

# DRS: Dynamic Resource Scheduling for Real-Time Analytics over Fast Streams

Tom Z. J. Fu<sup>1</sup> Jianbing Ding<sup>2</sup> Richard T. B. Ma<sup>1,3</sup> Marianne Winslett<sup>4</sup> Yin Yang<sup>5</sup> Zhenjie Zhang<sup>1</sup>

<sup>1</sup>Advanced Digital Sciences Center, Illinois at Singapore Pte. Ltd.

<sup>2</sup>School of Information Science and Technology, Sun Yat-sen University

<sup>3</sup>School of Computing, National University of Singapore

<sup>4</sup>Department of Computer Science, University of Illinois at Urbana-Champaign

<sup>5</sup>College of Science and Engineering, Hamad Bin Khalifa University

Email: {tom.fu,zhenjie}@adsc.com.sg, dingsword@gmail.com, tbma@comp.nus.edu.sg, winslett@illinois.edu, yyang@qf.org.qa

**Abstract**—In a data stream management system (DSMS), users register continuous queries, and receive result updates as data arrive and expire. We focus on applications with real-time constraints, in which the user must receive each result update within a given period after the update occurs. To handle fast data, the DSMS is commonly placed on top of a cloud infrastructure. Because stream properties such as arrival rates can fluctuate unpredictably, cloud resources must be dynamically provisioned and scheduled accordingly to ensure real-time response. It is essential, for the existing systems or future developments, to possess the ability of scheduling resources dynamically according to the current workload, in order to avoid wasting resources, or failing in delivering correct results on time.

Motivated by this, we propose DRS, a novel dynamic resource scheduler for cloud-based DSMSs. DRS overcomes three fundamental challenges: (a) how to model the relationship between the provisioned resources and query response time (b) where to best place resources; and (c) how to measure system load with minimal overhead. In particular, DRS includes an accurate performance model based on the theory of *Jackson open queueing networks* and is capable of handling *arbitrary operator topologies*, possibly with loops, splits and joins. Extensive experiments with real data confirm that DRS achieves real-time response with close to optimal resource consumption.

## I. INTRODUCTION

In many applications, such as analytics over microblogs, video feeds and sensor readings, data records are not available beforehand, but gradually and continuously arrive in the form of streams. A data stream management system (DSMS) handles such streams, and answers long-running, continuous queries to users. The results of such a query are delivered in the form of a stream of updates. Often, users are interested in performing streaming analytics in *real time*, meaning that each result update must reach the user within a given time period after the update occurs, i.e., the earliest possible time that it can be produced. For instance, consider a DSMS monitoring surveillance video streams in hospital wards. Events such as a patient falling should be detected promptly to alarm doctors and nurses in time.

To deal with fast, high-volume streams and stringent real-time response requirements, it is increasingly common to place the DSMS on top of a cloud infrastructure, which provides virtually unlimited computing resources on demand. Because

key properties of a data stream, including its volume, arrival rates, value distribution, etc., can fluctuate in an unpredictable manner, the DSMS should ideally dynamically provision cloud resources to each application, in order to satisfy the real-time constraints with minimum resource consumption. Meanwhile, inside an application, resources need to be carefully *scheduled* to different components to ensure optimal utilization. Misplacing resources may cause not only poor resource utilization, but instability of the system as a whole.

Figure 1 shows an example video stream processing application with two operators  $A$  (which extracts features from input video frames) and  $B$  (which recognizes objects from the extracted features), with the output of  $A$  fed to  $B$  as input. The record arrival rates for  $A$  and  $B$  are  $\lambda_A$  and  $\lambda_B$  respectively, where  $\lambda_A$  depends on the input, e.g., 24 frames per second, while  $\lambda_B$  depends on the output rate of  $A$ , i.e., the number of features extracted in unit time. Inside each operator, an input is first buffered into an *input queue* (i.e.,  $q_A$  in  $A$  and  $q_B$  in  $B$ ) before being processed by one of the parallel processors ( $A_1, \dots, A_n$  in  $A$ ,  $B_1, \dots, B_m$  in  $B$ ). Assuming the cloud provides identical processing units, each processor in  $A$  (respectively in  $B$ ) can process  $\mu_A$  ( $\mu_B$ ) inputs in a unit of time. Clearly, an operator must have sufficient processors to keep up with its input rate; otherwise, inputs start to fill its input queue, leading to increased latency due to waiting, and, eventually, errors when the queue reaches its size limit. Since the data arrival rate and processing rate for each processor are uncontrollable, the main resource scheduling issue is to determine the number of processors in each operator, in our example,  $n$  and  $m$ <sup>1</sup>.




Fig. 1. Example streaming analytics application.

A simple approach to scheduling resources is to monitor

<sup>1</sup>Although there are other types of cloud resources, such as storage and network bandwidth, we focus on computation-intensive applications where processors are the key resource.

the workload in each operator, and adjust the number of processors accordingly. This method is insufficient in multi-operator applications. For instance, consider the case that at some point, many recognizable objects appear in the video stream. Then, although the number of frames per second in the input (i.e.,  $\lambda_A$ ) remains stable, each frame now contains more extractable features, requiring more work at operator  $A$ . Hence,  $\mu_A$  decreases, which consequently overloads operator  $A$ , causing inputs to wait longer in its queue  $q_A$ , slowing down query response. Now, if we naively add processors to  $A$  to flush  $q_A$ , operator  $A$  then suddenly produces a large amount of outputs, leading to a burst in the input rate  $\lambda_B$  of operator  $B$ , overloading the latter. This problem is exacerbated when the application involves a complex network of operators. Figure 2 shows such an example, with splits ( $A$  to  $B, C$ ), joins ( $C, D$  to  $E$ ) and a feedback loop ( $E$  to  $A$ ). Such topological features are key enablers for certain applications, e.g., loops allows data reduction at the input based on the current query results, as we show with an example in Section V.




