

HybridMR: A Hierarchical MapReduce Scheduler for Hybrid Data Centers

Bikash Sharma*, Timothy Wood†, Chita R. Das*

*Department of Computer Science and Engineering, The Pennsylvania State University, University Park, PA
{bikash, das}@cse.psu.edu

†Department of Computer Science, The George Washington University, Washington, DC
timwood@gwu.edu

Abstract— Virtualized environments are attractive because they simplify cluster management, while facilitating cost-effective workload consolidation. As a result, virtual machines in public clouds or private data centers, have become the norm for running transactional applications like web services and virtual desktops. On the other hand, batch workloads like MapReduce, are typically deployed in a native cluster to avoid the performance overheads of virtualization. While both these virtual and native environments have their own strengths and weaknesses, we demonstrate in this work that it is feasible to provide the best of these two computing paradigms in a hybrid platform.

In this paper, we make a case for a hybrid data center consisting of native and virtual environments, and propose a 2-phase hierarchical scheduler, called *HybridMR*, for the effective resource management of interactive and batch workloads. In the first phase, *HybridMR* classifies incoming MapReduce jobs based on the expected virtualization overheads, and uses this information to automatically guide placement between physical and virtual machines. In the second phase, *HybridMR* manages the runtime performance of MapReduce jobs collocated with interactive applications in order to provide best effort delivery to batch jobs, while complying with the Service Level Agreements (SLAs) of interactive applications. By consolidating batch jobs with over-provisioned foreground applications, the available unused resources are better utilized, resulting in improved application performance and energy efficiency. Evaluations on a hybrid cluster consisting of 24 physical servers and 48 virtual machines, with diverse workload mix of interactive and batch MapReduce applications, demonstrate that *HybridMR* can achieve up to 40% improvement in the completion times of MapReduce jobs, over the virtual-only case, while complying with the SLAs of interactive applications. Compared to the native-only cluster, at the cost of minimal performance penalty, *HybridMR* boosts resource utilization by 45%, and achieves up to 43% energy savings. These results indicate that a hybrid data center with an efficient scheduling mechanism can provide a cost-effective solution for hosting both batch and interactive workloads.

Keywords—*Hybrid Data Center, Hadoop MapReduce, Resource Management, Scheduling, Virtualization, Performance, Energy*

I. INTRODUCTION

Virtualization has evolved as a key technology to support agile and dynamic IT infrastructure, which forms the base of large distributed systems like data centers and clouds. Virtualization enables autonomic management of underlying hardware, server sprawl reduction through workload consolidation, and dynamic resource allocations for better throughput and energy efficiency. Consequently, major cloud providers like Amazon EC2, RackSpace and Microsoft Azure, utilize server virtualization to efficiently share resources among customers, and allow for rapid elasticity.

Despite these numerous benefits, virtualization introduces an extra software layer to the system stack, incurring overheads to native performance. Analyses with generic benchmarks have shown the virtualization overheads to be around 5% for computation and 15% for I/O workloads [10]. While these virtualization overheads have continued to fall with the introduction of virtualization-aware hardware [22], the effects are still large enough that many companies like Google and Facebook still prefer physical machines (PMs) in their data centers to run their core applications like Web search [7]. For data analytics frameworks such as Hadoop MapReduce [2], which allow for efficient large scale distributed computation over massive data sets, virtualized cloud platforms seem like a natural fit for providing elastic scalability. However, virtualization in clouds is known to incur performance overheads, particularly when used for I/O-bound MapReduce activities [20], [21]. As a result, MapReduce users are often left with the choice of either maximizing performance with a native cluster or obtaining ease of use and resource efficiency with a virtualized environment.

Today's data centers offer two different modes of computing platforms - native and virtual clusters. Both these environments have their own strengths and weaknesses. For example, a native cluster is better for batch workloads like MapReduce from the performance perspective, lowers SLA violations, but usually suffers from poor utilization, and high hardware and power cost. A virtual cluster, on the other hand, is attractive for interactive workloads from consolidation and cost standpoints, but may not provide competitive performance like a native cluster, and incurs higher SLA infringements. Intuitively, a hybrid platform consisting of native and virtualized cluster should be able to exploit the benefits of both environments for providing a better cost-effective platform. In this paper, we explore this design alternative, which we call *hybrid data center*, and demonstrate its advantages for supporting both interactive and batch workloads, and achieving the right balance between all these design criteria, making it a desirable cluster configuration option.

Transactional applications like interactive web services and virtual desktop environments are prime candidates for virtualization. For supporting the SLA requirements of interactive applications, resources are generally over-provisioned, leading to poor utilization [15]. To exploit the potentials of both native and virtual cluster, we leverage the over-provisioning of bursty interactive applications by intelligently consolidating batch MapReduce jobs using the spare resources available on a virtualized platform. This allows for reaping the benefits of high consolidation in multi-tenant systems. This hybrid infrastructure presents different trade-offs across various design

metrics like performance, cost, energy and resource utilization between native, virtual and hybrid design choices.

For facilitating such a hybrid cluster platform, this paper presents the design and implementation of a 2-phase hierarchical scheduler, called *HybridMR*, that judiciously allocates virtual and physical resources to applications. Contrary to the traditional workload placement schemes that completely isolate batch MapReduce and interactive workloads, *HybridMR* consolidates the workload mix in a heterogeneous infrastructure to achieve higher performance, utilization and energy efficiency. The design of such a scheduler requires the knowledge of how different MapReduce jobs are impacted by virtualization overheads, estimates of their resource needs, and an understanding of how batch jobs will impact the performance of interactive jobs collocated on virtual machines (VMs) of the same host. *HybridMR* addresses these challenges through a 2-level scheduler design. Its first phase places MapReduce jobs on physical or virtual nodes depending on the expected virtualization overheads. Once a set of MapReduce jobs have been selected to run on the virtual cluster, along with interactive applications, the second phase scheduler decides *how much* resources can be safely assigned to them. For this, it makes use of statistical predictive models for understanding the runtime resource interference between the interactive and batch jobs, and employs dynamic resource management techniques to provide the best effort delivery to MapReduce jobs, while upholding the SLAs of interactive applications. *HybridMR* targets institutional and enterprise Intranet environments, where the data centers can run in a hybrid fashion, comprising of evolving cluster mix of native and virtual machines.

In summary, this paper’s main contributions are:

- We make a case for hybrid data centers with native and virtual machines for co-hosting both transactional and batch workloads towards exploiting the best of these two computing worlds – native performance with virtualization benefits.
- We have designed and implemented *HybridMR*, a 2-phase hierarchical scheduler, for the effective placement of MapReduce jobs on a virtualized platform, while upholding the SLAs of interactive applications.
- Through detailed analysis, we demonstrate the trade-offs and benefits of running Hadoop in a virtualized platform, and leverage the insights in the design of *HybridMR*. Evaluations on a hybrid data center consisting of 24 physical servers, and 48 virtual machines, with diverse workload mix of representative transactional and MapReduce applications, demonstrate that *HybridMR* can achieve up to 40% improvement in job completion times over virtual-only cluster, 45% improvement in resource utilization, and 43% savings in energy, both relative to native-only cluster. Further, we demonstrate the possibility and flexibility to dynamically vary the native and virtual cluster configurations to accommodate for the variations in the workload mix for maximizing the performance and energy benefits. These results demonstrate the potentials of a hybrid cluster configuration as a cost-effective solution, when compared to either-or native and virtual modes of computation.

The remainder of this paper is organized as follows: Section II provides empirical evaluations of MapReduce in a virtual environment. Section III describes the design of *HybridMR*. Section IV details the experimental testbed and evaluation results. Section V discusses related work. The

conclusions and future work are outlined in Section VI.

II. MAPREDUCE IN VIRTUAL ENVIRONMENT

Here, we address the following questions that are critical for running Hadoop MapReduce on a hybrid cluster:

Q1. What are the challenges, benefits and trade-offs of running Hadoop on a virtual versus equivalent native cluster?

