue to utilize all workers until it finishes or a new job arrives with a smaller deadline. EDF is known to be optimal in terms of preventing deadline misses as long as the system is preemptible and the cluster is not over-utilized. In practice, Hadoop has somewhat coarse grained preemption—each task runs to completion, but a job can be preempted between tasks. It is also difficult to predict whether a Hadoop cluster is over-utilized since tasks do not arrive on strict schedules as is typically assumed in Real Time Operating Systems work. Despite this, we still expect EDF to perform well when scheduling jobs since it will organize them to ensure they will not miss deadlines.

MP Scheduler: Our Maximum Progress (MP) Scheduler uses the models described in the previous section to enhance the EDF scheduler. Whenever a worker slot becomes free, we determine which job to assign to it based on the following criteria.

- First, MP examines each job in the queue to determine whether it can meet its deadline *with its currently allocated set of slots* using Equation 3. If one or more jobs are predicted to miss their deadline, then MP allocates the slot to whichever of those jobs has the closest deadline and returns.
- If all jobs are currently able to meet their deadlines, MP considers each job in the queue and uses Equation 2 to calculate its normalized task completion time if assigned the resources of the free slot. It finds the job with the smallest *normTCT* value, since that job is best matched for the available resources.
- Before assigning the slot, MP calculates the IO cost of running the selected job using Equation 5. If starting a new task of this type will cause any of the data nodes to become overloaded, then the job with the next highest *normTCT* is considered, and so on.

This algorithm ensures that the selected job is either currently unable to meet its deadline, or is the job that will make the most progress with the slot, without causing the data nodes to become overloaded.

VI. EVALUATION

A. Setup - machines, benchmarks

For our tests we use Xen 4.2.1 with Linux 3.7, running on a heterogenous cluster of Dell servers with Intel E5-2420 and Xeon X3450 CPUs with each having 16GB of RAM. The E5-2420 has six physical cores at 1.90GHz with 64KB of L1 cache and 256KB of L2 cache per core, and a shared 15MB L3 cache, and X3450 has four physical cores at 2.67GHz with 128KB of L1 cache and 1MB of L2 cache per core, and a shared 8MB L3 cache.

Our virtual cluster contains 13 physical servers with 7 servers running 4VMs per server, two for webserver with 1GB of RAM and 2VCPU each and another two for hadoop with 4GB of RAM with 2VCPU each. Four servers run 6 dedicated Hadoop VMs (each with their own disk). Two more servers run web clients.

The web server VM is *always* pinned to shared CPU cores and Hadoop VMs are pinned to either two dedicated or shared CPU cores depending on the server it runs. Xens Domain-0, which hosts drivers used by all VMs, is given the servers remaining cores.

Benchmarks: We use interactive workloads and batch workloads as our workloads. For transactional workloads, we use two applications: TPC-W, which models a three-tier online book store and Micro Web App, a PHP/MySQL application that emulates a multi-tier application and allows the user to adjust the rate and type of requests to control of CPU computation and I/O activities performed on the test system. For batch workloads, we choose the following Hadoop jobs. *PiEstimator*: estimates Pi value using 1 million points; *WordCount*: computes frequencies of words in 15GB

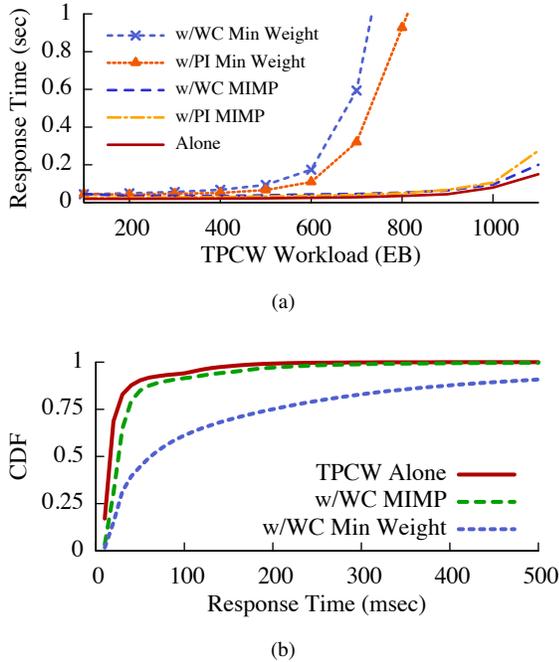


Figure 6. The MIMP scheduler provides response time almost equivalent to running a web application alone.

data; *Sort*: sorts 18GB data; *Grep*: finds match of randomly chosen regular expression on 6GB data; *TeraSort*: samples the 1GB input data and sort the data into a total order; *Kmeans*: clusters 6GB of numeric data. Both *Kmeans* and *Grep* are divided into two types of jobs.

