

Mitigating Bottlenecks in Wide Area Data Analytics via Machine Learning

Hao Wang and Baochun Li, *Fellow, IEEE*

Department of Electrical and Computer Engineering, University of Toronto

Abstract—Over the past decade, we have witnessed exponential growth in the density (petabyte-level) and breadth (across geo-distributed datacenters) of data distribution. It becomes increasingly challenging but imperative to minimize the response times of data analytic queries over multiple geo-distributed datacenters. However, existing scheduling-based solutions have largely been motivated by pre-established mantras (e.g., bandwidth scarcity). Without data-driven insights into performance bottlenecks *at runtime*, schedulers might blindly assign tasks to workers that are suffering from unidentified bottlenecks.

In this paper, we present *Lube*, a system framework that minimizes query response times by detecting and mitigating bottlenecks at runtime. *Lube* monitors geo-distributed data analytic queries in real-time, detects potential bottlenecks, and mitigates them with a bottleneck-aware scheduling policy. Our preliminary experiments on a real-world prototype across Amazon EC2 regions have shown that *Lube* can detect bottlenecks with over 90% accuracy, and reduce the median query response time by up to 33% compared to Spark’s built-in locality-based scheduler.

Index Terms—Wide area, data analytics, performance prediction, machine learning, bottleneck detection, task scheduling.

1 INTRODUCTION

WITH large volumes of data generated and stored at geographically distributed datacenters around the world, it has become increasingly common for large-scale data analytics frameworks, such as Apache Spark [47] and Hadoop [15] to span across multiple datacenters. Petabytes of data — including user activities, trending topics, service logs and performance traces — are produced on these geographically distributed datacenters every day, processed by tens of thousands data analytic queries.

Minimizing response times of geo-distributed data analytic queries is crucial, but far from trivial. Results of these analytics queries are typically used when making real-time decisions and online predictions, all of which depend upon the timeliness of data analytics. However, in contrast to data analytics in a single datacenter, the varying bandwidth on wide-area network (WAN) links and the heterogeneity of the runtime environment across geographically distributed datacenters impose new and unique challenges as query response times are minimized.

Known as *wide-area data analytics* in the literature, tasks (or data) are optimally placed across datacenters in order to improve data locality [26, 27, 36, 42, 43]. However, all previous works made the simplifying assumption that the runtime environment of wide-area data analytics is temporally stable, and that there are no runtime performance variations in these clusters. Naturally, this may not accurately reflect the reality. In addition, existing works have largely been motivated by a few widely accepted mantras, such as the scarcity of network bandwidth on access links from a datacenter to the Internet. With an extensive measurement study on analytic jobs, Ousterhout *et al.* [32] have convincingly pointed out that some of the widely held assumptions in the literature may not be valid in the context of a single cluster.

Delving into the fluctuating runtime environment of wide-area data analytics, this paper makes a strong case for analyzing and detecting performance bottlenecks in data analytics frameworks

at runtime. Shifting gears from a single cluster to the context of wide-area data analytics, we believe that the conclusion from [32] still holds: it may not always be the same resource — such as bandwidth — that causes runtime performance bottlenecks in wide-area data analytic queries. To generalize a step further, the types of resource that cause performance bottlenecks may even vary over time at runtime, as analytic queries are executed across datacenters. It becomes intuitive that, if we wish to reduce the query response times in wide-area data analytics, these performance bottlenecks need to be detected *at runtime*, and a new resource scheduling mechanism needs to be designed to mitigate them. Unfortunately, such a high-level intuition has not yet been well explored in the literature and remains a largely uncharted territory.

In this paper, we propose *Lube*, a new system that is designed to perform data-driven runtime performance analysis for minimizing query response times. *Lube* features a closed-loop design: the results of runtime monitoring are used for detecting bottlenecks, and these bottlenecks serve as input to the resource scheduling policy to mitigate them, again at runtime. Our original contributions in this paper are the following:

First, we propose effective and efficient techniques to detect resource bottlenecks at runtime. We investigate two bottleneck detection techniques, both driven by performance metrics collected in real-time. We start with a simple statistical technique, Autoregressive Integrated Moving Average (ARIMA) [10], and then propose machine learning techniques to further explore the implicit correlation between multiple performance metrics.¹ As one of the effective algorithms and a case study, we use the Sliding Hidden Markov Model (SlidHMM) [12], an unsupervised algorithm that takes time series as input and incrementally updates

1. For example, a higher network I/O will lead to higher JVM heap swap frequencies, since network send/receive semantics will trigger memory load/dump operations.




Fig. 1: A potential bottleneck in memory.

model parameters for detecting upcoming states.

Second, we propose a new scheduling policy that, when assigning tasks to worker nodes, mitigates bottlenecks by considering not only data locality (e.g., [42]), but also the severity of bottlenecks. The upshot of our new scheduling policy is the use of a technique similar to late binding in Sparrow [33], that holds a task for a short while before binding it to a worker node. This is designed to avoid the negative implications of false positives when detecting bottlenecks.

We have implemented a prototype of Lube on a Spark SQL cluster over Amazon EC2 with 37 instances across nine regions. Our experiments of the Big Data Benchmark [39] with a 1.1 TB dataset show that Lube is able to detect bottlenecks with an accuracy over 90% and reduces the median query response time by as much as 33% ($1.5\times$ faster).

2 LUBE: A BIRD’S-EYE VIEW

Data analytics over geo-distributed datacenters may suffer from a highly volatile runtime environment, due to the lack of load distribution when using resources, or varying bandwidth availability over wide-area network links [4]. As a result, resource bottlenecks are more likely to occur at runtime, when data analytic queries are executed over the wide area.

As a motivating example of such runtime bottlenecks, Figure 1 presents a heat map of real-time memory utilization on the Java Virtual Machine (JVM) heap, captured on a 5-node Spark SQL [7] cluster running the Big Data Benchmark [39]. As we can observe, within a specific time window (marked by $t_{current}$), memory is heavily utilized on node_1, while other nodes are largely idle on their memory utilization. This implies that memory becomes a bottleneck on node_1, because the Spark SQL scheduler assigned more tasks to this node with no knowledge that its memory may be overloaded at runtime.

Given the existence of resource bottlenecks, our ultimate objective is to reduce query response times by designing new task scheduling strategies that work around these bottlenecks. To achieve such an objective, we need to monitor performance attributes of data analytic queries at runtime and detect potential bottlenecks with very little overhead. To be more specific, we will need to design and implement the following components:

Lightweight performance monitors. A collection of performance monitors on each worker node is needed to capture process-level performance metrics in real-time. In Lube, rather than intrusively using code instrumentation, we choose to reuse existing lightweight system-level performance monitors on Linux (e.g., `jvmtop`, `iotop`, `iperf` and `nethogs`). **Online bottleneck detection.** With performance metrics collected in real-time, we will propose algorithms that analyze dependencies between performance metrics and detect potential bottlenecks at runtime.

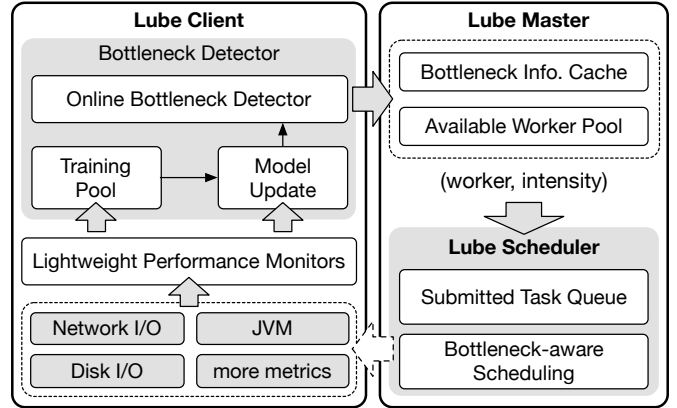


Fig. 2: *Lube*: a closed-loop architecture involving performance monitors, bottleneck detection, and task scheduling.