Q2. What system and application level changes are required to realize Hadoop’s deployment on hybrid data centers?

Q3. How does the ‘data sticky-ness’ of Hadoop affects the overall system performance in a virtualized Hadoop cluster?

First, we perform detailed analyses and present empirical evidences demonstrating the various challenges, benefits and trade-offs involved in the deployment of Hadoop on a virtual cluster. We also highlight the key design choices that impact Hadoop performance in a virtualized environment. Second, we present a split architecture of Hadoop that addresses the movement of large data across a virtual cluster during live VM migrations. For these studies, each VM is configured with 1 vCPU and 1024 MB Memory, and runs on dual-core, 4 GB RAM servers. We use a total of 1-24 PMs and 1-48 VMs, depending on the requirements of each experiment. Details about the experimental platform, including the MapReduce benchmarks used are described in Section IV.

In response to *Q1* and *Q2*, we start with comparing the performance of MapReduce benchmarks on native versus virtual environments. In Figure 1(a), the Y-axis represents the percentage increase in job completion time (JCT), with respect to a physical cluster. We observe that with the increase in the number of VMs per PM, the performance (JCT) of I/O-bound jobs like *Twitter*, *Wcount*, *DistGrep* and *Sort* in virtual cluster is 7-24% worse than the physical cluster. For CPU bound jobs like *PiEst* and *Kmeans*, the performance difference is within 8%. Moreover, across all benchmarks, the performance gap widens with the data size increase, as evident in Figure 1(b) (only *Sort* shown). To investigate this, we benchmark the performance of Hadoop Distributed File System (HDFS) in terms of read and write I/O, read and write throughput on a virtual cluster normalized with respect to the corresponding native cluster. We use the Hadoop TestDFSIO [2] benchmark, and the results are shown in Figure 1(c). We notice that the virtual cluster performs relatively worse than the native cluster, and again this performance gap broadens with the data size increase. The reasons for this are: (i) increase in contention for shared I/O resources across multiple VMs, when large amount of data is transferred between map and reduce phase; (ii) poor performance of HDFS when multiple VMs concurrently read/write data blocks from and to HDFS; and (iii) increase in the number of stragglers, hence more speculative tasks [16], causing inefficient resource utilization.

In another related experiment, we study the impact of Hadoop VMs crossing multiple hosts due to large size of the virtual cluster, limited server resources, and distributed data. For this, we create a 16 VMs Hadoop cluster. In *Cross-Host*, the 16 VMs are equally distributed across 8 dual-core PMs, while in *Same-Host*, all the 16 VMs are equally consolidated on 2 PMs. Figure 2(a) shows the JCT of *Sort* in these two scenarios. The poor performance of *Cross-Host* compared to *Same-Host* is due to the increase in the network I/O delay caused by remote VMs communication in *Cross-Host*. Further, the JCTs in each case increase as the input data scales up. An interesting observation here is that despite the *Cross-Host*

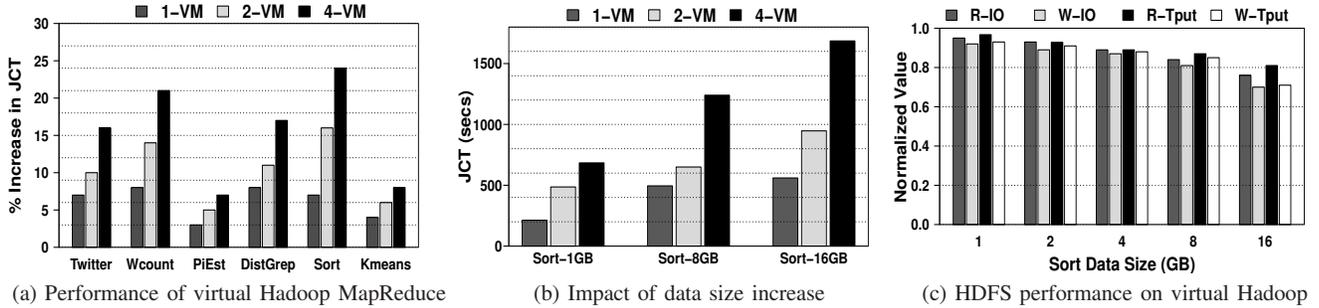


Figure 1: Illustration of virtualization overheads on Hadoop performance. Y-axis in (a), (c) normalized w.r.t. an equivalent physical cluster. Figure 1(c): R/W-IO: Read/Write IO in MB/sec (average IO rate); R/W-Tput: Read/Write-Throughput in MB/sec ((total number of bytes)/(sum of processing times)). In these experiments, 48 VMs are used.

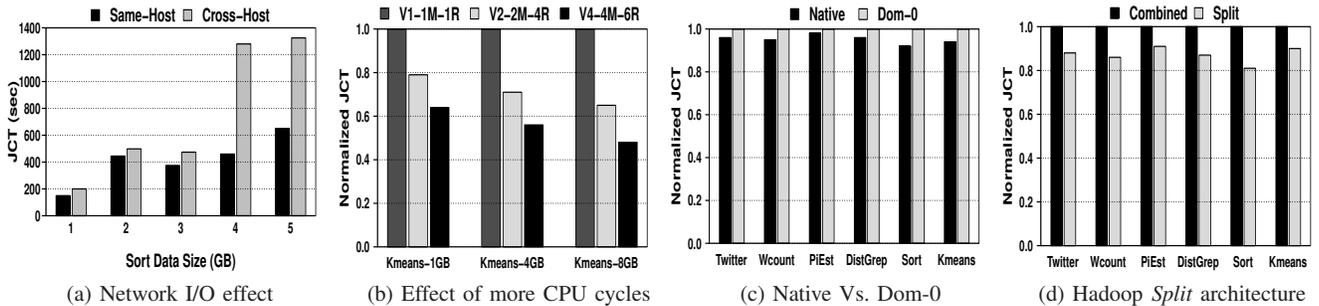


Figure 2: In (a), 16 VMs are used; 48 VMs are used in (b) and (c). In (b): V1, V2, V4 denote 1, 2, 4 VMs per PM. M, R denote number of Map and Reduce tasks. Y-axis in (b) and (c) is normalized w.r.t. Native; Y-axis in (d) is normalized w.r.t. Combined.

having access to more cores (each VM gets 1 vCPU) compared to *Same-Host* (each VM gets 0.25 vCPU), the performance of the former is still worse than the latter. Thus, when the data size and the number of concurrent running map/reduce tasks increase, the network communication overheads become the main resource bottleneck, and dictates the overall performance.

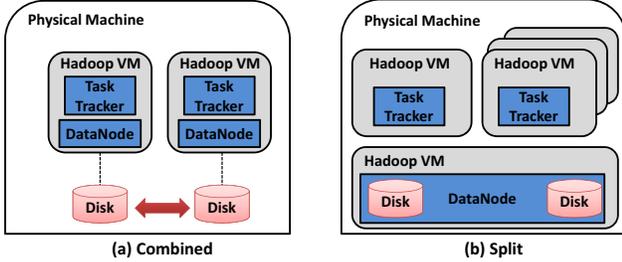
The performance of Hadoop MapReduce in a virtual environment can also be better in certain scenarios. For example, CPU-bound jobs like *kmeans* can benefit from having more VMs available, as shown in Figure 2(b). Further, this performance increase is higher for larger data sizes. This is because, a CPU-bound job can leverage higher number of map/reduce tasks to finish faster due to the availability of more compute cycles, and hence more slots, with multiple VMs on multi-core hosts. Note that in Figure 1(a), CPU-bound *kmeans* tends to perform worse with more number of VMs per PM. However, this is opposite to the results shown in Figure 2(b). This is because, for Figure 2(b) scenario, more number of map and reduce tasks are launched to exploit the free CPU cycles (*i.e.*, in form of Hadoop slots), compared to the scenario in Figure 1(a), where fewer number (2 by default) of concurrent mappers and reducers are present.