B. Minimizing Scheduler Interference

We start our evaluation by studying how our Minimal Interference Scheduler is able to provide greater performance isolation when mixing web and processor intensive tasks. We repeat a variation of our original motivating experiment, and adjust the number of TPC-W clients when running either *Pi* or Word Count Hadoop jobs on a shared server. As expected, Figure 6(a) shows that the response time when using Xen’s default scheduler quickly becomes unmanageable, only supporting about 500 clients before interference causes the response time to rise over 100ms. In contrast, our MI scheduler provides performance almost equivalent to running TPC-W alone, allowing it to support twice the throughput before response time starts to rise. A closer look at the response time CDF in Figure 6(b) illustrates that MIMP incurs only a small overhead when there are 700 clients.

C. Task Affinity Models

In this section, we illustrate the accuracy of our task completion time models and how they guide slot allocation.

1) *TCT models*: Figure 7 shows the training data and model curves generated by MIMP. Each Hadoop VM has two cores that are shared with a Micro Web App VM. We

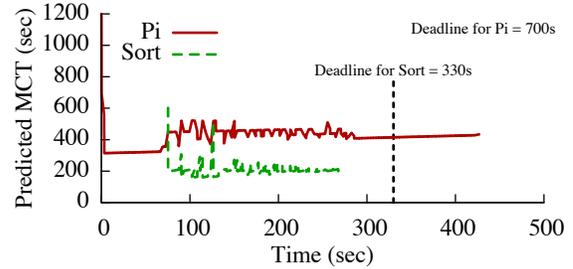


Figure 8. MIMP accurately predicts Map-phase completion time (MCT) and appropriately allocates slots to meet deadlines.

run a set of Hadoop jobs across our cluster using a randomly generated web workload ranging from 20 to 180% CPU utilization for each shared node. The x-axis represents the CPU utilization of the web VM before each task is started; we normalize the measured task completion time by the average TCT when running the same type of task on a node with no web workload.

These figures show the wide range of task completion times that are possible even for a fixed level of CPU availability. This variation occurs because the MI scheduler can give an unpredictable amount of CPU to the low priority Hadoop VM depending on fluctuations in the web workload.² Thus, it is quite difficult to make accurate predictions, although our models do still capture the overall trends.

When we apply these models to our case study workload described in Section VI-E, we find that 57% of the time our models over predict task completion time, and that the average over prediction is by 35%. The average under prediction is 29%. This is good since we would prefer our model to over predict task completion times, causing it to be more conservative, and thus less likely to miss deadlines.

2) *Total Map phase time prediction*: MIMP’s TCT models are used to predict whether a job will meet its deadline given its current slot allocation. Figure 8 shows how the predictions change as slots are allocated and removed from a job. We first start a *Pi* job at time 0 with a deadline of 700 seconds. Within 10 seconds, *Pi* has been allocated all of the available slots, so its predicted map phase completion time (MCT) quickly drops to about 370 seconds. At time 80 sec, a *Sort* job is started, causing the MIMP scheduler to divide the available slots between the two jobs. It reduces from *Pi*, but only enough to ensure that *Sort* will finish before its deadline. The predicted MCT of each job fluctuates as the number of slots it is given varies, but it remains accurate throughout the run.

D. Data Node I/O Saturation

Allocating more shared worker nodes to a Hadoop job will only increase performance if the data nodes that serve

²The variation is *not* simply related to differences in data node I/O levels, since even the *Pi* job (which does not make any storage accesses) sees a particularly high variation in TCT.