Bottleneck-aware scheduling. To react to detected bottlenecks, a bottleneck-aware scheduler will make task assignment decisions by considering both bottleneck severities and data locality. Besides, the scheduler should be able to tolerate inaccurate detections.

Figure 2 presents the closed-loop design architecture of Lube. On each worker node, a Lube client periodically collects runtime performance metrics, updates the machine learning model and reports detected bottlenecks to the Lube master; on the master node, the task scheduler makes task assignment decisions based on bottleneck intensities at the worker nodes, as well as data locality preferences of tasks. In return, the decisions made by the task scheduler will further influence the performance of data analytic queries at each worker node.

3 DETECTING BOTTLENECKS

Performance bottlenecks may emerge anytime and anywhere in wide-area data analytics. To mitigate performance bottlenecks in time, we will first need to detect them correctly at runtime. Lube performs online bottleneck detection on performance metrics collected in real-time.

We investigate two techniques to detect bottlenecks from the time series of performance metrics. One is a simple statistical model — the Autoregressive Integrated Moving Average (ARIMA) algorithm that approximates the future value by a linear function of past values and past errors; the other is an unsupervised machine learning model: the Sliding Hidden Markov Model (SlidHMM) algorithm that can autonomously learn the implicit correlation between multiple performance metrics.

3.1 Dataset Collection

We collect runtime performance metrics from both system-level monitors and monitoring interfaces of data analytic frameworks. Table 1 summarizes all features we have considered in the bottleneck detection. It should be noted that all features are obtained in a non-intrusive way by reusing existing toolkits and interfaces. Besides, it is easy to incrementally expand the training dataset by including performance metrics collected during the execution of queries in reality. The bottleneck detection models can be improved with the updated dataset. All metrics are recorded as a 2-tuple including the timestamp and the value. A master node

| Metrics | Description |
|---------------------|------------------------------------|
| disk_io | Disk read/write throughput |
| disk_iops | Disk I/O operations per second |
| net_io | Network throughput |
| cpu_util | CPU utilization per core |
| mem_util | Memory utilization |
| jvm_gc_util | JVM garbage collection utilization |
| jvm_heap_util | JVM heap utilization |
| pending_task_num | The number of waiting tasks |
| completed_task_num | The number of completed tasks |
| failed_task_num | The number of failed tasks |
| task_retry_num | The retry times of failed tasks |
| pending_stage_num | The number of waiting stages |
| completed_stage_num | The number of completed stages |

TABLE 1: The metrics used to detect the potential performance bottlenecks, which are collected by performance monitors of Linux and data analytic frameworks.

aggregates metrics from each worker node and updates the dataset on each worker node periodically.

In addition to the performance metrics, we also include the progress metrics collected from the data analytic frameworks. The progress metrics indicates the runtime execution status of workers. For example, a high `task_retry_num` on a worker implies that the worker is in severe bottleneck. Zero `pending_task_num` implies the worker is available or becoming available.

We perform data preprocessing on the raw data of performance metrics, because the range of values of features varies widely. For performance metrics such as `disk_io` and `net_io`, they are normalized by dividing the theoretical maximum performances defined by the hardwares. For utilization metrics such as `cpu_util` and `mem_util`, they are inherently normalized values between 0 and 1. Figure 3 shows heap maps of the normalized metrics.

3.2 ARIMA

Introduced by Box and Jenkins [10], the ARIMA model has been widely applied in time series analysis. As a combination of autoregressive (AR) model and the moving average (MA) model, the ARIMA model is defined by the following equation:

$$y_t = \theta_0 + \phi_1 y_{t-1} + \phi_2 y_{t-2} + \dots + \phi_p y_{t-p} + \epsilon_t - \theta_1 \epsilon_{t-1} - \theta_2 \epsilon_{t-2} - \dots - \theta_q \epsilon_{t-q}, \quad (1)$$

where y_t and ϵ_t denote the actual value and random error at time t respectively; ϕ_i ($i = 1, 2, \dots, p$) and θ_j ($j = 1, 2, \dots, q$) are the coefficients specified by the model. p and q are integers indicating the autoregressive (AR) and moving average (MA) polynomials respectively. For example, if $q = 0$, then Eq. (1) reduces to an AR model of order p ; if $p = 0$, then Eq. (1) reduces to a MA model of order q . A general ARIMA model is represented as $ARIMA(p, d, q)$, in which d is the degree of difference transformation for data stationarity.

The ARIMA methodology basically has three iterative steps. Firstly, we test the stationarity of the input time series by the Dickey-Fuller Test [17]. If the time series is not stationary, we transform it into a stationary time series by applying a suitable degree (defined by d) of differencing. Secondly, we examine the autocorrelation function (ACF) and the partial autocorrelation function (PACF) of the input time series to estimate appropriate p and q [10]. Thirdly, we check the residuals (actual values minus

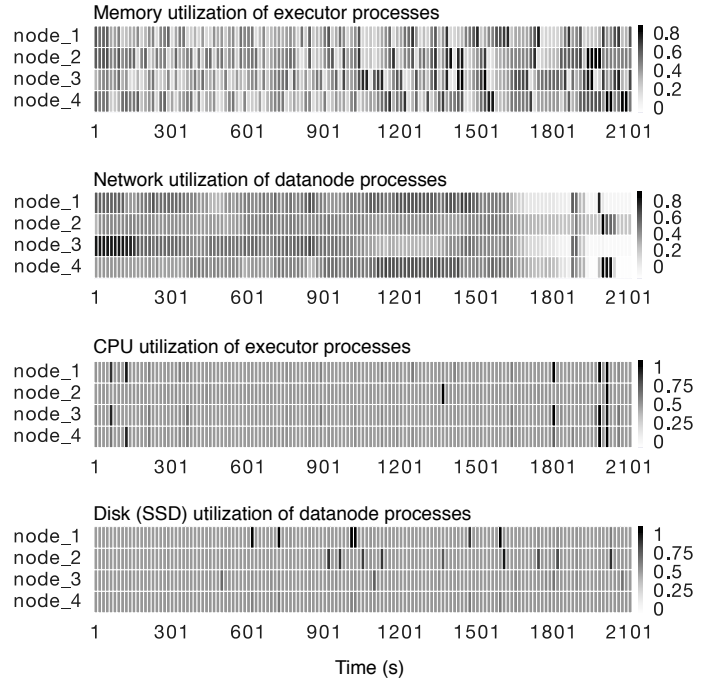


Fig. 3: Heat maps of performance metrics.

fitted values) to diagnose whether the model is adequate. If the model is not adequate, then we repeat the previous steps.

We build an univariate ARIMA model for each performance metric, rather than a vector ARIMA model for all metrics, as it usually becomes “overfitting” due to too many combinations of insignificant parameters [14]. To support online bottleneck detection, we periodically update the ARIMA model with continuously arriving performance metrics.

3.3 Sliding HMM

The Hidden Markov Model (HMM) [8] infers a sequence of hidden states that maps to the sequence of observation states. Through feeding a time series of observed performance metrics (O_1 to O_d) to HMM, we can infer the possible performance metrics O_k in the future (Figure 4). The HMM is usually defined as a three-tuple: (A, B, π) as the following notations (t is the time stamp):

$Q = \{q_1, q_2, \dots, q_N\}$, hidden state sequence.

$O = \{O_1, O_2, \dots, O_k\}$, observation state sequence.

$A = \{a_{ij}\}$, $a_{ij} = \Pr(q_j \text{ at } t+1 | q_i \text{ at } t)$, transition matrix.

$B = \{b_j(k)\}$, $b_j(k) = \Pr(O_k \text{ at } t | q_j \text{ at } t)$, emission matrix.

$\pi = \{\pi_i\}$, $\pi_i = \Pr(q_i \text{ at } t = 1)$, initial state distribution.

HMM learns the hidden states based on an expectation maximization algorithm, the Baum-Welch Algorithm [9]. This algorithm iteratively searches the model parameters (A, B, π) that maximizes the likelihood of $\Pr(O|\mu)$ — the best explanation of the observation sequence. Traces of JVM heap utilization in Figure 5 presents a clear periodical pattern. By learning the hidden states behind this pattern, the HMM infers the future performance metrics for bottleneck-aware scheduling.