Fig. 2. Example complex operator topology.

As we review in Section II, existing systems largely overlook the problem of dynamic resource scheduling. Consequently, to meet the real-time constraint, they either require manual tuning at runtime (which is infeasible for dynamic streams), overprovisioning resources to each operator (which wastes resources), or load shedding (which leads to incorrect results). Motivated by this, we design and implement DRS, a dynamic resource scheduling module. DRS generally applies to operator-based DSMSs, and allows operators to form an *arbitrary* topology, possibly with splits, joins and loops as shown in Figure 2. In particular, the support for loops can be a key enabler for certain applications, especially those involving iterations, as we show with an example in Section V. Meanwhile, from a semantics point of view, allowing arbitrary topologies is more general than two-step MapUpdate in Mupet [1], and the DAG model in TimeStream [2].

Our main contributions include effective and efficient solutions to three fundamental problems in dynamic resource scheduling: (a) how much resources are needed, (b) where to best place the allocated resources to minimize response time, and (c) how to implement resource scheduling in a real system with minimal overhead. In particular, our solutions to the first two problems are based on the theory of extended Jackson networks, which provides an educated estimate of system performance.

The rest of the paper is organized as follows. Section II surveys related work. Section III presents our performance model and optimization algorithm. Section IV describes the implementation of DRS. Section V contains an extensive set of experiments with real data. Section VI concludes with directions for future work.

## II. RELATED WORK

### A. Resource Scheduling in Cloud Systems

A cloud consists of a massive number of interconnected commodity servers. A key feature of the cloud is that its resources, such as CPU cores, memory, disk space and network bandwidth can be provisioned to applications on demand. In fact, most cloud infrastructure providers today offer pay-as-you-go options for resource usage. Hence, a fundamental requirement for a system to effectively use the cloud is *elasticity*, meaning that the system must be able to dynamically allocate and release cloud resources based on the current workload. Many traditional parallel and distributed systems, however, assume a fixed amount of resources available beforehand, rendering them unsuitable to be applied in a cloud platform. As a result, many novel elastic cloud-based paradigms and systems have emerged in the past decade.

The first wave of cloud-based systems were built for running a batch of (often slow) jobs offline. Notably, MapReduce [3] is a batch processing framework that hides the complexity of the cloud infrastructure, and exposes a simple programming interface to users consisting of two functions: *map* (e.g., for data filtering and transformation) and *reduce* (for aggregation and join). A plethora of MapReduce systems, improvements, techniques, and optimizations have been proposed in recent years, and we refer the reader to a comprehensive survey [4].

Resource scheduling has been a central problem in MapReduce like systems, and a plethora of schedulers have been developed and used in production, e.g., Fair Scheduler [5], Capacity Scheduler [6]. Since tasks running on nodes without relevant data incur costly network transmissions, delay scheduling [7] reduces such non-local tasks by forcing nodes to wait until either a local task appears, or a specified period has passed. These scheduling strategies, however, do not apply to our problem, because they are designed for offline, batch processing of (semi-) static data, where the goal is to minimize *total* job completion time; in contrast, we focus on real-time processing of streaming data, where each *individual* result update must be delivered on time.

Recently, much attention has been shifted to real-time interactive systems for big data analytics, such as Dremel [8], Impala, Presto [9], OceanRT [10], [11], C-Cube [12], SADA [13] and newer versions of Hive [14]. Such systems deal with static rather than streaming data; meanwhile, the term “real-time” here has a different meaning: that each query is executed quickly enough so that the user can wait online for its results. Hence, resource scheduling in these systems resembles offline systems, and their techniques do not apply to our problem for similar reasons. Another recent hot topic in cloud-based system research is cloud-based stream processing, which is most relevant to this work. We review them in Section II-C.

Finally, there exist generic scheduling solutions for provisioning to multiple applications competing for cloud resources. System such as Mesos [15], YARN [16] are prominent examples. Abacus [17] optimizes total utility by allocating resources via a truth-revealing auction. These methods generally assume that an application already knows the amount of resources it needs, and how to distribute these resources internally, which are the problems solved in this paper. Hence, they can be used in combination with the proposed solution.

## B. Traditional DSMSs

Stream processing has been an important research topic in both academia and industry. Earlier work focuses on DSMSs in a centralized setting, which resembles the traditional, centralized database management systems. For instance, STREAM [18] establishes formal semantics for queries over streams [19], and proposes efficient query processing algorithms, e.g., [20]. Similar systems include Aurora [21], GigaScope [22], TelegraphCQ [23], and System S [24]. Scheduling in such centralized systems means deciding the best order of operators to execute (by the central processor), e.g., in order to minimize memory consumption [25]. Hence, scheduling strategies in these systems such as [25] do not apply to our cloud-based setting, where operators are executed by multiple nodes in parallel, and computational resources are dynamically provisioned on demand.

Similarly, DSMSs built for traditional parallel settings, notably Borealis [26], also differ from cloud-based DSMSs in that the former assume that a fixed amount of computational resources available beforehand, rather than dynamically allocated. Hence, to our knowledge, no scheduling technique along this line of research applies to our problem. Next we review cloud-based DSMSs.