We next explore leveraging Xen’s split architecture to run Hadoop in the privileged domain (Dom-0). Figure 2(c) shows that Dom-0 provides near native performance, with overheads less than 5% on average. This opens up the possibility of a ‘flexibly virtualized’ cluster, where some machines can be easily transitioned from running VMs to quasi-native applications in Dom-0 [22]. This supports the case for hybrid clusters, whose configurations can be flexibly and on-demand adjusted.

Finally, in response to Q3, the other concern in the deploy-

ment of Hadoop in a virtualized environment is the challenge to deal with the inherent ‘data sticky-ness’ of Hadoop, *i.e.*, how to handle the movement of large amount of data living on Hadoop VMs, which may lead to inefficiency in terms of time and network communication overheads, thus, defeating many of the benefits of virtualization. This also reduces the effectiveness of resource elasticity due to the tight coupling between the compute and data storage layer. A prospective workaround to this problem is to place the TaskTracker (compute node) and DataNode (storage node) on separate VMs [35]. This *split* architecture (see Figure 3) maintains data locality, and removes the constraint of moving large amounts of data around the cluster during VM migrations, since the data stays resident in the HDFS, while the number of map/reduce (compute) VMs can vary as required. We did some preliminary evaluations to assess the performance implications of this *split* architecture. Figure 2(d) shows the results. We observe on an average 12.8% improvement in JCT over the default Hadoop (*Combined*). We adopt this *split* Hadoop architecture in this paper.

To summarize, the above analyses demonstrate that the performance of virtualized Hadoop is governed by a multitude of design factors, and thus, coming up with an optimal configuration both at the job and cluster level is non-trivial, thereby making scheduling and resource management difficult. Some key observations are: (i) Hadoop performance on a virtual cluster is relatively worse than the equivalent physical cluster. The magnitude of performance degradation varies depending on the nature of the job. For example, jobs which are I/O and network intensive like *Sort* suffer more performance degradation than CPU-bound jobs like *PIest*; (ii) input data size affects the magnitude of performance difference in direct proportion; (iii)



Compute and storage layer are combined on a single VM in (a), while they are split across separate VMs in (b). Efficient virtual networking between VMs in a host allows (b) to maintain compute-data locality.

Figure 3: Hadoop Split Architecture.

Hadoop’s performance difference between native and Dom-0 is marginal; and (iv) *Split* Hadoop architecture is an optimized deployment on a virtual environment.

III. DESIGN OF HYBRIDMR

This section describes the overall architecture of *HybridMR* (as shown in Figure 4), which operates in two phases. In the first phase, *HybridMR* attempts to classify incoming MapReduce jobs based on the expected virtualization overheads, and uses that information to automatically guide placement between physical and virtual machines. The second phase performs dynamic resource management to minimize interference and improve performance of collocated interactive and batch jobs. Specific details of these phases are described below:

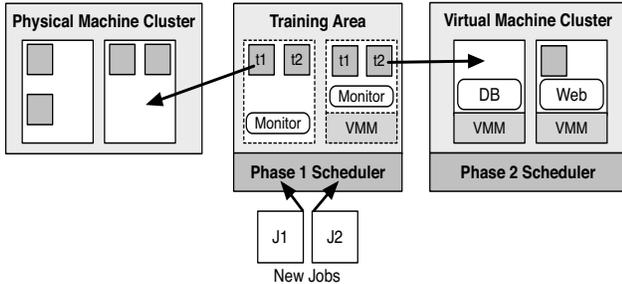


Figure 4: Overview of *HybridMR*.

A. Phase 1 Scheduler

The main goal of this phase is to differentiate between workloads that should be scheduled on physical machines or virtual machines running in a hybrid data center. Since, our objective is to harness the spare resources on VMs running interactive applications, the interactive applications by default are assigned to the virtual cluster, and the placement of the MapReduce jobs is governed by this phase. Thus, when a MapReduce job arrives, it is initially started separately on a small training cluster containing both physical and virtual environments, respectively (see Figure 4). We utilize statistical profiling techniques to estimate the JCTs of MapReduce jobs (refer Section III-A1). By comparing the estimated JCTs of the two instances of the job, corresponding to its run on native and virtual machines, the level of performance overhead (as quantified by JCT) incurred by the virtualization layer is estimated. If the overhead is not significant, then the job is selected for deployment on the virtual cluster, else it is run on a separate physical cluster. To estimate the JCT, we leverage the following job profiling methodology:

1) *Job Profiling*: We profile MapReduce jobs to estimate their JCTs prior to execution. MapReduce jobs are data-intensive and massively parallel, hence their JCTs are predominantly dependent on two factors: (i) input data set size; and (ii) resource set size, *i.e.*, number of nodes in the cluster. The estimation scheme accumulates a database of past execution history of jobs (in terms of job completion times, corresponding to different input data set sizes and cluster sizes). During training, a job is run on a representative small cluster and with a smaller data set, and extrapolation techniques are used to estimate the run-time if the job were run on the full cluster and data set. The profile database maintains separate run-times for map and reduce phases to account for the differences across phases. The exact association of a job’s profile with the cluster size and input size is described next.

Cluster size: To quantify the dependence of cluster size on job completion time, we measure the total time as well as the time taken by the map and reduce phases separately, against different cluster sizes. In Figure 5(a), we observe that for a given data size, the JCT of a job follows an inverse relation to the cluster size. Similar inverse relation follows for the completion time of map phase as shown in Figure 5(b). However, the reduce phase completion time follows a piecewise non-linear relation with the cluster size as shown in Figure 5(c). This is due to the fact that the reduce phase’s dependence on data size is highly erratic [12].

Data size: For a given cluster size, the JCT is almost linearly proportional to the input data size, as demonstrated in Figure 5(d). This indicates that by simple linear extrapolation, we can accurately estimate the JCT of a new job instance with a different data size. However, when sufficient historical job profiles are not available, we can still profile on a subset of data, and then use linear regression based extrapolation to derive the JCT corresponding to the actual input size.

Thus, by building the job profiles using different data sizes and/or cluster sizes, this module estimates the JCT. The end-to-end mechanism for job profiling based on the above logic is outlined in Algorithm 1. To elaborate, for an arriving MapReduce job, Algorithm 1 searches the database for matching cluster size and/or data size. On a match, the corresponding profiled JCT estimate is retrieved. When there is no match, an appropriate extrapolation technique based on the map and reduce phase is used to derive the estimated JCT.

We build the database of job profiles, averaged across 3 runs of each job instance. Note that similar procedure for profiling MapReduce jobs has been explored in [12], [33]. There are also complex techniques [18], [33], where profiling is done for fine-grained stages of MapReduce framework. Also, the technique can be extended for online profiling [12], [33]. However, we observed that the above simple profiling technique works well in practice for our purpose, incurring low overheads, while producing reasonable prediction accuracy. For example, Figure 6(a) shows the actual and estimated JCTs obtained from profiling of `Sort` on 10 GB. We got similar results for other MapReduce benchmarks, and our profiling scheme results in an average error of 10.8% with a standard deviation of 9.7%.