To support bottleneck detection in runtime, HMM must be updated online. However, such online updates incur a heavy cost

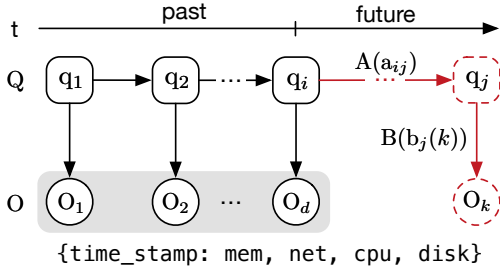


Fig. 4: The Hidden Markov Model.

in both time and space, as the Baum-Welch algorithm needs to re-calculate both old and new time series input. Hence, we propose to use the Sliding Hidden Markov Model (SlidHMM) [12], which is a sliding version of the classic HMM, and is particularly designed for online characterization of high-density time series. The core of SlidHMM is a sliding window that accepts new observations and evicts outdated ones. A moving average approximation replaces the outdated observations during SlidHMM’s training phase. Different from the traditional HMM, SlidHMM updates incrementally with the partial calculation on a fixed window of observations. Thus, it improves the efficiency of bottleneck detection in both time and space.

3.4 Accuracy-Overhead Trade-Off

A natural conflict lies between the accuracy of bottleneck detection and the overheads of algorithms. More training iterations lead to more accurate detection, but introduces more overheads and latency. Designed to coexist with costly data analytics engines, bottleneck detection should work under the radar.

Considering the ultimate goal of bottleneck detection is to enable bottleneck-aware scheduling, it is fairly unnecessary to enforce an accuracy of one decimal place. A number of bottleneck severity levels are enough for bottleneck-aware scheduling. To this end, we shape the input performance metrics from continuous values to discrete ones (*e.g.*, integers from 1 to 100) and classify

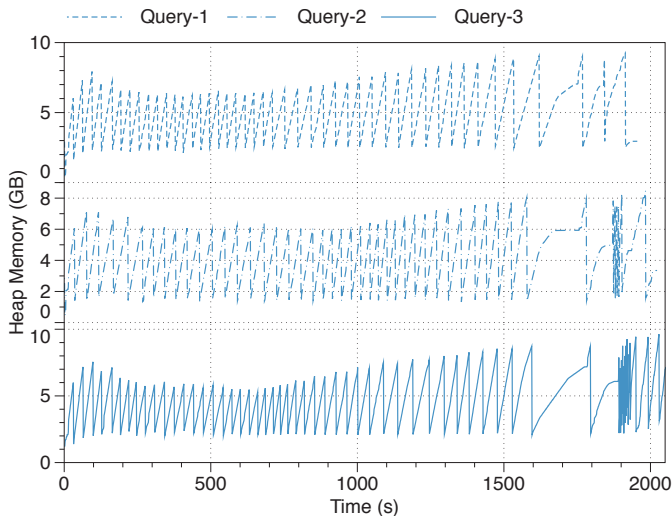


Fig. 5: The JVM heap utilization traces of a Spark executor process.

the output into several bottleneck severity levels. It results in a coarse inference of bottleneck severity, though, which is adequate for our error-tolerant scheduling strategy.

4 BOTTLENECK-AWARE SCHEDULING

To justify the imperatives of bottleneck-aware scheduling, we visualized the performance metrics collected from a Spark SQL cluster running real-world workloads in a geographically distributed fashion on Amazon EC2. Figure 3 reveals the necessity and feasibility of performing bottleneck-aware scheduling in wide-area data analytics: **a single worker node is bottlenecked continuously while all nodes are rarely bottlenecked in chorus.** A bottlenecked node slows down the running tasks, and if we keep assigning tasks to the bottlenecked node, performance will be further degraded. Meanwhile, there usually exist available nodes to take over the tasks assigned to bottlenecked nodes.

Unfortunately, neither existing resource management platforms (*e.g.*, Mesos [23] and YARN [6]) nor scheduling solutions (*e.g.*, Iridium [36] and Sparrow [33]) support online detection of such performance bottlenecks. The built-in schedulers of Spark and Hadoop make decisions only based on data locality, with the objective of reducing network transmission times [46].

Bottlenecked nodes consume extra time to process tasks. To minimize the response times of data analytic queries, we propose a simple task scheduler to coordinate with our bottleneck detection algorithms and mitigate bottlenecks at runtime. A node will be marked as available if no upcoming bottlenecks have been detected; a task has several levels of locality preferences in descending order. When assigning a task, this scheduler jointly considers data locality and bottleneck severity. Essentially, it searches for an available node that satisfies the highest locality preference level of the task compared to all available nodes. Considering that bottlenecks may not be correctly detected, we introduce a *late-binding* algorithm to our bottleneck-aware scheduler. Sparrow [33] applies this algorithm to work around incorrect samplings. The intuition of *late-binding* is that the worker nodes first verify the correctness of bottleneck detection, and then launch the assigned tasks. If such verification fails, the task will be reassigned.

Our bottleneck-aware scheduling algorithm has a prospective bird-view over bottleneck severities of all nodes, thanks to the online bottleneck detection. Basically, the bottleneck-aware scheduling algorithm is seeking a worker that not only best matches the locality preferred by the task, but also keeps bottleneck-free in future period of time. Pseudocode for this algorithm is shown in Algorithm 1. Note that a task usually has a list of preferred localities sorted by preference levels in descending order. However, if no bottleneck-free worker meets the task’s locality preference, the task will be assigned to a random picked worker that is bottleneck-free. In case of potential false positives in bottleneck detection, our scheduler will call an error-tolerance algorithm first before the task is sent to the matched worker.

4.1 Error Tolerance

Our bottleneck detection algorithms are not always accurate. It is impossible to have one hundred percent accuracy. A false positive detection further worsens the bottleneck severity, since a node detected as bottleneck-free will be assigned with more tasks. Similar to Sparrow [33], we introduce a *late-binding* algorithm to overcome detection inaccuracies in bottleneck-aware scheduling.

Algorithm 1: Bottleneck-Aware Scheduling Algorithm

```

Input: taskList,           ▷ pending task list
         workerList,        ▷ running worker list
         btOfWorker,       ▷ bottleneck severity hashmap
Output: taskToWorker     ▷ assignment decisions

1 for task in taskList do
2   for locality in task.localityList do
3     ▷ search a worker that matches the locality
4     worker = search(workerList, locality);
5     btLevel = btOfWorker[worker];
6     if btLevel is low then           ▷ bottleneck-free
7       ▷ for error-tolerance, call Algorithm 2
8       binding = lateBind(task, worker);
9       add binding to taskToWorker;
10      break;                       ▷ go to the next task
11     else                             ▷ in bottleneck
12      continue;                     ▷ try next locality
13   if no matched worker then
14     ▷ randomly pick a free worker
15     randFreeWorker = randFree(btOfWorker);
16     ▷ for error-tolerance, call Algorithm 2
17     binding = lateBind(task, randFreeWorker);
18     add binding to taskToWorker;

```

The intuition of *late-binding* is that the workers do not immediately launch the assigned tasks and instead place reservations for the requested resources (*i.e.*, CPU cores and memory). Meanwhile, the scheduler probes performance metric observations and justifies whether the detected bottlenecks are false positive. The scheduler will adopt a back-off strategy when false positives happen. For example, if the observed performance metrics increase rapidly, the detected bottleneck-free will be identified as false positive. The scheduler then releases the resource reservation on the nodes and puts back the assigned tasks to the list of pending tasks. Pseudocode for this algorithm is shown in Algorithm 2. For false negative cases that mistakenly mark a worker node as non-available, the worker will still receive tasks if there are no bottleneck-free worker nodes because the Lube scheduler is work-

Algorithm 2: Late-Binding Algorithm

```

Input: task,           ▷ a task to be binded
         taskList,      ▷ pending task list
         worker,        ▷ detected as bottleneck-free
         obsOfWorker   ▷ observed performance hashmap
Output: binding      ▷ task-to-worker

1 reserve worker.cpu and worker.mem;
2 while tick in timePeriod do   ▷ probe for a time period
3   obsVal = obsOfWorker[worker]; ▷ observation
4   add obsVal to trend;
5   if trend is ascending then   ▷ false positive
6     add task to taskList;       ▷ back to pending list
7     release worker.cpu and worker.mem;
8     return empty binding;
9   else
10  return (task, worker);        ▷ late-binding

```

conserving. In other words, such false negative worker nodes will not be under-utilized, as if there were pending tasks in the queue.