## C. Cloud-Based Stream Processing

There are two general methodologies for processing streams in a cloud: using an operator-based DSMS, and discretizing stream inputs into mini-batches [27]. The former derives from traditional DSMSs described in Section II-B, whereas the latter reduces stream processing to batch execution, explained in Section II-A. In general, mini-batch systems are optimized for throughput, at the expense of increased query response time, since each input must wait until a full batch is formed. While it is possible to minimize this extra latency by having extremely small batches, doing so would lead to high overhead, defeating the purpose. We focus on operator-based DSMSs since our target applications have real-time constraints, in which response time is key.

Two popular open source operator-based DSMSs are Storm [28] and S4 [29]. Their main difference is that Storm guarantees the correctness of its results (e.g., through its Trident component), while S4 does not. Both systems rely on manual configurations for resource scheduling. Hence, to avoid slow responses due to operator overloading, the user has to either overprovision resources to every operator, which is wasteful, or continuously tuning the system, which is infeasible for dynamic streams.

Many research prototypes of operator-based DSMSs are proposed, such as TimeStream [2], which features efficient fault recovery, and Samza [30]. None of these systems, however, addresses the resource scheduling problem. In the following we present DRS, the first effective resource scheduler for cloud-based operator DSMSs.

## III. DYNAMIC RESOURCE SCHEDULING

Section III-A clarifies assumptions in DRS. Section III-B presents the DRS performance model, which estimates query

response time given a resource allocation scheme. Section III-C describes the DRS dynamic resource scheduling algorithm. Table I summarizes frequently used notations throughout the paper.

TABLE I. TABLE OF NOTATIONS.

Symbol	Meaning
$N$	Total number of operators in an application
$\lambda_i$	Mean arrival rate of inputs to $i$ -th operator
$\lambda_0$	Mean arrival rate of inputs to the application
$\mu_i$	Mean processing rate of inputs to $i$ -th operator
$k_i$	# of processors allocated to the $i$ -th operator
$\mathbf{k}$	A Vector $(k_1, \dots, k_N)$ containing all $k_i$ s.
$T_{\max}$	Real-time constraint parameter: each input of the application is expected to be fully processed within $T_{\max}$ time.
$K_{\max}$	Resource constraint parameter: maximum number of available processors that can be allocated to the operators.
$t$	an input tuple to the streaming application
$T$	A random variable on the total sojourn time of a tuple $t$

## A. Assumptions

We focus on stream analytics applications, which are usually memory-based and computation intensive. For such applications, *processors* are the main type of resource, each of which contains a CPU (or one of its cores) and a certain amount of RAM. Disk space is not critical as streaming inputs are computed on-the-fly. Although networking delay can also affect query latency, we do not explicitly model it, because (a) it is often correlated with computational costs and (b) it can be affected by uncontrollable factors, such as other transmission-heavy applications on the same server or in the same subnetwork. Further, data centers today are increasingly equipped next-generation networking hardware that provide significantly higher bandwidth and lower latency, such as 10G Ethernet and InfiniBand, whose prices have been dropping rapidly. In contrast, processor speed in terms of CPU clock rate and RAM latency has stagnated in the past few years. Hence, we assume processors to be the bottleneck of the system, not network bandwidth.

For the ease of presentation, we further assume that all processors in the cloud have *identical computational power*. Nevertheless, the proposed models and algorithms can also support settings with heterogeneous processors, and we explain how this is done whenever necessary. Meanwhile, we assume that *load balancing* is achieved in every operator, i.e., each processor inside the same operator performs roughly equal amount of work. How to achieve load balancing is an orthogonal topic, and it is under active research, e.g., [31]. Under these assumptions, the processing speed of an operator depends mainly on the number of processors therein.

The goal of DRS is to fully process each input of the application in real time. Specifically, an input tuple to the application, e.g., a video frame in Figure 1, may lead to multiple intermediate results, e.g., features extracted by operator  $A$ , and objects recognized by operator  $B$ . We say an input tuple  $t$  is fully processed, if and only if every intermediate result derived from  $t$  has been processed by its corresponding operator. We use the term *total sojourn time* to refer to the duration from the time that  $t$  first arrives at the system to the time that  $t$  is fully processed. Our goal is then to ensure that the expected total sojourn time of each input  $t$  is no more than a user-specified duration, denoted by  $T_{\max}$ .

## B. Performance Model

Given an application's operator network, e.g., the one in Figure 2, the current resource allocation and characteristics of the streaming data, the DRS performance model estimates the total sojourn time of an average input of the application, explained at the end of last subsection. The current resource allocation is represented by the number of processors assigned to each operator. Formally, we define  $N$  as the number of operators in an application and a resource allocation is modeled by a vector  $\mathbf{k} = (k_1, k_2, \dots, k_N)$ , where  $k_i (1 \leq i \leq N)$  corresponds to the number of processors allocated to the  $i$ -th operator.