2) *Placement of MapReduce Jobs*: When a job is submitted to the system, depending on the type of the job (transactional or batch), and its desired completion time, the heuristic steps outlined in the Algorithm 2 determine its initial placement on the native or virtual nodes. Note that, Phase I simply steers the initial placement of the job, whereas the exact resource

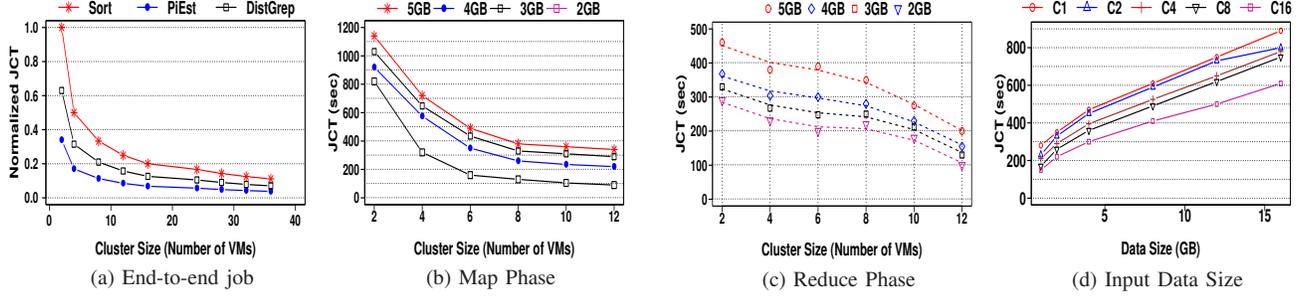


Figure 5: Dependence of job completion time on cluster size (end-to-end, map and reduce phase), and input data size. In (d), C1, C2, C4, C8, C16 represent virtual clusters with 1, 2, 4, 8, 16 nodes, respectively. `Sort` is used in (b), (c) and (d).

Algorithm 1 MapReduce Job Profiling Algorithm.

Input: Q : queue of incoming jobs, each specifying the input data size and/or cluster size; $DB_{profile}$: profile database, containing historic observations of job completion times (JCTs) (end-to-end and separate for map and reduce phases), along with the corresponding cluster sizes and input data set sizes.

- 1: **for** each job J_i in $Q=J_1, J_2, \dots, J_n$ **do**
- 2: **if** `LOOKUP_CLUSTER_SIZE($DB_{profile}, J_i$)` \neq `NULL` & `LOOKUP_DATA_SIZE($DB_{profile}, J_i$)` \neq `NULL` **then**
- 3: $JCT_{estimated} = \text{Retrieve}(DB_{profile}, J_i(csize), J_i(dsize))$
- 4: **else if** $DB_{profile}$ does not contain exact match for J_i 's input configuration **then**
- 5: **if** $DB_{profile}$ contains different data size values for the same cluster size (see Figure 5 (d)) **then**
- 6: Do linear extrapolation to get $JCT_{estimated}$
- 7: **else if** $DB_{profile}$ contains different cluster size values for the same data size **then**
- 8: Do separate Map and Reduce phase based extrapolation to get $JCT_{estimated}$ (see Figures 5 (a), (b), (c))
- 9: **end if**
- 10: **end if**
- 11: **return** $JCT_{estimated}$
- 12: **end for**

Algorithm 2 Job Placement Algorithm.

Input: Q : queue of incoming jobs; $Inputload$: number of clients for MapReduce job; $P_CLUSTER$: cluster of physical nodes; $V_CLUSTER$: cluster of virtual nodes; $JCT_{desired}[]$: vector of jobs desired completion times.

- 1: **for** each job J_i in $Q=J_1, J_2, \dots, J_n$ **do**
- 2: **if** $J_i \in$ transactional workload **then**
- 3: Place J_i on $V_CLUSTER$
- 4: **else if** $J_i \in$ batch MapReduce workload **then**
- 5: Profile J_i using Algorithm 1 to obtain the vector of estimated job completion time ($JCT_{estimated}[]$).
- 6: **if** $JCT_{estimated}[i] \geq JCT_{desired}[i]$ **then**
- 7: Place J_i on $P_CLUSTER$
- 8: **else**
- 9: Place J_i on $V_CLUSTER$
- 10: **end if**
- 11: **end if**
- 12: **return** jobs assigned to $P_CLUSTER$ and $V_CLUSTER$
- 13: **end for**

configuration of VMs or PMs, where the job will run, is determined by the Phase II scheduler, as described next.

B. Phase II Scheduler

The goal of the Phase II Scheduler is to efficiently manage the resources of the virtual cluster across transactional

and MapReduce jobs in order to comply with the SLA of interactive jobs, while providing the best effort performance guarantees to the MapReduce applications. Figure 7 shows the overall architecture of the Phase II Scheduler. It is composed of two main components: (i) a *Dynamic Resource Manager (DRM)*; and (ii) an *Interference Prevention System (IPS)*. The DRM monitors the available capacity on each VM to guide the placement of MapReduce jobs within the virtual cluster. DRM records the resource consumption and completion times of each task. This information is used to build a model for each MapReduce job that correlates the resource allocation to the constituent task completion time, allowing the scheduler to make intelligent placement decisions. The IPS is an online monitor that observes the performance of the interactive applications within the cluster to detect when they are not receiving sufficient resources to meet their demands. If any interference is detected, then the responsible map/reduce tasks are properly handled. For example, the VM running the task can have its resource share decreased, can simply be paused, or it can even be migrated to a different host to mitigate the observed interference. Even if the VM is aborted, the correctness of the corresponding MapReduce job is not affected, since the MapReduce master would initiate speculative execution [16], assuming the task as a prospective straggler.

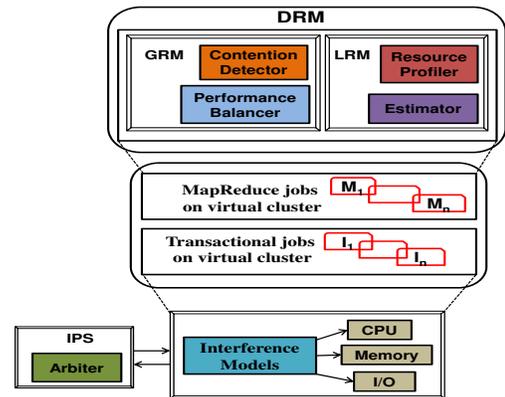


Figure 7: Components of the Phase II Scheduler.

To illustrate the level of interference which may be caused due to collocated VMs and/or contending tasks within the same VM, we conduct a study, where 4 VMs are deployed on a quad-core physical server, and run a mix of CPU and I/O bound MapReduce jobs. Each VM is pinned to a core,

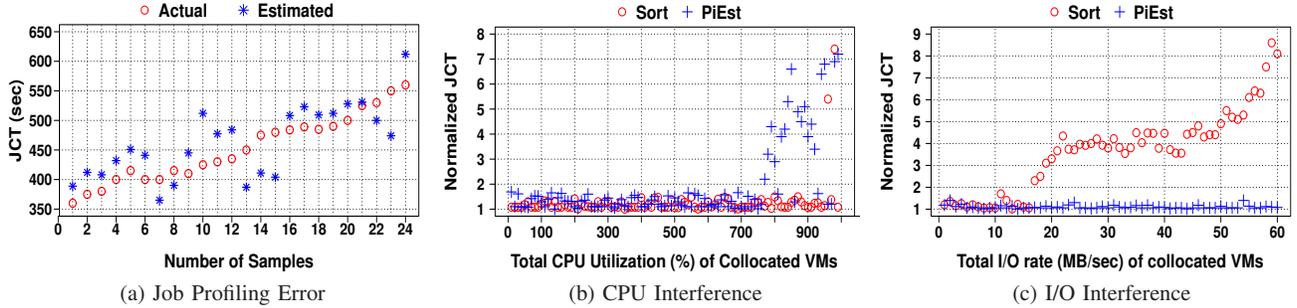


Figure 6: (a) Profiling error in Phase I; (b) and (c) show slowdown of JCT due to CPU and I/O interference from collocated VMs. Y-axis in (b) and (c) are normalized corresponding to the case when the jobs (Sort and PiEst) are run in isolation.

and 8 threads are concurrently run, sharing a single disk. We benchmark the completion time of CPU-bound PiEst and I/O-bound Sort. We first run each job on a single VM in isolation and note the JCT. We then measure the corresponding JCT when these jobs are run together with 3 other instances of PiEst and Sort, running on other VMs. Figure 6(b) shows the JCT of PiEst normalized to the corresponding JCT when run in isolation. The X-axis denotes the total CPU utilization, represented as the percentage of a single thread. We observe the slowdown in the run-time when the CPU interference from other background jobs increases, while I/O-bound Sort remains unaffected. Similar observation follows for I/O-bound Sort (Figure 6(c)), where we can see exponential increase in JCT due to increased I/O contention from other collocated VMs. Thus, understanding the run-time interference, and performing dynamic resource adjustment is required to mitigate such performance degradation. This approach is adopted by the Phase II Scheduler.