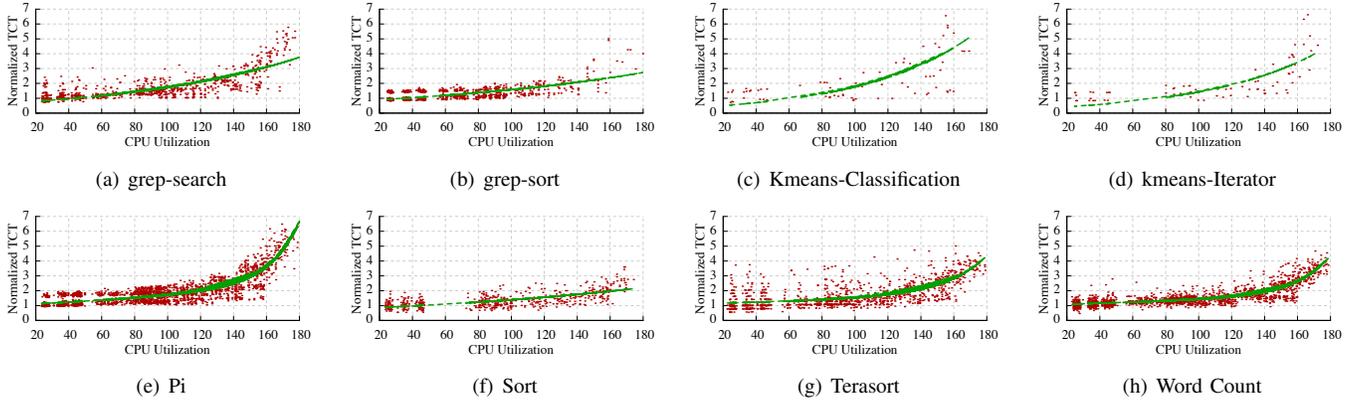


Figure 7. MIMP trains different models for each Hadoop job type.

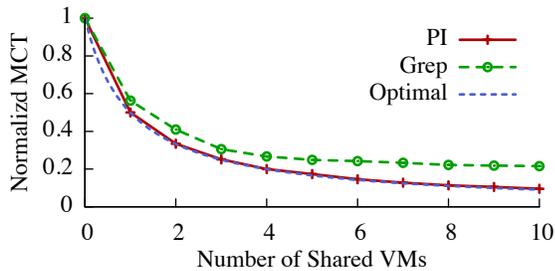


Figure 9. Data intensive Hadoop tasks see diminishing returns once the data nodes become saturated.

	FIFO	EDF	MIMP
#Miss Deadline	67	2	1
Total TCT(h)	72.17	72.61	67.95
#Failed Jobs	17	0	0
#Failed Tasks	1080	1	0
Avg Lateness(s)	217.8	6.2	7.9

Table II
WORKLOAD PERFORMANCE STATISTICS

that job are not overloaded. To evaluate this behavior, we measure the map phase completion time (MCT) when the number of available shared nodes is increased from 0 to 10; there are also three dedicated nodes that run map tasks and act as data nodes. Each shared node runs a web server that leaves 164% of the CPU available for the Hadoop VM.

Figure 9 shows the impact of the data node IO bottleneck on the Pi and Grep jobs. We normalize the map phase completion time by the measured time without any shared nodes. Initially, both job types perform similarly, but the Grep job soon sees diminishing returns as the data nodes become saturated. In contrast, Pi is an entirely CPU bound job, so it is able to scale optimally, i.e., ten shared servers lowers the completion time by a factor of nearly ten.

E. Case study

To evaluate our overall system, we perform a large scale experiment where we run a trace of Hadoop jobs on a shared cluster and evaluate the performance of three different job scheduling algorithms. We use a total of 20 Hadoop

VMs each with two cores: 6 dedicated hosts, 6 with a light web workload (20-35% CPU utilization), 6 with a medium load (85-95%), and 2 that are highly loaded (130-170%). We generate these workloads based on our observations of the DIT and Wikipedia data traces, although we use a higher overall utilization level than was found in those traces since this puts more stress on making intelligent scheduling decisions. The web workloads are generated using httpperf clients connected to our Micro Web App benchmark.

We generate a random Hadoop trace composed of our six representative job types. Jobs of each type arrive following a Poisson distribution; the mean inter-arrival period is used as the deadline for that type of job, with the exception of KMeans jobs which we set to have no deadline. The job trace lasts 2.5 hours and contains 174 jobs in total.