Regarding data characteristics, the important variables are the rate that tuples arrive at each operator, and how fast they can be processed by one processor. Networking delay is not explicitly expressed in our model, and we discuss this issue further at the end of this subsection. Note that our model assumes neither deterministic tuple arrival rates nor processing times; in other words, instantaneous arrival rates and processing times can fluctuate. On the other hand, in order to make the problem tractable, we do assume that the system remains in a relatively steady state during the span that DRS performs modeling and resource scheduling. This means that the average tuple arrival rate and processing time at each operator remains stable, and we obtain these quantities through the measurement module of the system, described in Section IV. Specifically, for the  $i$ -th operator ( $1 \leq i \leq N$ ), we use  $\lambda_i$  to denote the mean arrival rate of its inputs, and  $\mu_i$  to denote the mean processing rate of each of its processors. For instance, the case of  $k_i = 3$ ,  $\lambda_i = 10$  and  $\mu_i = 3$  means that on average, 10 tuples arrive at the  $i$ -th operator in unit time, and each of its 3 processors processes 3 tuples in unit time. For an operator with multiple input streams, i.e., join operators,  $\lambda_i$  is the total arrival rates of all its input streams, and  $\mu_i$  is the average processing rate of the operator, regardless of which input stream the tuple comes from.

Additionally, we define  $\lambda_0$  as the mean arrival rate of inputs that flow into the application's operator network from outside of it. When there are clear "source" operators in the operator network whose inputs come entirely from outside the network,  $\lambda_0$  is simply the total arrival rates of these sources. In general, however, there may not be a simple relationship between  $\lambda_0$  and the set of  $\lambda_i$ 's,  $1 \leq i \leq N$ . For example, in Figure 2,  $\lambda_0$  is the arrival rate of tuples that come (from outside the system) to operator  $A$ ; the input arrival rate  $\lambda_A$  for  $A$  on the other hand is the sum of  $\lambda_0$  and the arrival rate of  $A$ 's other input stream, produced by operator  $E$ .

We use random variable  $T$  to denote the total sojourn time of an input to the application. Our goal is to estimate  $E[T]$ , i.e., the expected value of  $T$ . The basic idea for estimating  $E[T]$  is to model the system as an open queueing network (OQN) [32], and apply known results in queueing theory. In OQN, the total sojourn time of an input tuple  $t$  is computed by summing up its total service time (i.e., total time spent on processing  $t$  and intermediate results derived from  $t$ ) and total queuing delay (total time that  $t$  and its derived tuples wait in operator queues). This closely matches our setting. The challenge, however, is that there are numerous OQN models in the queueing literature, and selecting an appropriate one is non-trivial. On one hand, complex queueing network models generally do not have known solutions; among the

ones that do, most have only numerical solutions (rather than analytical ones), which renders effective optimization hard; on the other hand, an overly simplified model may rely on strong assumptions, such as deterministic tuple arrival rates, which do not hold in our setting. After comparing various options and testing them through experiments, we chose to build our model based on a combination of one of Erlang's models [33] and the Jackson network [32], [34]. The former enables effective analysis of each individual operator, and the latter helps to aggregate these analyses to estimate  $E[T]$  for the whole network. Our model has an analytical solution, and it involves only mild limitations, which will be discussed shortly.

We first focus on a single operator, say the  $i$ -th. We use  $T_i$  to denote the time between the arrival of an input of the operator and the time when the operator finishes processing it. We model the operator as an M/M/ $k_i$  system [33], where  $k_i$  is the number of processors for operator  $i$ . According to the Erlang formula [33],  $E[T_i]$  is calculated by:

$$E[T_i](k_i) = \begin{cases} \frac{\left(\frac{\lambda_i}{\mu_i}\right)^{k_i} \pi_0}{k_i! \left(1 - \frac{\lambda_i}{\mu_i k_i}\right)^2 \mu_i k_i} + \frac{1}{\mu_i} & \text{for } k_i > \frac{\lambda_i}{\mu_i}; \\ +\infty & \text{for } k_i \leq \frac{\lambda_i}{\mu_i}, \end{cases} \quad (1)$$

where  $\pi_0$  is a normalization term, given by:

$$\pi_0 = \left[ \sum_{l=0}^{k_i-1} \frac{\left(\frac{\lambda_i}{\mu_i}\right)^l}{l!} + \frac{\left(\frac{\lambda_i}{\mu_i}\right)^{k_i}}{k_i! \left(1 - \frac{\lambda_i}{\mu_i k_i}\right)} \right]^{-1}. \quad (2)$$

Intuitively, since new tuples arrive at an average rate  $\lambda_i$ , and each processor processes tuples at an average rate  $\mu_i$ , when  $k_i \leq \frac{\lambda_i}{\mu_i}$ , the processors cannot keep up with incoming tuples. Consequently, the number of tuples in the operator queue increases with time, leading to infinite queuing delay. When,  $k_i > \frac{\lambda_i}{\mu_i}$ , tuples are expected to be handled faster than they arrive. However, due to the randomness of the arrival and process rates, the queue may still grow when the arrival rate is temporarily higher than the processing rate. Clearly, the expected service time for each tuple is  $\frac{1}{\mu_i}$ . The expected queuing delay is captured by the complicated term in Equation (1).