1) *Dynamic Resource Manager (DRM)*: The goal of the DRM is to build interference prediction models based on the run-time resource profiles of collocated tasks of MapReduce jobs and other co-hosted interactive jobs in order to estimate the task slowdown caused by resource contention. These estimations are used to orchestrate the optimal resource allocations via dynamic resource management to minimize the interference, and improve performance. DRM consists of two main components (a) a *Global Resource Manager (GRM)*; and (b) a *Local Resource Manager (LRM)*. The GRM consists of two sub-components: (i) a *Contention Detector* that dissects the cause of resource contention and identifies both resource-deficit and resource-hogging tasks; and (ii) a *Performance Balancer* that leverages the run-time resource estimations from each LRM to suggest the resource adjustments based on the global coordinated view of all tasks. The LRM also consists of two sub-components: (i) a *Resource Profiler* that collects and profiles the run-time resource usage of tasks at each node; and (ii) an *Estimator* that constructs statistical estimation models of a task’s run-time performance as a function of its resource allocations. Due to space constraints, we are unable to provide specific details. For further details about each of these components, please refer to our prior work [31].

DRM performs the following two functions: (i) Resource bottleneck detection: Here, the GRM based on the resource usage feedback from all LRMs identifies *which task* is experiencing bottleneck for *which resource*, and on *which TaskTracker*. The rationale and scheme behind this detection mechanism is outlined in [31]; and (ii) Resource bottleneck mitigation: The

LRMs after getting information about both resource deficit and resource hogging tasks from the GRM, invoke their Estimator module to estimate the tasks completion times. The Estimator builds predictive models for the completion time of a task as a function of its resource usage/allocation. The difference between the estimated and the current resource allocation is the resource imbalance that has to be dynamically allocated to the resource deficit task.

The Estimator module is responsible for building predictive performance interference models for tasks’ run-time resource usage/allocations (in terms of CPU, memory and I/O resources). In a virtualized environment, the interference is contributed by both collocated tasks within the same VM, and the co-hosted VMs on the same tasks host. It employs statistical regression models to obtain the estimated optimal resource allocation for the next epoch of a task’s run, which operates in two stages. The first stage inputs a time-series of progress scores of a given map/reduce task and outputs a time-series of estimated completion times corresponding to multiple epochs in its life time. The second stage uses a time-series of past resource usage/allocations across the task’s run-time starting from its start till the current time of estimation. The estimation model then outputs the estimated resource allocation for the next epoch as a function of a task’s cumulative run-time. Separate estimation models are built for CPU, memory and I/O resource profiles of a task. We use the linear regression model for CPU, and piece-wise linear regression model for memory, as derived from our previous work [31]. In this paper, we also estimate the I/O interference through an exponential regression model, similar to [13]. For controlling the I/O bandwidth allocation to an individual task, we use the recently released Linux kernel support for I/O throttling using cgroups [4]. Complete details about the construction of models, along with the specifics of dynamic resource allocation mechanism, are detailed in [31].

2) *Interference Prevention System (IPS)*: For the transactional workloads, collocated with the MapReduce jobs in the virtualized environment, we build separate performance interference models, because their application characteristics are different from MapReduce jobs. We explored building both linear regression model as well as non-linear exponential regression model to quantify the CPU interference. However, we stick to a linear regression model as it fits well in our case in terms of offered simplicity and accuracy (refer [31]). For memory, a piece-wise linear regression model well captures the interference [31]. For I/O, we use a non-linear exponential regression based interference model. Specific details of the

Algorithm 3 Arbitration Algorithm.

```
1: Obtain the list of map/reduce tasks,  $TASK\_LIST_{interference}$ ,  
   interfering with collocated interactive jobs using Estimator  
2: Evaluate the estimated interference for each task in  
    $TASK\_LIST_{interference}$  using DRM's prediction model  
   TaskInterference[]=GetEstimatedInterference()  
   # Get list of all available VMs in  $V\_CLUSTER$   
3: AvailableVMs[] = GetAllAvailableVMs()  
4: while InterferenceTasks[]  $\neq$  NULL do  
5:   Get best available VM using BestFit bin-packing heuristic [12]  
6:    $VM_{best}$ =GetVMByBestFit()  
7:   Use Min-Min heuristics [13] to schedule the 'least-interfering'  
   task on  $VM_{best}$   
8:   Adjust the resources of  $VM_{best}$   
9: end while
```

regression models used are similar to [11].

The IPS uses these performance interference models for interactive applications to dynamically orchestrate their resource allocations in the virtual cluster. The IPS continuously tracks the performance of interactive jobs. If at some instant, the performance falls outside allowable SLA bounds, the IPS invokes its Arbiter module to determine the specific task(s) and resource(s) that are causing this performance degradation. The Arbiter employs its arbitration scheme (Algorithm 3) to mitigate the interference, and restores the performance of the interactive applications. The Arbiter takes into consideration the interference from other collocated VMs, co-scheduled MapReduce jobs and packing of VMs, while deciding on the exact migration of VMs to other hosts, or reassignment of interfering tasks to other VMs. Algorithm 3 describes the details of the arbitration scheme.

IV. EXPERIMENTAL ANALYSIS

Infrastructure: Our experimental testbed contains a mix of physical and virtual nodes constituting native and virtual clusters, respectively. The native cluster contains 24 physical nodes. Each node has a dual 64-bit, 2.4 GHz AMD Opteron processor, 4 GB RAM, and Ultra320 SCSI disk drives, connected with 1 Gbps Ethernet. This cluster is installed with Hadoop v0.22.0 with 2 replicas per HDFS block, and 2 map, 2 reduce slots per node. The Hadoop FairScheduler [5] scheduling policy is used. The virtualized cluster contains 24 physical nodes, which are virtualized with Xen v3.4.2 hypervisor [9], and host 2 VMs per PM to create an equivalent 48 nodes virtual cluster. Each VM runs RedHat AS4 with kernel v2.6.18, and the VMs are connected by 1 Gbps Ethernet. Each VM is configured with 1 GB RAM and 1 virtual CPU.

Benchmarks: We use a workload mix consisting of both transactional applications as well as batch workloads. For transactional workloads, we use three representative applications: (i) RUBiS [28], which models an online auction site; (ii) TPC-W [32], which models a three-tier online book store; and (iii) Olio [26], which is a Web 2.0 social events application. For batch workloads, we use the following representative MapReduce jobs:

- **Twitter:** uses 25 GB twitter data [8] to rank users (Memory + I/O-bound).
- **Wcount:** computes frequencies of words in 20 GB of text data (Memory + I/O-bound).
- **PiEst:** estimates π using 10 million points (CPU-bound).
- **DistGrep:** finds match of randomly chosen regular expressions on 20 GB of text data (I/O-bound).

- **Sort:** sorts 20 GB of text data (I/O-bound).
- **Kmeans:** clusters 10 GB of numeric data (CPU-bound).

The above MapReduce jobs are chosen based on their popularity and being representative of real MapReduce benchmarks, with diverse resource mix. We use the Yokogawa's WT210 power meter to measure the server power. For better precision, we run all our experiments 3 times, and use the average value. Unless otherwise specified, all experiments are performed in the hybrid cluster consisting of 24 PMs and 48 VMs.