Scheduler Comparison: Table II shows the performance statistics of each job scheduler when processing this trace. Unsurprisingly, FIFO performs very poorly, missing deadlines for 67 out of 174 jobs, with an additional 17 jobs failing to complete at all. The EDF scheduler performs much better, but still misses two jobs, with an average lateness of 6.2 seconds. The total task completion time (i.e., the sum of all successful task execution times) is 72.61 hours for EDF; FIFO is slightly lower only because it has failed tasks which do not add to the total.

MIMP provides the best performance, missing only one job deadline by 7.9 seconds. Most importantly, it achieves this while using 4.66 hours less total execution time than EDF. This is possible because MIMP makes smarter decisions about which tasks to run on which nodes, better matching them to the available resources.

Job Distribution: To understand why MIMP improves performance, we now examine how each scheduler assigns workers to jobs. Figure 10(b) shows the number of slots assigned to each job during a 10 minute portion of the trace. This shows that EDF assigns all task slots to whatever job has the earliest deadline, even though some slots may be better suited for a different job type. In contrast, MIMP tends to run multiple jobs at the same time, allocating the

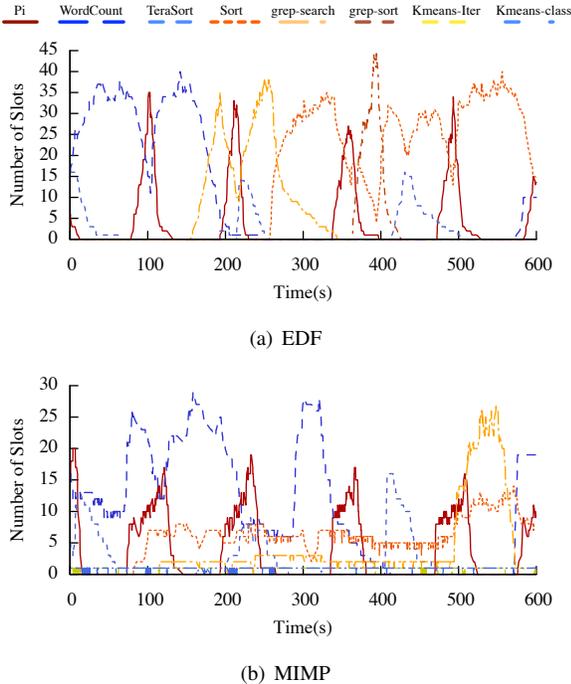


Figure 10. The slot allocation for the first 10 minutes - MIMP vs. EDF

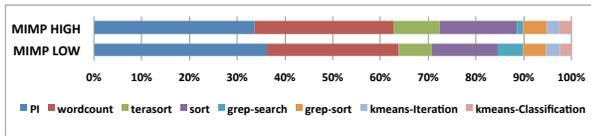


Figure 11. Job Distribution on high vs. low load server

best fitting slots to each one. While this can result in longer job completion times, MIMP still ensures that all jobs will meet their deadlines and improves overall efficiency due to resource affinity.

Figure 11 breaks down the percent of tasks completed by the highly loaded and lightly loaded servers when using MIMP. Since each job has a different arrival period, some jobs (such as Pi) have more total tasks than infrequent jobs (such as k-means). However, the results still follow our expectations based on the models shown in Figure 7. For example, Pi and grep-search have particularly high normalized TCT when the web server utilization rises, so MIMP runs relatively fewer of those tasks on the highly loaded servers. In contrast, sort’s completion time does not change much, so MIMP runs more of those tasks on the highly loaded servers.

VII. RELATED WORK

Virtual Machine Scheduling: Several previous works propose to improve the efficiency of Xen’s Credit CPU scheduler. For example, [3] proposes a modification to the scheduler that asynchronously assigns each virtual CPU to a physical CPU in order to reduce CPU sleep time. Lin et

al. [13] developed *VSched* that schedules batch workloads within interactive VMs without compromising the usability of interactive applications. The drawback of this system is that it was not designed to run on a cluster. Xi et al., use techniques from Real-Time scheduling to give stricter deadline guarantees to each virtual machine [21].