5 IMPLEMENTATION

In order to implement a practical closed-loop framework that enables online bottleneck-aware scheduling, we specifically set our implementation goals as follows:

- **Low-latency:** To detect bottlenecks accurately and make scheduling decisions effectively, the closed-loop of Lube must operate at fine time granularity. Therefore, modules and pipelines integrated in this loop must act in low latency.
- **Extensible:** Data analytics tasks consume various resources. The framework should provide simple interfaces to plug in customized performance monitors, in case of missing potential bottlenecks in different resources.
- **Compatible:** The framework should be able to open the bottleneck detection service to prevailing data analytics engines. It must provide compatible loose-coupled APIs.
- **Low-overhead:** All components in Lube framework must perform under the radar. They must be lightweight enough, introducing negligible impact on running analytics tasks.

To achieve these goals, we implement the software stack and open APIs of Lube as Figure 6. In essence, components are Python lightweight daemon processes that interact with each other asynchronously. Redis [37], an in-memory cache, works as the underlying low-latency and reliable messaging channels. Messages exchanged between different components are tagged with an epoch timestamp. To guarantee consistent timing on the whole cluster, a NTP [3] clock synchronization service is deployed on each node. Before cached in redis, messages are all serialized in a lightweight data-interchange format, JSON [2], .

On each worker node, performance monitors and the bottleneck detection module communicate in a publish/subscribe messaging paradigm. The monitors are continuously publishing performance metric streams to local redis server. As a subscriber, the local bottleneck detection module receives runtime performance metrics and detects potential upcoming bottlenecks. Subsequently, the detected bottleneck severities will be published to a remote redis server on master node. For error-tolerant (Algorithm 2), both observed performance metrics (*i.e.*, val_ob) and detected bottlenecks (*i.e.*, val_inf) will be directed to the redis server on master node.

On the master node, a redis server is deployed to cache detected bottleneck severities and observed performance metrics from each worker node. When a task is ready to run, the Lube Scheduler will make task assignment decisions based on runtime bottleneck severities cached on the redis server.

Our implementation reuses the redis APIs as our interfaces, because the redis APIs have simple and clear definitions as well as a broad support for various programming languages. Inherently, Lube is implemented in a well-decoupled structure with high flexibility and extensibility for further development.

5.1 Pluggable Performance Monitors

A performance monitor is implemented as a daemon process that monitors specific metrics in the background. As an instance of Python multiprocessing.Process, a performance monitor process is optimized for multiple CPU cores and non-blocking. Though implementation based on Python threading consumes

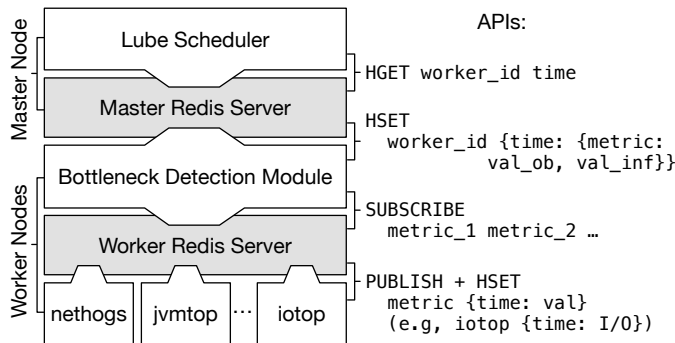


Fig. 6: The software stack and APIs of Lube.

less memory, it suffers high context-switching latency introduced by the Global Interpreter Lock (GIL) of Python.

Plugging in another customized performance monitor is easy. We have provided a base class `PerfMonitor`, with which performance monitor developers only need to: 1) create a child class; 2) pass in a monitor command (e.g., `iotop`) and a filter function (for filtering out redundant output). By tens of lines of code, one can plug a new performance monitor into the Lube framework.

In our prototype of Lube, we have developed performance monitors based on `nethogs`, `jvmtop` and `iotop`, which monitor network I/O, JVM performance and disk I/O at process-level respectively. The output of performance monitors is organized as a `HashMap` following the API style of redis:

```
{'monitor_name':
    timestamp, metric_1, metric_2, ...}
```

The overhead of performance monitors is negligible. Monitoring toolkits like `nethogs`, `jvmtop` and `iotop` are well-tested and widely-deployed on Linux. Our experiments in Section 6 also show the overhead introduced by monitoring per second can be ignored.

5.2 Bottleneck Detector

The Bottleneck Detector is also a daemon process, running locally on each worker. As a subscriber of the local redis server, the Bottleneck Detector periodically updates the machine learning model based on new monitored performance metrics, and asynchronously pushes predicted bottleneck severities to the remote redis server on master node.

The core wrapped by the Bottleneck Detector process is bottleneck detection algorithms. We build statistics model ARIMA from Python `statsmodels.tsa` library and the machine learning model SlidHMM from Python `scikit-learn` library. The models are updated online and serialized as a file saved persistently by Python pickle. The serialization is necessary for potential failures on worker nodes, and avoid training a model from scratch everytime.

The bottleneck information is recorded in a UNIX-style number, each digit indicating the severity level of a specific performance metric. This digit-based structure is easy to extend, efficient to transfer and safe to cache. In our prototype, we define a 4-digit number as follows:

```

n m d c
| | | |
network: 0-4 <-+ | +-> cpu : 0-4
memory : 0-4 <----+ ----> disk: 0-4
```

As an instance, 4201 indicates the intensity levels of the bottleneck in network, memory, disk and CPU are 4, 2, 0 and 1 respectively (the highest level is 4).

5.3 Lube Scheduler

We have to integrate the Lube Scheduler natively into the codebases of data analytics engines. As one of the core components in data analytics engines, the schedulers have few open APIs exposed to the third-party schedulers. Our prototype scheduler is implemented for Spark-1.6.1, one of the most prevailing data analytics engines. Note that our APIs of bottleneck detection are aligned with standard redis APIs, implementing bottleneck-aware schedulers on other data analytics engines like Hadoop [15] and Cosmos [11] will not take much effort.

To implement an efficient and compatible Lube Scheduler, we fully reuse the original codebase of Spark-1.6.1. We implement the Lube Scheduler as a subclass of the Spark default task scheduler, inheriting most functions and data structures we do not need to refactor for the bottleneck-aware scheduling. The default task scheduler maintains two queues to respectively track available worker nodes and pending tasks. Each available worker is represented as a slot of CPU cores and memory capacities. A worker node enters the queue when it has free CPU cores and memory capacities. A task and a worker node will both be dequeued after a task assignment decision is made by the scheduler. The default scheduler applies a delay scheduling strategy [46] for data locality. To make a task assignment decision, the scheduler first dequeues a task from the queue of pending tasks, and iteratively looks for a worker that matches the task's locality preference. Based on this intuition, to enable bottleneck-aware scheduling, we extend the scheduler a little bit — to further check the bottleneck severities after the node and execute late-binding as Algorithm 2.

6 EVALUATION

In this section, we present evaluation results on Lube. We demonstrate that Lube (i) provides accurate bottleneck detection for Spark jobs, (ii) significantly reduces completion times of queries, and (iii) is robust and scalable to large-scale clusters and workloads. The highlights of our evaluation are as follows.