Next we aggregate all  $E[T_i]$ 's to obtain an estimate of  $E[T]$  for the entire operator network. According to the theory of Jackson networks [32], [34],  $E[T]$  is computed by a weighted average of the  $E[T_i]$ 's:

$$E[T](\mathbf{k}) = E[T](k_1, k_2, \dots, k_N) = \frac{1}{\lambda_0} \sum_{i=1}^N \lambda_i E[T_i](k_i). \quad (3)$$

This completes the DRS performance model. Since our model relies on Erlang's formula and Jackson open queueing network, it inherits two limitations. First, the model implicitly assumes that both the inter-arrival times of external tuples (that come from outside the system) and the service time of the operator are *i.i.d.* samples from random variables following the exponential distribution. Second, Jackson network does not explicitly model pipelining between different operators. Hence, our model may give an inaccurate estimate of  $E[T]$ , when the service time or tuple arrival distribution deviates significantly

from the expected exponential distribution, or when pipelining affects total processing time considerably. In the meantime, our model does not explicitly consider networking costs, due to the fact that measuring the networking delay between two nodes requires complex inter-node protocols, e.g., for clock synchronization, which can be prohibitively expensive in a real-time application. Therefore, when networking delay becomes a dominant factor in the total sojourn time of an average input, our model tends to produce an underestimation of the true result. Nevertheless, as we show in the experiments, the value of  $E[T]$  predicted by our model is sufficiently accurate, when the underlying application is computation intensive, which is one important assumption made in Section III-A. Further, even when the prediction is inaccurate, it is still strongly correlated with the exact value of  $E[T]$ , meaning that DRS remains capable of identifying the best resource allocation with the predicted value. In the rest of the section, we show how DRS schedules resources based on the performance model.

### C. Scheduling Algorithm

In a nutshell, DRS (a) monitors the current performance of the system (more details in Section IV), (b) checks whether the performance falls (or is about to fall) under the real-time constraint, or when the system can fulfil the constraint with less resources, and (c) reschedules resources when (b) returns a positive result. The main challenge lies in (b), which needs to answer two questions, including *how many* processors are needed to fulfil the real-time requirement, and *where* to place them in the operator network. We first focus on the latter question. Specifically, given a number (say,  $K_{\max}$ ) of processors, we are to find an optimal assignment of these processors to the operators of the application that obtains the minimum expected total sojourn time. The problem can be mathematically formalized as follows:

$$\begin{aligned} \min_{\mathbf{k}} E[T](\mathbf{k}) \\ \text{s.t. } \sum_{i=1}^N k_i \leq K_{\max}, k_i \text{ is interger}, i = 1, 2, \dots, N; \end{aligned} \quad (4)$$

A naive solution to the above optimization problem is to view it as an integer program, and apply a standard solver. However, current integer programming solvers are prohibitively slow, especially considering that DRS itself has to run in real time. In the following we describe a novel algorithm that solves Program (4) with negligible cost.

The key property used in the proposed algorithm is that  $E[T_i](k_i)$ , defined in Equation (1), is a convex function of  $k_i$ , the number of processors assigned to the  $i$ -th operator. This property has already been proved in [32]. It follows from the convexity of  $E[T_i](k_i)$  implies that marginal benefit for incrementing  $k_i$  drops monotonously as  $k_i$  becomes larger. Formally, for all  $k'_i > k_i$ , we have:

$$E[T_i](k_i) - E[T_i](k_i + 1) > E[T_i](k'_i) - E[T_i](k'_i + 1) \quad (5)$$

Now observe from Equation (3) that  $E[T]$  is a weighted sum of the  $E[T_i]$ s, and each weight  $\lambda_i$  is independent of the value of  $k_i$ . Hence,  $E[T]$  is also a convex function of the  $k_i$ s, meaning that incrementing each  $k_i$  also has diminishing

marginal benefit with respect to  $E[T]$ . Based on this observation, we design a greedy algorithm, listed in Algorithm 1. The idea is to start from the smallest possible value of each  $k_i$  (lines 1-4) and iteratively add one processor to the operator that leads to the largest decrease in  $E[T]$  (lines 8-15). According to Equation (1), each  $k_i$  must be larger than  $\frac{\lambda_i}{\mu_i}$ , since otherwise,  $E[T_i](k_i)$  becomes infinitely large, leading to infinity on  $E[T]$  as well.

---

#### Algorithm 1 AssignProcessors

---

**Input:**  $K_{\max}$ ,  $\lambda_0$ ,  $\{\lambda_i, i = 1, \dots, N\}$ ,  $\{\mu_i, i = 1, \dots, N\}$ .

**Output:**  $\mathbf{k} = (k_1, k_2, \dots, k_N)$

```

1: for all  $i \leftarrow 1, \dots, N$  do
2:    $k_i \leftarrow \left\lceil \frac{\lambda_i}{\mu_i} \right\rceil$  /* Initialize each  $k_i$  */
3: end for
4: if  $\sum_{i=1}^N k_i > K_{\max}$  then
5:   throw an exception saying that the number of processors
     are not sufficient for the application.
6: end if
7: while  $\sum_{i=1}^N k_i < K_{\max}$  do
8:   for all  $i \leftarrow 1, \dots, N$  do
9:      $\delta_i \leftarrow \lambda_i \cdot [E[T_i](k_i) - E[T_i](k_i + 1)]$ 
10:  end for
11:  /* find the operator with the largest marginal benefit. */
12:   $j \leftarrow \arg \max_i \delta_i$ 
13:   $k_j \leftarrow k_j + 1$ 
14: end while
15: return  $\mathbf{k} = (k_1, k_2, \dots, k_N)$ 

```

---

Since  $E[T]$  is convex, the above greedy algorithm always finds the optimal solution, similar to the case of the server reallocation problem [35]. This is restated as follows:

*Theorem 1:* Algorithm 1 always returns exact optimal solution to Program 4.

The proof is given in [36].