Performance Benefits of HybridMR: We begin with demonstrating the performance benefits of *HybridMR*'s Phase I Scheduler in the initial placement of MapReduce jobs. Towards this, we conduct an experiment, where we measure the average response times of interactive applications and JCTs of MapReduce jobs, when scheduled by the Phase I Scheduler, normalized against the corresponding JCTs measured with random placement of jobs (i.e., first-come-first-served discipline). In Figure 8(a), we plot this normalized JCT, labeled as 'Performance Gain'. We observe that through efficient placement by the Phase I Scheduler, both interactive and batch jobs benefit in performance, and the magnitude of gain varies depending on the workload mix characteristics. We consider 3 workload mixes: *wmix-1*, *wmix-2* and *wmix-3*, where each represents 50% interactive + 50% batch jobs, 20% interactive + 80% batch jobs, and 80% interactive + 20% batch jobs, respectively. Intuitively, this performance gain is achievable due to the fact that certain MapReduce jobs, which are resource intensive and have more stringent performance bounds, when placed on native Hadoop cluster reduce the interference across other MapReduce and interactive jobs.

We next analyze the potential benefits of Phase II Scheduler towards improving the performance of MapReduce jobs scheduled on the virtual cluster. In this experiment, we use 48 VMs virtual Hadoop cluster. We compare the performance of each of the six MapReduce jobs when run standalone (*Single*). We measure the JCT when Phase II Scheduler-based resource orchestration is enabled ($JCT_{HybridMR}$) versus without it ($JCT_{default}$). We compute percentage reduction in JCT as $((JCT_{default} - JCT_{HybridMR})/JCT_{default}) * 100$. Figure 8(b) shows the percentage reduction in JCT of the six MapReduce jobs. Here, each legend corresponds to the case when only CPU, Memory, I/O or all three of them are managed by *HybridMR*. Across all jobs, *CPU+Memory+I/O* yields the maximum reduction in JCTs, while the reduction with *CPU*, *Memory* or *I/O* alone varies in accordance with the benchmark resource characteristics. Specifically, we observe an average and a maximum JCT reduction of 22% and 29.1%, respectively, with the *CPU+Memory+I/O* mode. Note that larger jobs (with respect to both large input size and long running job) like *Sort* and *Kmeans* benefit more when compared to other relatively shorter jobs like *PiEst* and *DistGrep*. This is due to the fact that larger jobs run in multiple map/reduce phases, and thus, *HybridMR* has more opportunities to dynamically coordinate resources and improve their JCTs. Similar benefits are observed in the JCT of each job when run in the presence of other 5 concurrent running MapReduce jobs (*Multiple jobs*). Figure 8(c) shows the results. Note that since *Multiple jobs* scenario induces more interference than corresponding *Single job*, higher performance benefits are observed with *Multiple jobs*. We observe an average and a maximum reduction of 28.5% and 40.8% in the JCTs with

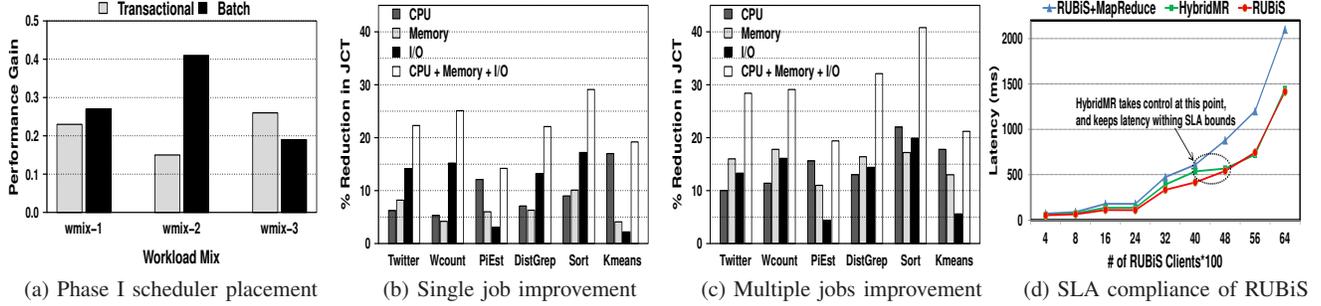


Figure 8: Performance benefits of *HybridMR* on a virtualized platform.

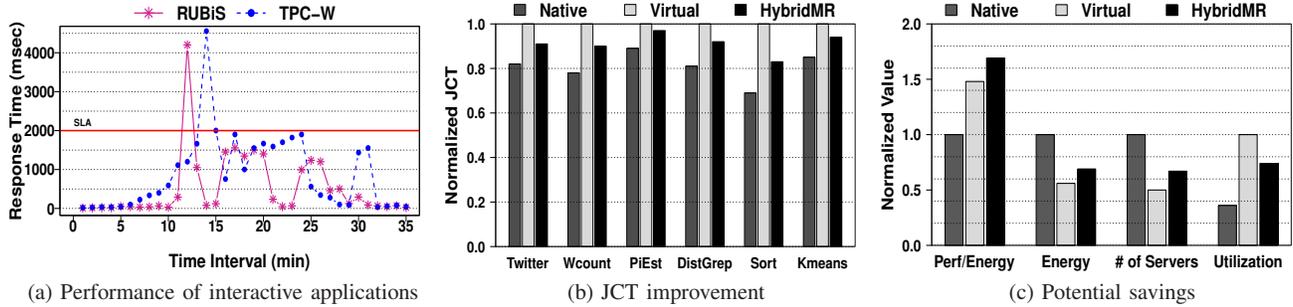


Figure 9: *HybridMR* achieves better balance between Native and Virtual. Y-axis in (b) and (c) is normalized w.r.t. max. value.

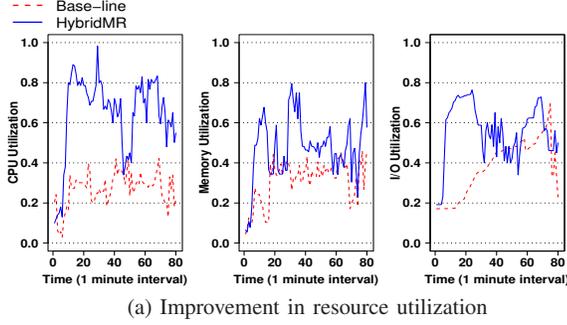
CPU+Memory+I/O. In Figure 10(a), we also observe high CPU, Memory and I/O utilization achieved by *HybridMR*.

Next, we analyze the impact of *HybridMR* on the SLAs of interactive applications like RUBiS. We perform an experiment, where we increase the number of RUBiS clients and observe the impact on the latency of web server. Figure 8(d) shows the results. Here, the top *RUBiS+MapReduce* curve corresponds to the case, where the RUBiS virtualized cluster co-hosts other MapReduce jobs assigned in the FIFO order by the default MapReduce scheduler. The bottom *RUBiS* curve denotes the case, where RUBiS runs in isolation with no MapReduce jobs. The middle *HybridMR* curve depicts the scenario with the *HybridMR* scheduler. Here, we observe that for lower client workloads, *HybridMR* is able to co-host MapReduce jobs with RUBiS and maintain its latency within bounds. However, when *HybridMR* detects that client latency is exceeding a given threshold, its Phase II Scheduler can adaptively migrate or pause the execution of co-hosted interfering MapReduce tasks to bring down the latency similar to the bottom *RUBiS* curve. This shows the dynamic adaptation of *HybridMR* in keeping up with the SLA of interactive applications.