- **Bottleneck detection accuracy:** Experiments on running different workloads show that Lube with SlidHMM achieves over 90% bottleneck detection accuracy.
- **Bottleneck detection robustness:** We show that SlidHMM can hold a bottleneck detection accuracy around 90% when the scales of the cluster and the query increase.
- **Performance improvement:** Lube speeds up the median query response times from 26.88% (1.4×) to 33.46% (1.5×) at negligible overheads.
- **Error-tolerance effectiveness:** Lube Scheduler achieves approximated speedup rate with different bottleneck detection accuracies by ARIMA and SlidHMM.
- **Scalability:** We evaluate Lube with a variety of queries on clusters in different scales. The results show Lube scales well and maintains stable performance benefits upon the increasing scales of clusters and workloads.

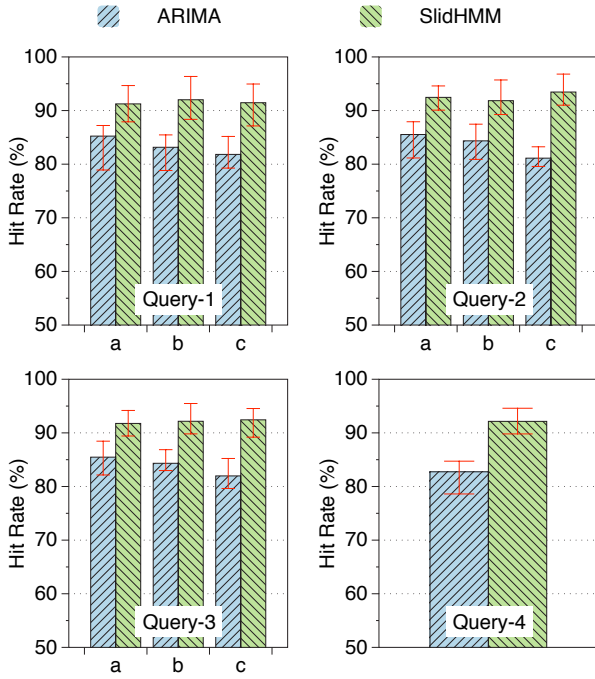


Fig. 7: The accuracy of bottleneck detection.

6.1 Methodology

EC2 deployment: We deploy Lube across 9 EC2 regions in N. Virginia, N. California, Oregon, Ireland, Frankfurt, Tokyo, Seoul, Singapore and Sydney [5]. We launch 37 m4.2xlarge instances in total, with 4 workers in each region and 1 master in one region (N. Virginia). Each instance has a 8 cores CPU, 32 GB memory, 1000 Mbps network² and a 100 GB SSD disk.

Software environment: All instances run on Ubuntu-14.04 with Oracle Java-1.8.0. The whole data analytics platform is build with Spark-1.6.1, HDFS-2.6.4 and Hive-1.2.1. For Spark, we configure each worker with 6 CPU cores and 24 GB memory; For HDFS, we configure each data chunk with 3 replicas, the same as the default.

Workloads: We use the Big Data Benchmark [39] as our workload, with datasets from Intel HiBench [24] and Common Crawl [1] document corpus. The dataset is around 1.1 TB (36 datanodes, 30s GB per each). Derived from Pavlo *et al.* [34], Big Data Benchmark contains a mix of Hive and Spark SQL queries for evaluating performance of large-scale parallel queries. There are four major workloads: Query 1 and Query 2 are both exploratory SQL queries; Query 3 is a join query; Query 4 is a user-defined-function (UDF) query executing an external python script to count URLs. Query 1-3 each has three scale levels *a*, *b*, *c*, from small to large.

Evaluation Metrics: As for bottleneck detection accuracy, we use *hit rate* to represent the accuracy of our bottleneck detection algorithms. A *hit rate* is defined as the proportion of detected bottlenecks that are observed by monitors among all detected bottlenecks. As for benefits of bottleneck-aware scheduling, we compare the query response times and task completion times on

2. EC2 only guarantees intra-region bandwidth. The inter-region traffic runs on public links that is highly fluctuating and intensely competitive.

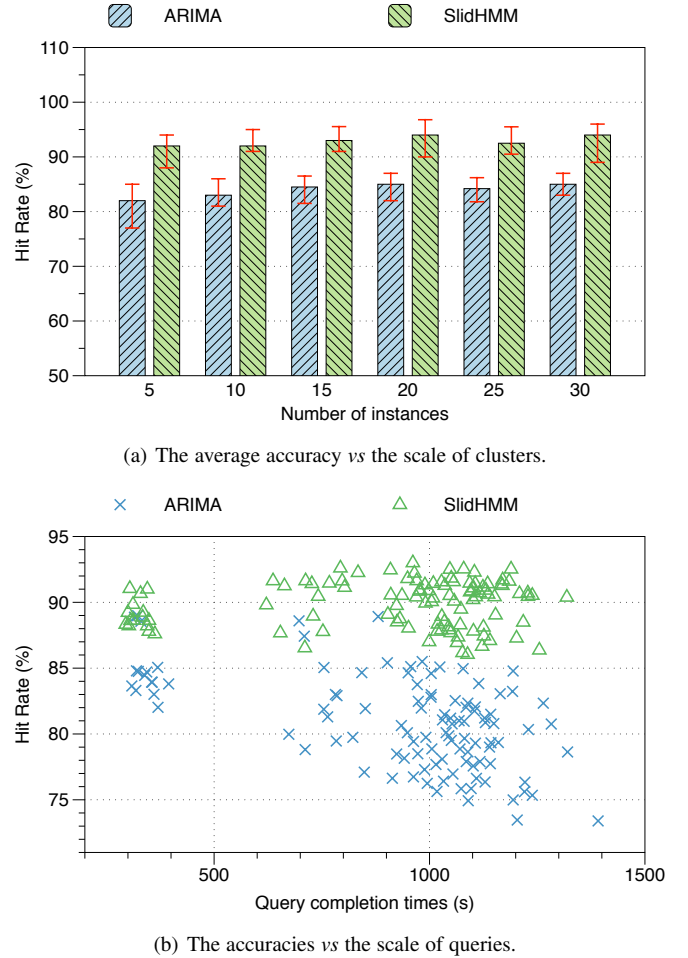


Fig. 8: Analysis of Lube's bottleneck detection accuracy.

pure Spark and Spark with Lube enabled. In addition, we measure the overhead by running Lube with default Spark task scheduler instead of the Lube Scheduler. For effectiveness of Lube's error-tolerance, we compare query response times and task completion times under different detection accuracies.

6.2 Bottleneck Detection Accuracy

We use *hit rate* to represent the accuracy of bottleneck detection accuracy. The *hit rate* is formulated as:

$$hit\ rate = \frac{\#((t, detection) \cap (t, observation))}{\#(t, detection)} \times 100\%$$

which calculates the proportion of verified bottleneck detections. We run each workload of Big Data Benchmark with the ARIMA algorithm and the SlidHMM algorithm for 15 times respectively, tracing both detected bottleneck severity sequences and observation sequences. By calculating *hit rate* offline, we plot the bottleneck detection accuracies under different queries in Figure 7.

Figure 7 presents the accuracies of bottleneck detection under different settings. We collect the time stamps and the bottleneck sequences during the running of the Big Data Benchmark for 15 times respectively. During the collection, it should be noted that the Lube scheduler is disabled in order to observe the bottlenecks. By comparing the time sequences of detected bottlenecks and




Fig. 9: CDF of task completion times.

observed bottlenecks, we calculate the *hit rate* offline. In our experiments, the average *hit rate* of the SlidHMM is 92.1%, while it’s 83.57% for ARIMA. The *hit rate* of ARIMA tends to decrease with the increase of query scale. As a linear combination of autoregression and moving average, ARIMA ignores nonlinear patterns in the performance metric sequences, which may lower the accuracy of bottleneck detection.

6.3 Bottleneck Detection Robustness

We evaluate accuracies of both ARIMA and SlidHMM on clusters consisting of different numbers of instances, when running each query of the Big Data Benchmark. We plot the average accuracies achieved by ARIMA and SlidHMM under different scales of clusters as Figure 8(a). The accuracies of both ARIMA and SlidHMM are stable in terms of an increasing cluster scale, credited to the distributed design of Lube’s bottleneck detection module running locally on each worker node.