Next we focus on the question on how to determine the minimum number of processors that are expected to achieve real-time processing, i.e., the expected total sojourn time  $E[T]$  is no larger than a user-defined threshold  $T_{\max}$ . This can be modeled with the following optimization problem.

$$\begin{aligned} \min_{\mathbf{k}} \sum_{i=1}^N k_i, \\ \text{s.t. } E[T](\mathbf{k}) \leq T_{\max}, k_i \text{ is interger}, i = 1, 2, \dots, N; \end{aligned} \quad (6)$$

Similar to Program (4), both constraints and objective of Program (6) are convex in terms of  $\mathbf{k}$ . Hence, we solve Program (6) with a greedy strategy similar to Algorithm 1. Specifically, we start by initializing each  $k_i$  with minimal requirement, as in lines 1-4 of Algorithm 1. The algorithm repeatedly adds one processor to the operator with the maximum marginal benefit as in lines 8-15 of Algorithm 1, until  $E[T]$  is no larger than  $T_{\max}$ . We omit the proof of correctness for this algorithm, since it is nearly identical to that of Algorithm 1.

In practice, the solution of Program (6) may not give us the precise amount of resources necessary for meeting the real-time requirement at all times, for two reasons. First, the

total sojourn time can be different for every input, and  $E[T]$  is merely its expected value. Second, the performance model described in Section III-B outputs only an estimate of  $E[T]$ , rather than its precise value. To address this problem, DRS starts with the number of processors suggested by the solution of Program (6), monitors the actual total sojourn time  $E[\hat{T}]$ , and continuously adjusts the number of processors based on the measured value of  $E[\hat{T}]$ . In next section, we discuss the system design and implementation issues with DRS.

#### IV. SYSTEM DESIGN

An overview of the system architecture is presented in Figure 3, which generally consists of two layers, the DRS layer and the CSP (cloud-based streaming processing) layer. Specifically, DRS layer is responsible for performance measurement, resource scheduling and resource allocation control, while CSP layer contains the primitive streaming processing logic, e.g. running instances of *Storm* [28] and *S4* [29], and cloud-based resource pool service, e.g. YARN [16] and Amazon EC2.

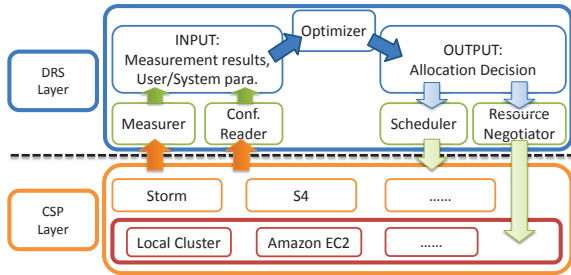


Fig. 3. The architecture overview

While the core of DRS layer is responsible to the optimization of resource scheduling based on the model derived in the previous section, the system to support such functionality is not that straightforward to build. Given the heterogeneous underlying infrastructure and the complicated streaming processing applications running on the CSP layer, it is crucial to collect the accurate metrics from the infrastructure, aggregate the statistics, make online decisions and control the resource allocation in an efficient manner.

To seamlessly combine the optimization model and the concrete streaming processing system, we build a number of independent functional modules, which bridge the gap between the physical infrastructure and abstract performance model. As is shown in Figure 3, on the input side of the optimizer component, we have *measurer* module and *configuration reader* module, which generate the statistics needed by the optimizer based on the data/control flow from CSP layer. On the output end of the workflow diagram, the *scheduler* module and *resource negotiator* module transform the decisions of the optimizer into executable commands for different streaming processing platforms and resource pools. The technical details and key features of the modules are discussed in [36].

#### V. EMPIRICAL STUDIES

To test the effectiveness of DRS, we have implemented  $it^2$  and integrated it into *Storm* [28], which provides the underlying CSP layer. The overview of the important concepts

and architectural aspects of Storm, and the description of how we implement the measurer, scheduler and resource negotiator modules of DRS in Storm are provided in [36].

#### A. Testing Applications

We implement two real-time stream analytics applications: video logo detection (VLD) and frequent pattern detection (FPD) from different domains.

**Logo Detection from a Video Stream.** Given a set of query logo images, the logo detection application identifies these images from the input video stream. Although much work has been done to improve the accuracy and efficiency of VLD, performing it in real time remains a major challenge, due to the high computational complexity.



Fig. 4. The topology of real-time video logo detection application.

Figure 4 illustrates the topology of the real-time VLD application, which is a chain of operators containing a spout, a feature extractor, a feature matcher, and an aggregator. The spout extracts frames from the raw video stream. The output rate of frames may vary from time to time due to the generation algorithm and the original video contents. We employ scale-invariant feature transform (SIFT) [37] algorithm to extract features from each frame. This step is time-consuming, involving convolutions on the 2-dimensional image space. Moreover, the number of result SIFT features may vary dramatically on different frames, causing significant variance on the computation overhead over time. The feature matcher measures  $L_2$  distance between its input SIFT features to those pre-generated logo features, and outputs matching pairs with distance lower than a pre-defined threshold. Finally, the aggregator judges whether a logo appears in a video frame by aggregating all input matching feature pairs, i.e., if the number of matched features in a video frame exceeds a threshold, the logo is considered to appear in the frame.

**Frequent Pattern Detection over a Microblog Stream.** This application maintains the frequent patterns [38] on a sliding window over a microblog stream from Twitter. For each input sentence, we append an additional label “+/-”, indicating it is entering/leaving the dedicated window. Given a set of input item groups in the sliding window and a threshold, we define a maximal frequent pattern (MFP) to be the itemset satisfying: (a) the number of item groups containing this itemset, called its occurrence count, is above the threshold; and (b) the occurrence count of any of its superset is below the threshold.