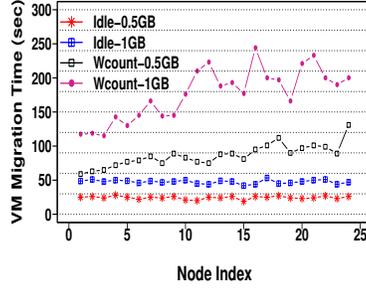
In a related analysis, we demonstrate the effectiveness of *HybridMR* in maintaining the SLAs of interactive RUBiS and TPC-W applications, when run collocated with other MapReduce jobs. In Figure 9(a), we observe that in the time interval 1-10 minutes, the response times of both RUBiS and TPC-W are below the user defined SLA (2 sec. in our experiment). However, around 12th (RUBiS) and 14th (TPC-W) time instants, the response time exceeds the SLA. *HybridMR* quickly identifies this, and is able to migrate the collocated map/reduce tasks from the VMs running RUBiS and TPC-W to mitigate the interference. Consequently, the response time falls within limits again. Similar behavior follows for the Olio benchmark, but we skip the analysis in the interest of space.

Cross-Platform Performance Comparison: We next demonstrate the benefits of *HybridMR* in scheduling a mix of interactive and batch jobs. In this experiment, we evaluate three design choices for cluster configuration. In the first design choice (*Native*), the workload mix is scheduled on a 24-node native Hadoop cluster. In the second design choice (*Virtual*), the workload mix is scheduled on a 24-node virtual cluster (hosted on 12 PMs, each PM contain 2 VMs). The third design choice (*HybridMR*) corresponds to the case, where the workload mix is scheduled on a 24-node hybrid cluster, consisting of 12 PMs + 12 VMs (consolidated on 6 PMs with 2 VMs each). Thus, *Native*, *Virtual* and *HybridMR* require 24, 12, and 18 number of physical servers, respectively. From Figure 9(b), we observe that the performance of MapReduce jobs is the worse for *Virtual* because of the possible virtualization overheads. *Native* achieves the best performance as expected. *HybridMR*'s performance is intermediate between *Native* and *Virtual*. However, as we can observe in Figure 9(c), *HybridMR* achieves the highest *Performance/Energy*, which is a very important data center design metric. In terms of other metrics like energy, utilization, and total number of physical servers, *HybridMR* strikes a better balance between the two extreme choices – *Native* and *Virtual* design alternatives.

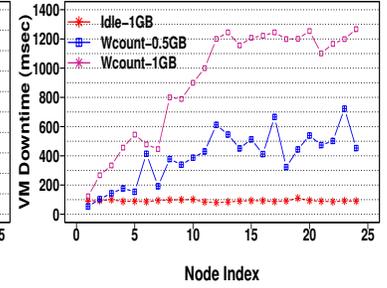
Impact of live migration of Hadoop VMs: In the IPS module (Section III-B2), we leverage live migration of VMs as a mechanism to redistribute interfering map/reduce tasks. To understand the performance implications of Hadoop VMs migration, we perform some empirical analyses. We run *Wcount* on different data size on a 24 VMs Hadoop cluster, and migrate randomly the 24 VMs during the execution. Figure 10(b) and Figure 10(c) demonstrate the migration time and downtime of each VM during migration. We observed: (i) larger the amount of memory involved, longer is the migration time, while the downtime dependence on memory is ad-hoc; (ii) migration



(a) Improvement in resource utilization



(b) Impact of VM migration



(c) Impact of VM downtime

Figure 10: In (a), *HybridMR* boosts resource utilization; (b) and (c) show performance impact of live migration of Hadoop VMs.

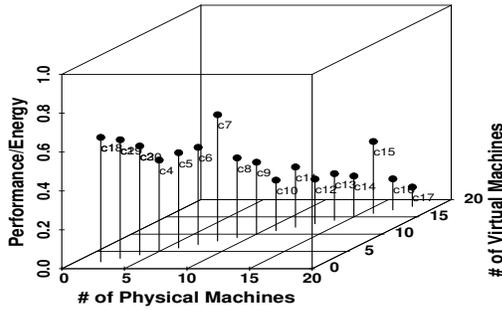


Figure 11: Hybrid configuration design trade-off analysis.

time of VMs running *Wcount* is relatively longer than the corresponding idle Hadoop cluster; and (iii) downtime of each VM running *Wcount* shows wide variation, possibly due to skewness of each Hadoop VM, due to difference in resident data blocks, concurrent map/reduce tasks, and interference from other colocated VMs. From this analysis, we conclude that live migration of Hadoop VMs incurs some overheads, especially the downtime. However, due to the inherent fault-tolerant characteristics of Hadoop (*i.e.*, replication), the downtime period can be adjusted by regenerating the data blocks from other available copies. However, the MapReduce jobs still run to finish in the face of the migration induced downtime.

Design Trade-off Analysis: We next discuss the flexibility of *HybridMR* in terms of splitting a given physical infrastructure into different hybrid cluster configurations, each consisting of a fixed number of VMs and PMs for running a workload mix to meet different performance and energy requirements. To highlight this empirically, we split our testbed (consisting of 24 PMs hosting 48 VMs) into 20 different cluster configurations ($C_1 - C_{20}$), where each C_i corresponds to randomly selected PMs and VMs, and runs workload mix of interactive and MapReduce jobs. For each configuration, we measure the average completion times (performance), average energy consumption and average utilization across all jobs and nodes. In Figure 11, each configuration corresponds to a particular $\langle \# \text{ of PMs}, \# \text{ of VMs}, \text{Performance/Energy} \rangle$ tuple in the 3D space. For example, configuration C_7 (with 12 PMs, 12 VMs) and C_{17} (with 24 PMs, no VMs) gives the best and worst *Performance/Energy*, respectively. Such analysis can provide some useful insights. For example, a cluster administrator may

use some hints from a similar empirical analysis for deciding on how to partition the infrastructure into hybrid cluster to meet a desired performance-energy envelope.

V. RELATED WORK

The prior work related to this paper is summarized under the following three categories:

Resource Scheduling in Hadoop MapReduce: Scheduling techniques for dynamic resource management of MapReduce jobs have been recently addressed [23], [27], [29], [31], [33]. Two popular recent resource managers for Hadoop MapReduce are Mesos [19] and YARN [3]. Most of the works in this category primarily focus on different scheduling policies for MapReduce jobs to reduce their completion times, improve cluster resource utilization and energy efficiency.

MapReduce and virtualization: There has been a recent thrust to improve the performance of MapReduce on virtualized platforms. Amazon Elastic MapReduce [1] is a publicly offered web service that runs on virtualized cloud platform. Serengeti [6] is an open source project initiated by VMware to enable rapid deployment of Hadoop on a virtual platform. Cardosa et al. [12] proposed techniques for MapReduce provisioning in virtualized cloud clusters, with emphasis on energy reduction. Managing MapReduce workloads using Amazon virtual spot instances is studied in [14]. Resource allocation in Hadoop MapReduce using dynamic prioritization is proposed in [29]. Preliminary evaluations of Hadoop MapReduce performance on virtualized clusters is recently done in [20], [35]. Virtual machine scheduling heuristic to improve Xen scheduler, targeting MapReduce workloads, is addressed in [21]. Harnessing unused CPU cycles in interactive clouds for batch MapReduce workloads has been recently motivated in [15]. Bu et al. [11] proposed an interference-aware and locality-aware scheduling algorithm for optimizing MapReduce in virtualized environment. VMware DRS [34] manages VMs through dynamic load balancing.

Interference based resource management in clouds: With the advent of cloud computing, resource management in virtualized clouds has emerged as an important research avenue. There exists quite a few literatures in this context. For example, Q-Clouds [25] leverages online feedback control technique to dynamically manage the resource allocation to VMs. TRACON [13] is an interference-aware scheduling algorithm for data-intensive applications in virtual environment. Automatic resource provisioning in MapReduce cloud clusters is explored in [12]. Younggyun et al. [36] presents an empirical study on the performance interference effects in virtual environments.