Figure 8(b) shows that SlidHMM achieves accuracy around 90% for long-running queries, which last more than 1000 seconds. However, the accuracy of ARIMA decreases with the increase of query completion times.

6.4 Scheduling Effectiveness and Overheads

The results show that with an accuracy of over 90% in bottleneck detection, Lube speeds up median query response times from 26.88% ($1.37\times$) to 33.46% ($1.5\times$). Lube achieves a faster query response and maintains a low overhead from the task level to query level.

At the task level, Figure 9 plots the task completion times CDF of pure Spark, Lube-ARIMA and Lube-SlidHMM. For Query 1, the average (75th percentile) task completion time of pure Spark is 150.928 seconds (246.19 seconds). Lube-ARIMA

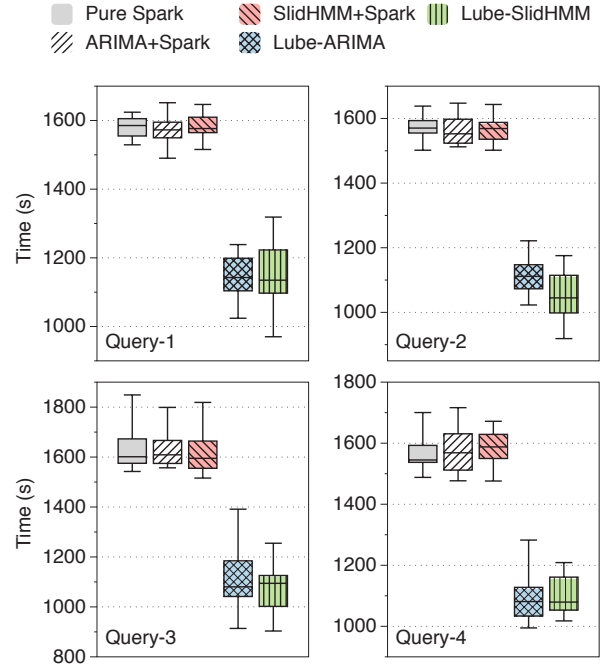


Fig. 10: Query response times.

saves 12.454 seconds (22.075 seconds) for average seconds (75th percentile) tasks compared to pure Spark, while Lube-SlidHMM saves 14.783 seconds (27.469 seconds) for average (75th percentile) tasks. Our bottleneck-aware scheduler brings a substantial improvement to the completion times of long tasks.

At the query level, we measure query response times under different control groups. Pure Spark is the baseline; Lube-ARIMA and Lube-SlidHMM show the reduction of query response times; and, the Spark default scheduler with Lube-ARIMA (ARIMA + Spark) and Lube-SlidHMM (SlidHMM + Spark) are the control group to evaluate Lube’s overhead. Figure 10 shows that running Lube-ARIMA or Lube-SlidHMM with the Spark default scheduler does not introduce much overhead since the query response times under these three settings are similar. In addition, for median query response times, Lube reduces 26.88% to 33.46% ($1.37\times$ to $1.5\times$ faster) of time with the ARIMA algorithm, while reduces 28.41% to 33.18% ($1.4\times$ to $1.5\times$ faster) of time with the SlidHMM algorithm. From these results, we can conclude that Lube reduces the overall query response times.

In addition, though there is a gap of nearly 10% in accuracies of ARIMA algorithm and SlidHMM algorithm, Lube achieves close performance in reducing task completion times and query response times with these two algorithms. Since the Lube Scheduler applies a late-binding algorithm for error-tolerance, 10% accuracy loss has little impact on query performance.

6.5 Scalability

We evaluate the scalability of Lube by analyzing the reduction of query completion times when scale up the input data sizes and cluster instance numbers. Figure 11 presents the scalability of Lube in three dimensions: (i) the cluster size, (ii) the query input data size, and (iii) the query complexity.

The reduction rate of query completion times is calculated in this way:

$$\text{reduction rate} = \frac{\text{average}(t_{\text{pure}}) - t_{\text{lube}}}{\text{average}(t_{\text{pure}})} \times 100\%,$$

in which the $\text{average}(t_{\text{pure}})$ is the average completion times of queries running on pure Spark, the t_{lube} is the completion times of queries running on Lube-ARIMA or Lube-SlidHMM.

Figure 11(a) plots the reduction rates of query completion times when queries are executed on different sizes of clusters. The completion time reduction rates achieved by both Lube-ARIMA and Lube-SlidHMM float within a steady range of 26% to 34%. The scale-up of the cluster has negligible effects on the reduction rates.

Figure 11(b) plots the reduction rates of completion times when queries take different scales of input data. Lube-SlidHMM is more robust to the increase of input data sizes, compared to Lube-ARIMA, which tends to have less performance improvement (*i.e.*, query completion time reductions).

Figure 11(c) plots the relationships between the reduction rates and the $\text{average}(t_{\text{pure}})$, which implies the inherent complexity of queries. There’s no obvious decrease in reduction rates for both Lube-ARIMA and Lube-SlidHMM when the $\text{average}(t_{\text{pure}})$ increases.

7 RELATED WORK

There exists large volumes of existing research on optimizing the performance of wide-area data analytics. Clarinet [41] pushes wide-area network awareness to the query planner, and selects a query execution plan before the query begins. Graphene [22] presents a Directed Acyclic Graph (DAG) scheduler with awareness of DAG dependencies and task complexity. Iridium [36] optimizes data and task placement to reduce query response times and WAN usage. Geode [42] minimizes WAN usage via data placement and query plan selection. SWAG [25] adjusts the order of jobs across datacenters to reduce job completion times. These works develop their solutions based on a few widely-accepted mantras, which are shown to be skeptical in a systematic analysis on the performance of data analytics frameworks [32]. The *blocked time analysis* proposed in [32] calls for more attention to temporal performance variations.

However, there is still very little existing effort on optimizing the performance of data analytics with the awareness of variations in the runtime environment. Hadoop speculative task execution [45] duplicates tasks that are slow or failed, but not knowing the exact bottlenecks may lead to worse performance. As far as we know, Lube is the first work that leverages machine learning techniques to detect runtime bottlenecks and schedules tasks with awareness of performance bottlenecks. Lube is orthogonal to schedulers of the existing wide-area data analytic frameworks. These schedulers can enable bottleneck awareness at runtime by integrating with Lube.

Machine learning techniques have been actively applied to predict and classify data analytics workloads. NearestFit [13] establishes accurate progress predictions of MapReduce jobs by a combination of nearest neighbor regression and statistical curve fitting techniques. Ernest [40] applies a linear regression model to predict the performance of large-scale analytics. CherryPick [4] uses Bayesian Optimization to build cloud performance models for specific applications and cloud configurations.