Fig. 5. The topology of the stream frequent pattern detection application.

Figure 5 illustrates the operator topology. There are two spouts, which generate an event tuple as an itemset enters/leaves the current processing window, respectively. The pattern generator generates candidate patterns, i.e., itemsets. These candidates include an exponential number of possible

<sup>2</sup>The source code is available online: <https://github.com/ADSC-Cloud/resa/>

non-empty combinations of items. Hence, its computation varies, according to the number of items in recent transactions.

The detector maintains the state records containing (a) the occurrence counts and (b) MFP indicator, of all the candidate itemsets. When a state change happens to some itemset, e.g., from MFP to non-MFP, the detector outputs a notification to the reporter, and also to itself through the loop back link. Since (a) each processor in the detector maintains only a portion of the state records; and (b) a state change can affect the states of other itemsets stored at a different processor, the loop ensures that the state change notifications be sent to all the instances. Finally, the reporter presents the updates of the detection results to the user. In our implementation, the reporter simply write its inputs to an HDFS file.

### B. Experiment Setup

The experiments were run on a cluster of 6 Ubuntu Linux machines interconnected by a LAN switch. Each machine is equipped with an Intel quad-core CPU 3.4GHz and 8GB of RAM. Following common configurations of Storm, we allocated one machine to host the Nimbus and the Zookeeper Server; the remaining 5 machines host executors for the experimental applications. We also configured each of these 5 machines so that one machine can host at most 5 executors. The main purpose of this constraint is to mitigate the interference caused by other executors running on the same machine, and the resource contention due to the over-allocation of executors on a single machine. As a result, there are 25 executors in total.

For both applications, namely video logo detection (VLD) and frequent pattern detection (FPD), we allocated two executors as spouts, and one executor for DRS. The remaining  $25 - 3 = 22$  executors are used as bolts, i.e.,  $K_{\max} = 22$ . For VLD, the input data are a series of videos clips of the soccer games, and we selected 16 logos as the detection targets. The frame rate simulates a typical Internet video experience, which is uniformly distributed in the interval  $[1, 25]$  with a mean of 13 frames/second. For FPD, we use a real dataset containing 28,688,584 tweets from 2,168,939 users collected from Oct. 2006 to Nov. 2009. We set the sliding window to 50,000 tweets, and simulated the arrival of tweets to the topology following the Poisson process with an average arrival rate of 320 tweets per second.

### C. Experimental Results

For both applications, we run two sets of experiments: (a) with re-balancing<sup>3</sup> disabled, i.e., we keep DRS running passively, meaning that it continues to monitor the system performance and recommend new (if better) resource allocation configurations, but does not perform re-scheduling; (b) re-balancing is disabled at the beginning, and then enabled at a later time. These experiments aim to test the quality of the performance model and evaluate the effectiveness of the resource scheduling algorithm of DRS.

**Experiments with re-balancing disabled.** In this set, each experiment lasts for 10 minutes. Figure 6 shows the mean and standard deviation of the total sojourn times under 6 different

allocations for each application. The x-axis ( $x_1:x_2:x_3$ ) denotes an resource configuration (in a partial order of  $x_1, x_2, x_3$ ), where  $x_1, x_2, x_3$  are the number of executors allocated to the operators *SIFT Feature Extractor*, *Feature matcher*, and *Matching aggregator* in Figure 4, or the *Pattern generator*, *detector*, and *reporter* in Figure 5. The two configurations with “\*”, (10:11:1) for VLD and (6:13:3) for FPD are the recommended allocations by the passively running DRS.

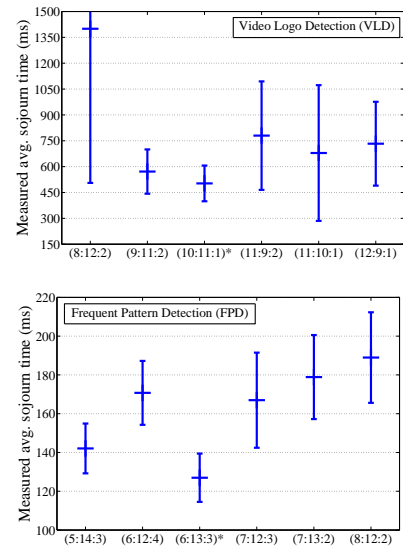


Fig. 6. The mean and standard deviation of the complete sojourn times under different resource configurations with re-balancing disabled, where the configurations with “\*” are the recommended allocations by the passively running DRS.

From Figure 6, we make the following observations. The resource configurations (10:11:1) for VLD and (6:13:3) for FPD, achieve the best performance according to the measured average sojourn time. This turns out to be consistent with the recommendations provided by the passively running DRS, which validates the accuracy and effectiveness of our DRS performance model and resource scheduling algorithm.

In particular, these two configurations not only obtain the smallest average sojourn times, but also the minimum standard deviation, meaning that these two allocations lead to the smallest performance oscillations. Different configurations, including the 5 closest ones in terms of the  $L_1$  distance (i.e., the remaining 5 in the experiment) to the best configurations (10:11:1) for VLD and (6:13:3) for FPD, all exhibit considerably worse performance. These results demonstrate that it is not trivial to find the optimal resource allocation especially when the application topology becomes more complicated (e.g. more than three bolt operators), and hence reveal the importance and usefulness of the DRS.