Other work has looked at avoiding interference between these tasks by careful VM placement [7], [23] or dedicating resources [11]. Paragon [7] proposes a heterogeneous and interference-aware data center scheduler. The system prefers to assign the applications on the heterogeneous hardware platform that the application can benefit from and have less interference with the co-scheduled applications. MIMP extends our preliminary study [25] to reduce interference through minor changes to Xen’s scheduler and then uses the residual resources for big data applications.

To improve I/O performance, Xu et al. [22] propose the use of *vTurbo* cores that have a much smaller time-slice compared to normal cores, reducing the overall IRQ processing latency. Cheng et al. [2] improves I/O performance for Symmetric MultiProcessing VMs by dynamically migrating the interrupts from a preempted VCPU to a running VCPU thereby avoiding interrupt processing delays. Our current focus is on shared environments where disk I/O is not the bottleneck for interactive applications, but view this as important future work.

Hadoop Scheduling & Modeling: Job scheduling in MapReduce environments has focused on topics like fairness [10], [24] and dynamic cluster sharing among users [17]. Bu et al. [1] propose a new hadoop scheduler based on the existing fair scheduler. They present an interference and locality-aware task scheduler for MapReduce in virtual clusters and design a task performance prediction model for an interference-aware policy. Morton et al. [15] provide a time-based progress indicator for a series of Map Reduce jobs, which can be used to predict job completion times. Polo et al. provide a system to dynamically adjust the number of slots provided for map reduce jobs on each host to maximize the resource utilization of a cluster and to meet the deadline of the jobs [16]. [19] decides the appropriate number of slots allocated to map and reduce based on the upper and lower bounds of batch workload completion time obtained from the history of job profiling. Our work is distinct from prior work in that we estimate the job completion time of the batch jobs that are running in clusters with unpredictable resource availability due to other foreground applications.

Previous works [5], [9], [12], [14] have shown heterogeneous cluster designs wherein a core set of dedicated nodes running batch jobs are complemented by residual resources from volunteer nodes, or in some cases using “spot instances” from EC2 [4]. The closest work to ours is by Clay et al [5]. They present a system that determines the appropriate cluster-size to harness the residual resources

of under utilized interactive nodes to meet user-specified deadlines and minimize cost and energy. Our work extends this by focusing on how groups of jobs should be scheduled across a shared cluster in order to both minimize interference and meet job deadlines.

VIII. CONCLUSIONS

Virtualization allows servers to be partitioned, but resource multiplexing can still lead to high levels of performance interference. This is especially true when mixing latency sensitive applications with data analytic tasks such as Hadoop jobs. We have designed MIMP, a Minimal Interference, Maximal Progress scheduling system that manages both VM CPU scheduling and Hadoop job scheduling to reduce interference and increase overall efficiency.

MIMP works by exposing more information to both the Hadoop job scheduler and the Xen CPU scheduler. By giving these systems information about the priority of different VMs and the resources available on different servers, MIMP allows cluster utilization to be safely increased. MIMP allows high priority web applications to achieve twice the throughput compared to the default Xen scheduler, and has response times nearly identical to running the web application alone. Despite the increased variability this causes in Hadoop task completion times, MIMP is still able to meet more deadlines than an Earliest Deadline First scheduler and lowers the total execution time by nearly five hours in one of our experiments.

Acknowledgments: We thank the reviewers for their helpful suggestions. This work was supported in part by NSF grant CNS-1253575 and the National Natural Science Foundation of China Grant No. 61370059, 61232009, and Beijing Natural Science Foundation Grant No. 4122042.