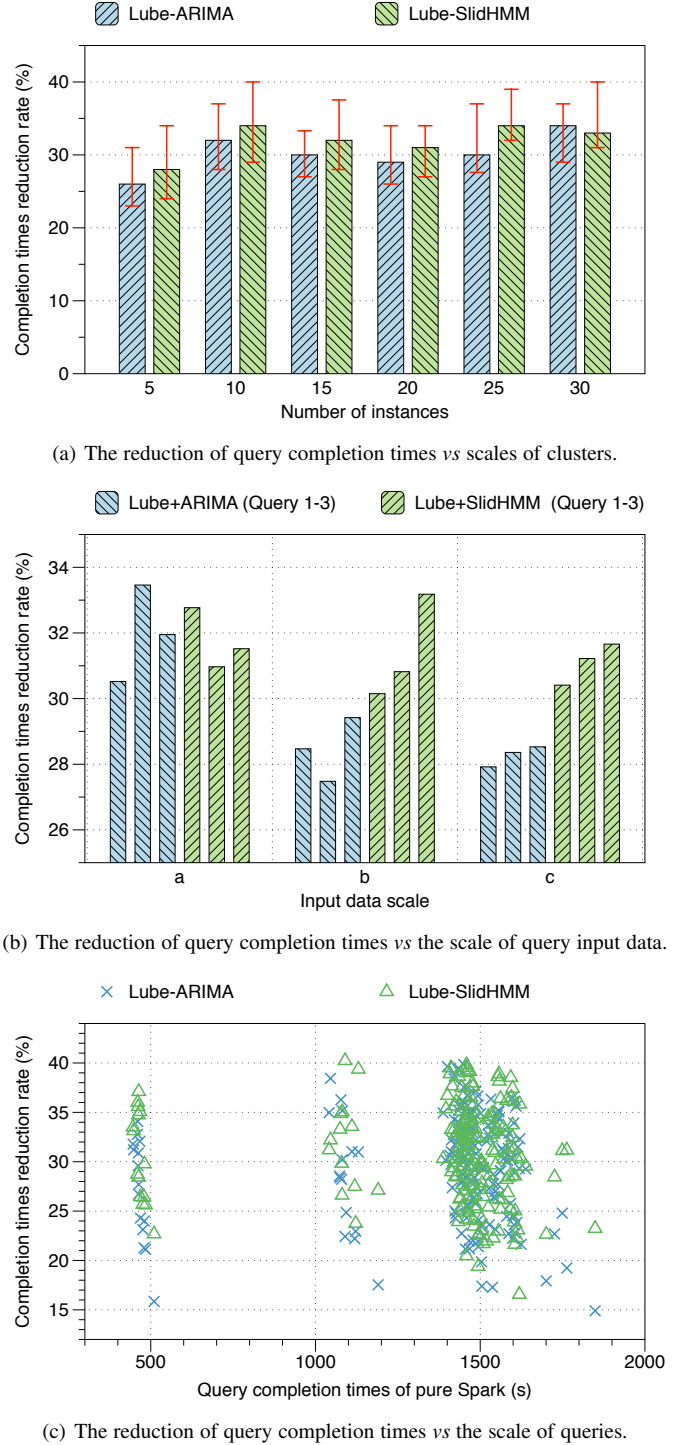


Fig. 11: Analysis of query performance improvements of Lube.

A Kernel Canonical Correlation Analysis (KCCA) model is applied to predict query properties including completion time and resource demands [19]. ASPr [28] applied Artificial Neural Networks (ANNs) to forecast overhead of function calls and loops in script languages.

8 DISCUSSIONS

Our preliminary experiments have highlighted the performance of Lube in reducing query response times achieved through detecting

and mitigating bottlenecks at runtime. While this motivates the research on performing data-driven runtime performance analysis to optimize data analytics frameworks, there are a few aspects that need additional discussions.

Selection of runtime metrics. It is the selection of runtime metrics that determine the efficacy of the runtime performance analysis. There exist an enormous volume of runtime metrics from multiple hierarchies within wide-area data analytics frameworks. To efficiently detect and mitigate bottlenecks in low-level resources (*e.g.*, CPU, memory, disk I/O and network I/O *etc.*), we have studied several performance monitors and various combinations of performance metrics. However, the space of selecting appropriate metrics has still not been fully explored. We will put more efforts in the exploration of runtime metrics and the practice of feature selection techniques such as LASSO Path [18] and principal component analysis (PCA) [35].

Bottleneck detection models. Lube achieves a substantial improvement by applying two simple models, ARIMA and SlidHMM. The emerging data-driven techniques have broadened the horizon of data analytics optimization methodologies. We would like to further explore the latest data-driven techniques, such as Generative Adversary Network (GAN) [20] and Reinforcement Learning [38]. For example, DeepRM [30] builds a deep reinforcement learning model for strategies of cluster resource management. However, the surprising accuracy of machine learning models makes us wonder the practical boundary of their effectiveness, which is imperative for robust and reproducible solutions.

Fine-grained scheduling. The scheduling policy applied by Lube is bottleneck-aware but coarse-grained due to the absence of task resource demand. Recent work on profiling tasks of data analytic applications characterized resource utilization of tasks by an analytical or statistical model [16, 21, 29, 31]. Besides, users can also provide the application resource requirements [23]. By jointly considering the task resource demand and the runtime bottleneck severity, Lube can enforce a fine-grained task scheduling policy that mitigates the bottleneck of a type of resource required by particular tasks. This fine-grained scheduling policy is expected to achieve higher resource utilization and to further reduce query completion times.

WAN conditions. Most recent work considered the heterogeneity and the variance of wide-area network bandwidths [4, 22, 25, 36, 40, 42]. A few approaches have been applied to measure network conditions in these works. Lube captures the local network throughput by measuring network I/O on each node, which only revealed a coarse-grained awareness of network. It also attempts to measure the pair-wise WAN bandwidth by a cron job running iperf on each node. In our future work, we plan to exploit the capabilities of Software-Defined Networking (SDN) to complement the global wide-area network conditions at runtime.

9 CONCLUSION

In this paper, we have presented Lube, a closed-loop framework that mitigates bottlenecks at runtime to improve the performance of wide-area data analytics. Lube monitors runtime query performance, detects bottlenecks online and mitigates them with a bottleneck-aware scheduling policy. Experiments across nine EC2 regions show that Lube achieves over 90% bottleneck detection

accuracy and, compared to the default Spark scheduler, reduces the median query response time by up to 33%.

As a highlight, this work delves into runtime performance of wide area data analytics and delivers a responsive strategy as a workaround. It introduces an orthogonal solution to today's scheduling strategies in the wide area — detecting and mitigating bottlenecks online. Lube is instrumental to improving performance profiling techniques, machine learning on time series analysis, and optimized resource scheduling strategies. For those who work on scheduling strategies in distributed data analytics engines, our pluggable performance sensor module will provide valuable insights on the real-time performance metrics they care about.


10 ACKNOWLEDGMENTS

This work is supported by a research contract with Huawei Technologies Co. Ltd. and an NSERC Collaborative Research and Development (CRD) grant.

REFERENCES

- [1] Common Crawl Web Crawl Data. <http://commoncrawl.org>, 2016. [Online; accessed 1-July-2016].
- [2] JSON Website. <http://json.org>, 2016. [Online; accessed 1-May-2016].
- [3] NTP: The Network Time Protocol. <http://www.ntp.org>, 2016. [Online; accessed 1-July-2016].
- [4] O. Alipourfard, H. H. Liu, J. Chen, S. Venkataraman, M. Yu, and M. Zhang. CherryPick: Adaptively Unearthing the Best Cloud Configurations for Big Data Analytics. In *Proc. USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, pages 4–2, 2017.
- [5] Amazon. Amazon Web Services. <https://aws.amazon.com/about-aws/global-infrastructure/>, 2016. [Online; accessed 1-July-2016].
- [6] Apache. Apache Hadoop Official Website. <http://hadoop.apache.org/>, 2016. [Online; accessed 1-May-2016].
- [7] Apache. Spark SQL. <https://spark.apache.org/sql/>, 2016. [Online; accessed 1-July-2016].
- [8] L. E. Baum and T. Petrie. Statistical inference for probabilistic functions of finite state Markov chains. *The annals of mathematical statistics*, 37(6):1554–1563, 1966.
- [9] L. E. Baum, T. Petrie, G. Soules, and N. Weiss. A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *The Annals of Mathematical Statistics*, 41(1):164–171, 1970.
- [10] G. E. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung. *Time series analysis: forecasting and control*. John Wiley & Sons, 2015.
- [11] R. Chaiken, B. Jenkins, P.-Å. Larson, B. Ramsey, D. Shakib, S. Weaver, and J. Zhou. SCOPE: easy and efficient parallel processing of massive data sets. *Proc. the VLDB Endowment*, 2008.
- [12] T. Chis. Sliding Hidden Markov Model for Evaluating Discrete Data. In *10th Proc. Computer Performance Engineering European Workshop (EPEW)*, volume 8168, 2013.
- [13] E. Coppa and I. Finocchii. On data skewness, stragglers, and MapReduce progress indicators. In *Proc. the 6th ACM Symposium on Cloud Computing*. ACM, 2015.
- [14] J. G. De Gooijer and R. J. Hyndman. 25 years of time series forecasting. *International Journal of Forecasting*, 22(3):443–473, 2006.
- [15] J. Dean and S. Ghemawat. MapReduce: Simplified Data Processing on Large Clusters. In *Proc. USENIX Symposium on Operating System Design and Implementation (OSDI)*, 2004.
- [16] C. Delimitrou and C. Kozyrakis. Quasar: Resource-efficient and QoS-aware Cluster Management. In *Proc. the 19th International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, 2014.
- [17] D. A. Dickey and W. A. Fuller. Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74(366a):427–431, 1979.
- [18] J. Friedman, T. Hastie, and R. Tibshirani. *The Elements of Statistical Learning*. Springer series in statistics New York, 2001.
- [19] A. Ganapathi, H. Kuno, U. Dayal, J. L. Wiener, A. Fox, M. Jordan, and D. Patterson. Predicting multiple metrics for queries: Better decisions enabled by machine learning. In *IEEE 25th International Conference on Data Engineering (ICDE)*, pages 592–603, 2009.