We share our motivation of hybrid compute clusters with [15], [22], [24]. We believe our work differs from others in the following manner. First, a detailed empirical evaluation and performance analysis study of Hadoop MapReduce on virtual environment have not been addressed in prior literature. Second, scheduling of heterogeneous workloads (mix of transactional and batch MapReduce jobs) on a hybrid cluster (consisting of native and virtual environments) to exploit the spare resources available due to over-provisioning of interactive applications, has not been explored before. Third, our proposed scheduler, *HybridMR*, uniquely focuses on the performance enhancement of MapReduce jobs, collocated with interactive jobs in a virtual environment while complying with the SLAs of the foreground jobs. Fourth, no previous studies have paid much attention to the benefits and design trade-offs of hybrid data centers consisting of both native and virtual servers, and hosting heterogeneous workload mix.

VI. CONCLUSIONS

This paper presents a 2-phase hierarchical scheduler, called *HybridMR*, for hybrid data centers, consisting of a mix of native and virtual machines to leverage the benefits of both paradigms. In the first phase, *HybridMR* profiles incoming MapReduce jobs to gauge the estimated virtualization overheads, and utilizes this information to automatically guide placement between physical machines and virtual machines. In the second phase, *HybridMR* builds run-time resource prediction models, and performs dynamic resource orchestration to minimize the interference within and across collocated interactive and MapReduce applications. Detailed evaluations on a hybrid compute cluster consisting of 24 physical servers and 48 virtual machines, with a diverse workload mix of interactive and batch MapReduce benchmarks, demonstrate that *HybridMR* achieves up to 40% improvement in job completion times of MapReduce jobs over a virtual cluster, while satisfying the SLAs of interactive applications. Further, *HybridMR* provides 45% improvement in resource utilization and around 43% energy savings compared to a native cluster, with minimal performance penalty. Furthermore, we empirically demonstrate that it is possible to dynamically change the native and virtual cluster configurations to accommodate variations in workload mix for maximizing the performance-energy envelope. These results establish that a hybrid cluster configuration is a better cost-effective solution than the either-or approach to native and virtual modes of computation.

As part of the future work, we plan to investigate the applicability of *HybridMR* with other Hadoop implementations, including iterative MapReduce (Twister [17]) and in-memory version (Spark [37]). We are also pursuing the the dual problem of scheduling interactive jobs on virtual and physical machines in a heterogeneous compute environment [30].

ACKNOWLEDGMENT

We thank the anonymous reviewers, Adwait Jog, Amin Jadidi, Mahshid Sedghi, Nachiappan Chidambaram Nachiappan, and Onur Kayiran for their valuable comments towards improving this paper. This research is supported in part by NSF grants #1213052, #1205618, #1152479, #1147388, #0702617 and research grants from Google and Intel.

REFERENCES

- [1] Amazon MapReduce. <http://aws.amazon.com/elasticmapreduce/>.
- [2] Apache Hadoop. <http://hadoop.apache.org>.

- [3] Apache Hadoop YARN. <http://tinyurl.com/hadoop-yarn>.
- [4] Cgroup Support for I/O Throttling. <http://tinyurl.com/io-throttle>.
- [5] Hadoop Fair Scheduler. <http://tinyurl.com/scheduler-fair>.
- [6] Hadoop on VMware. <http://serengeti.cloudfoundry.com/>.
- [7] No Virtualization for Clouds. <http://tinyurl.com/novirtualization>.
- [8] Twitter Traces. <http://an.kaist.ac.kr/traces/WWW2010.html>.
- [9] Xen Hypervisor. <http://www.xen.org>.
- [10] P. Barham, B. Dragovic, and et al. "Xen and the Art of Virtualization". *SIGOPS Oper. Syst. Rev.*, 37(5):164–177, 2003.
- [11] X. Bu, C. Xu, and J. Rao. "Interference and Locality-Aware Scheduling for MapReduce Service in Virtual Clusters". Technical report, University of Colorado, 2012.
- [12] M. Cardosa, P. Narang, A. Chandra, H. Pucha, and A. Singh. "STEAM-Engine: Driving MapReduce Provisioning in the Cloud". In *HIPC*, 2011.
- [13] R. C. Chiang and H. H. Huang. "TRACON: Interference-aware Scheduling for Data-intensive Applications in Virtualized Environments". In *SC*, 2011.
- [14] N. Chohan, C. Castillo, and et al. "See Spot Run: Using Spot Instances for Mapreduce Workflows". In *HotCloud*, 2010.
- [15] Ross Benjamin Clay. "Enabling MapReduce to Harness Idle Cycles in Interactive-User Clouds". Master's thesis, NC State University, 2011.
- [16] J. Dean and S. Ghemawat. "MapReduce: Simplified Data Processing on Large Clusters". *Communications of the ACM*, 51(1):107–113, 2008.
- [17] J. Ekanayake, H. Li, and et al. "Twister: A Runtime for Iterative MapReduce". In *HPDC*, 2010.
- [18] H. Herodotou and S. Babu. "Profiling, What-if Analysis, and Cost-based Optimization of MapReduce Programs". *PVLDB*, 4(11):1111–1122, 2011.
- [19] B. Hindman, A. Konwinski, and et al. "Mesos: A Platform for Fine-grained Resource Sharing in the Data Center". In *NSDI*, 2011.
- [20] S. Ibrahim, H. Jin, and et al. "Evaluating MapReduce on Virtual Machines: The Hadoop Case". In *CloudCom*, 2009.
- [21] H. Kang, Y. Chen, J. L. Wong, R. Sion, and J. Wu. "Enhancement of Xen's Scheduler for MapReduce Workloads". In *HPDC*, 2011.
- [22] T. Kooburat and M. Swift. "The Best of Both Worlds with On-demand Virtualization". In *HotOS*, 2011.
- [23] G. Lee, B.G. Chun, and R. Katz. "Heterogeneity-Aware Resource Allocation and Scheduling in the Cloud". In *HotCloud*, 2011.
- [24] H. Lin, X. Ma, and et al. "MOON: MapReduce On Opportunistic eNvironments". In *HPDC*, 2010.
- [25] R. Nathuji, A. Kansal, and A. Ghaffarkhah. "Q-clouds: Managing Performance Interference Effects for QoS-aware Clouds". In *EuroSys*, 2010.
- [26] Olio. Web 2.0 Social Events Application. <http://incubator.apache.org/olio>.
- [27] J. Polo, C. Castillo, and et al. "Resource-aware Adaptive Scheduling for Mapreduce Clusters". In *Middleware*, 2011.
- [28] RUBiS. E-commerce Application. <http://rubis.ow2.org>.
- [29] T. Sandholm and K. Lai. "MapReduce Optimization using Regulated Dynamic Prioritization". In *SIGMETRICS*, 2009.
- [30] B. Sharma, V. Chudnovsky, J. L. Hellerstein, R. Rifaat, and C. R. Das. "Modeling and Synthesizing Task Placement Constraints in Google Compute Clusters". In *SOCC*, 2011.
- [31] B. Sharma, R. Prabhakar, S. Lim, M. T. Kandemir, and C. R. Das. "MROrchestrator: A Fine-Grained Resource Orchestration Framework for MapReduce Clusters". In *CLOUD*, 2012.
- [32] TPC-W. E-commerce Benchmarking. <http://www.tpc.org/tpcw/>.
- [33] A. Verma, L. Cherkasova, and R. H. Campbell. "ARIA: Automatic Resource Inference and Allocation for Mapreduce Environments". In *ICAC*, 2011.
- [34] VMware. Distributed Resource Scheduler. <http://www.vmware.com/files/pdf/VMware-Distributed-Resource-Scheduler-DRS-DS-EN.pdf>.
- [35] VMware. Elastic Hadoop for Cloud. <http://tinyurl.com/elastic-hadoop>.
- [36] K. Younggyun, R. Knauerhase, and et al. "An Analysis of Performance Interference Effects in Virtual Environments". In *ISPASS*, 2007.
- [37] M. Zaharia, M. Chowdhury, and et al. "Resilient Distributed Datasets: A Fault-tolerant Abstraction for In-memory Cluster Computing". In *NSDI*, 2011.