REFERENCES

- [1] X. Bu, J. Rao, and C.-z. Xu. Interference and Locality-aware Task Scheduling for MapReduce Applications in Virtual Clusters. In *Proc. of HPDC*, 2013.
- [2] L. Cheng and C.-L. Wang. vBalance: using interrupt load balance to improve I/O performance for SMP virtual machines. In *Proc. of SOCC*, 2012.
- [3] T. Chia-Ying and L. Kang-Yuan. A Modified Priority Based CPU Scheduling Scheme for Virtualized Environment. *International Journal of Hybrid Information Technology*, 6(2), 2013.
- [4] N. Chohan, C. Castillo, M. Spreitzer, M. Steinder, A. Tantawi, and C. Krintz. See spot run: using spot instances for mapreduce workflows. In *Proc. of HotCloud*, 2010.
- [5] R. B. Clay, Z. Shen, and X. Ma. Accelerating Batch Analytics with Residual Resources from Interactive Clouds. In *Proc. of MASCOTS*, 2013.
- [6] J. Dean and S. Ghemawat. MapReduce: Simplified Data Processing on Large Clusters. *Commun. ACM*, 51, 2008.
- [7] C. Delimitrou and C. Kozyrakis. Paragon: QoS-aware Scheduling for Heterogeneous Datacenters. In *Proc. of ASPLOS*, 2013.
- [8] A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel. The cost of a cloud: research problems in data center networks. *SIGCOMM Comput. Commun. Rev.*, 39(1), Dec. 2008.
- [9] H. Herodotou, F. Dong, and S. Babu. No one (cluster) size fits all: automatic cluster sizing for data-intensive analytics. In *Proc. of SOCC*, 2011.
- [10] M. Isard, V. Prabhakaran, J. Currey, U. Wieder, K. Talwar, and A. Goldberg. Quincy: fair scheduling for distributed computing clusters. In *Proc. of SOSP*, 2009.
- [11] E. Keller, J. Szefer, J. Rexford, and R. B. Lee. NoHype: virtualized cloud infrastructure without the virtualization. In *Proc. of ISCA*, 2010.
- [12] G. Lee, B.-G. Chun, and H. Katz. Heterogeneity-aware resource allocation and scheduling in the cloud. In *Proc. of HotCloud*, 2011.
- [13] B. Lin and P. A. Dinda. Vsched: Mixing batch and interactive virtual machines using periodic real-time scheduling. In *Proc. of Super Computing*, 2005.
- [14] H. Lin, X. Ma, J. Archuleta, W.-c. Feng, M. Gardner, and Z. Zhang. MOON: MapReduce On Opportunistic eNvironments. In *Proc. of HPDC*, 2010.
- [15] K. Morton, A. Friesen, M. Balazinska, and D. Grossman. Estimating the progress of MapReduce pipelines. In *Proc. of ICDE*, 2010.
- [16] J. Polo, C. Castillo, D. Carrera, Y. Becerra, I. Whalley, M. Steinder, J. Torres, and E. Ayguadé. Resource-aware adaptive scheduling for mapreduce clusters. In *Proc. of Middleware*, 2011.
- [17] T. Sandholm and K. Lai. Dynamic proportional share scheduling in Hadoop. In *Proc. of JSSPP*, 2010.
- [18] B. Sharma, T. Wood, and C. R. Das. HybridMR: A Hierarchical MapReduce Scheduler for Hybrid Data Centers. In *Proc. of ICDCS*, 2013.
- [19] A. Verma, L. Cherkasova, V. Kumar, and R. Campbell. Deadline-based workload management for MapReduce environments: Pieces of the performance puzzle. In *Proc. of NOMS*, 2012.
- [20] C. A. Waldspurger. Memory resource management in VMware ESX server. *Proc. of OSDI*, 36(SI), Dec. 2002.
- [21] S. Xi, J. Wilson, C. Lu, and C. Gill. RT-Xen: towards real-time hypervisor scheduling in xen. In *Proc. of EMSOFT*, 2011.
- [22] C. Xu, S. Gamage, H. Lu, R. Kompella, and D. Xu. vTurbo: accelerating virtual machine I/O processing using designated turbo-sliced core. In *Proc. of Usenix ATC*, 2013.
- [23] Y. Xu, Z. Musgrave, B. Noble, and M. Bailey. Bobtail: Avoiding Long Tails in the Cloud. In *Proc. of NSDI*, 2013.
- [24] M. Zaharia, D. Borthakur, J. Sen Sarma, K. Elmeleegy, S. Shenker, and I. Stoica. Delay scheduling: a simple technique for achieving locality and fairness in cluster scheduling. In *Proc. of EuroSys*, 2010.
- [25] W. Zhang, S. Rajasekaran, and T. Wood. Big Data in the Background: Maximizing Productivity while Minimizing Virtual Machine Interference. In *Proc. of Workshop on Architectures and Systems for Big Data*, 2013.
- [26] Z. Zhang, L. Cherkasova, A. Verma, and B. T. Loo. Performance Modeling and Optimization of Deadline-Driven Pig Programs. *ACM TAAS*, 8, 2013.