- [20] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative Adversarial Nets. In *Advances in neural information processing systems (NIPS)*, pages 2672–2680, 2014.
- [21] R. Grandl, G. Ananthanarayanan, S. Kandula, S. Rao, and A. Akella. Multi-resource packing for cluster schedulers. *ACM SIGCOMM Computer Communication Review*, 44(4):455–466, 2015.
- [22] R. Grandl, S. Kandula, S. Rao, A. Akella, and J. Kulkarni. Graphene: Packing and Dependency-aware Scheduling for Data-Parallel Clusters. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, page 81, 2016.
- [23] B. Hindman, A. Konwinski, M. Zaharia, A. Ghodsi, A. D. Joseph, R. H. Katz, S. Shenker, and I. Stoica. Mesos: A Platform for Fine-Grained Resource Sharing in the Data Center. In *Proc. USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2011.
- [24] S. Huang, J. Huang, J. Dai, T. Xie, and B. Huang. The HiBench Benchmark Suite: Characterization of the MapReduce-based Data Analysis. In *Proc. International Conference on Data Engineering Workshops (ICDEW)*, 2010.
- [25] C. Hung, L. Golubchik, and M. Yu. Scheduling Jobs Across Geo-Distributed Datacenters. In *Proc. ACM Symposium on Cloud Computing (SoCC)*, 2015.
- [26] V. Jalaparti, P. Bodik, I. Menache, S. Rao, K. Makarychev, and M. Caesar. Network-Aware Scheduling for Data-Parallel Jobs: Plan When You Can. In *ACM SIGCOMM Computer Communication Review*, pages 407–420, 2015.
- [27] K. Kloudas, M. Mamede, N. Preguiça, and R. Rodrigues. Pixida: Optimizing Data Parallel Jobs in Wide-Area Data Analytics. *VLDB Endowment*, 9(2):72–83, 2015.
- [28] J. Li, X. Ma, K. Singh, M. Schulz, B. R. de Supinski, and S. A. McKee. Machine learning based online performance prediction for runtime parallelization and task scheduling. In *IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, 2009.
- [29] C.-K. Luk, S. Hong, and H. Kim. Qilin: Exploiting Parallelism on Heterogeneous Multiprocessors with Adaptive Mapping. In *Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture*. ACM, 2009.
- [30] H. Mao, M. Alizadeh, I. Menache, and S. Kandula. Resource Management with Deep Reinforcement Learning. In *Proc. the 15th ACM Workshop on Hot Topics in Networks*. ACM, 2016.
- [31] V. S. Marco, B. Taylor, B. Porter, and Z. Wang. Improving Spark Application Throughput Via Memory Aware Task Co-location: A Mixture of Experts Approach. In *Proc. the 18th ACM/IFIP/USENIX Middleware Conference*. ACM, 2017.
- [32] K. Ousterhout, R. Rasti, S. Ratnasamy, S. Shenker, and B.-G. Chun. Making sense of performance in data analytics frameworks. In *Proc. USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2015.
- [33] K. Ousterhout, P. Wendell, M. Zaharia, and I. Stoica. Sparrow: distributed, low latency scheduling. In *Proc. the 24th ACM Symposium on Operating Systems Principles (SOSP)*, 2013.
- [34] A. Pavlo, E. Paulson, A. Rasin, D. J. Abadi, D. J. DeWitt, S. Madden, and M. Stonebraker. A comparison of approaches to large-scale data analysis. In *Proc. the ACM International Conference on Management of data (SIGMOD)*, pages 165–178, 2009.
- [35] K. Pearson. LIII. On Lines and Planes of Closest Fit to Systems of Points in Space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572, 1901.
- [36] Q. Pu, G. Ananthanarayanan, P. Bodik, S. Kandula, A. Akella, P. Bahl, and I. Stoica. Low latency Geo-Distributed Data Analytics. In *Proc. ACM SIGCOMM 2015*. ACM, 2015.
- [37] redis. Redis Website. <http://redis.io/>, 2016. [Online; accessed 1-May-2016].
- [38] R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT press Cambridge, 1998.
- [39] UC Berkeley AMPLab. The Big Data Benchmark. <https://amplab.cs.berkeley.edu/benchmark/>, 2016. [Online; accessed 1-July-2016].
- [40] S. Venkataraman, Z. Yang, M. Franklin, B. Recht, and I. Stoica. Ernest: efficient performance prediction for large-scale advanced analytics. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*. USENIX Association, 2016.
- [41] R. Viswanathan, G. Ananthanarayanan, and A. Akella. Clarinet: WAN-aware optimization for analytics queries. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*. USENIX Association, 2016.
- [42] A. Vulimiri, C. Curino, P. Godfrey, T. Jungblut, J. Padhye, and G. Varghese. Global Analytics in the Face of Bandwidth and Regulatory Constraints. In *Proc. USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2015.
- [43] A. Vulimiri, C. Curino, P. Godfrey, K. Karanasos, and G. Varghese. WANalytics: Analytics for A Geo-Distributed Data-Intensive World. In *Proc. Conference on Innovative Data Systems Research (CIDR)*, pages 1087–1092, 2015.
- [44] H. Wang and B. Li. Lube: mitigating bottlenecks in wide area data analytics. In *Proc. the 9th USENIX Conference on Hot Topics in Cloud Computing*, pages 1–1. USENIX Association, 2017.
- [45] T. White. *Hadoop: The Definitive Guide*. " O'Reilly Media, Inc.", 2012.
- [46] 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. ACM European Conference on Computer Systems*, pages 265–278, 2010.
- [47] M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M. Franklin, S. Shenker, and I. Stoica. Resilient Distributed Datasets: A Fault-Tolerant Abstraction for In-Memory Cluster Computing. In *Proc. USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2012.



Hao Wang received both of his B.E. degree in Information Security and M.E. degree in Software Engineering from Shanghai Jiao Tong University, Shanghai, China, in 2012 and 2015 respectively. His research interests include large-scale data analytics, distributed computing, machine learning and datacenter networking.



Baochun Li (F) received his B.E. degree from the Department of Computer Science and Technology, Tsinghua University, China, in 1995 and his M.S. and Ph.D. degrees from the Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, in 1997 and 2000.

Since 2000, he has been with the Department of Electrical and Computer Engineering at the University of Toronto, where he is currently a Professor. He holds the Nortel Networks Junior Chair in Network Architecture and Services from October 2003 to June 2005, and the Bell Canada Endowed Chair in Computer Engineering since August 2005. His research interests include cloud computing, large-scale data processing, computer networking, and distributed systems. In 2000, Dr. Li was the recipient of the IEEE Communications Society Leonard G. Abraham Award in the Field of Communications Systems. In 2009, he was a recipient of the Multimedia Communications Best Paper Award from the IEEE Communications Society, and a recipient of the University of Toronto McLean Award. He is a Fellow of IEEE and a member of ACM.