To take a close look at how DRS provides resource configuration recommendations, correctly, Figure 7 shows the relationship between the measured average sojourn times and the estimated average sojourn time, which is derived by the performance model described in Section III-B, of the six resource allocation configurations for both VLD and FPD, with re-balancing disabled.

As shown in Figure 7, the points representing the measured and the estimated average sojourn time are showing the strictly

<sup>3</sup>This is a term used by Storm, and it has the same meaning as re-scheduling.

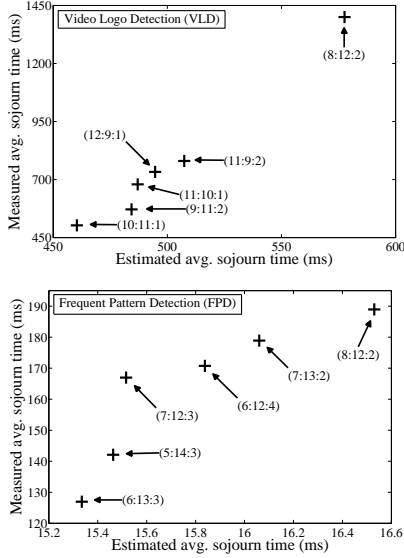


Fig. 7. Comparing average sojourn time estimated by the model and measured in the experiment.

monotonicity, which signifies that the performance model is capable of suggesting the best resource allocation configuration. Moreover, the performance model outputs accurate estimates for VLD; though, with some slight underestimation comparing to the measured values, which is expected, as our model does not consider network overhead. It is worth noting that the estimates are accurate even though the underlying conditions are not satisfied for the Jackson network theory and Erlang model. For example, the frame rate is uniformly (rather than exponential as required) distributed. Meanwhile, the operator input queues do not follow strict FIFO rule; instead, tuples are hashed to processors. Different operators are also run in parallel, which leads to pipelining. The model is clearly robust to these variations of the conditions.

For FPD, the estimated sojourn times show larger deviations to the measured ones. This is mainly because the model does not consider network transmission cost, which takes a dominant portion of the total query latency in this particular application. In other words, the FPD is de facto the type of data intensive rather than the computation intensive application that we focus on. Nevertheless, our model still correctly indicates the relative order of the performance of different resource allocation configurations. Meanwhile, since the estimates are strongly correlated with the true values, a polynomial regression can be used straightforwardly to make accurate predictions of the true latency value given the estimated one.

To further validate the above explanation, we carried out a separate experiment over a synthetic topology with a simple chain of three operators. Each operator simply performs some computations (such as empty for-loops) with varying load (e.g., number of loops). We used 30 executors running on 6 physical machines, connected in the same subnetwork. The results are reported in Figure 8.

As shown in Figure 8, We tried 6 different workloads in terms of total CPU time (excluding the queue time) of the three bolts, from 0.567 millisecond, to 309.1 milliseconds (x-




Fig. 8. The degree of underestimation (the ratio of the measured to the estimated average sojourn time) v.s. the total CPU time of the three bolts of the synthetic chain topology

axis, log-scale), and the y-axis shows the ratio of the measured average sojourn time to the estimated value. It shows a clear decreasing trend of the degree of underestimation (ratio of the measured to the estimated average sojourn time) as the total CPU time of the three bolts increases.

**Experiments with re-balancing disabled first and then enabled.** In this set of experiments, we investigate the performance of the real re-scheduling operation activated and executed by the DRS when it detects the non-optimal resource allocation configurations. For each experiment, it lasts for 27 minutes and the re-balancing function is disabled from the beginning till the end of the 13th minute, and becomes enabled afterwards. In this way, we are able to have a clear view of the performance (in terms of the average sojourn time) across the re-scheduling events.

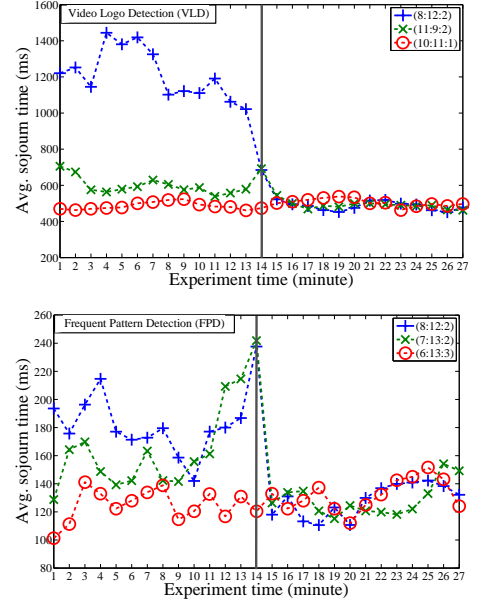


Fig. 9. The average sojourn times of three different allocations in the initial state for each application, where re-balancing function is disabled from the beginning till the end of the 13th minute, and becoming and keeping enabled since the 14th minute.

Figure 9 shows three curves for each of the applications. In particular, each curve represents an initial allocation. For both applications, the two curves initially with the non-optimal allocations, experience the re-scheduling events at the 14th minute, while the one with the optimal allocation as its initial state does not. From Figure 9, we can see that optimizer triggers the re-scheduling action as early as possible, which responds quickly to the less promising resource scheduling



plan. After the re-scheduling, all the curves with different initial allocations, were scheduled with the unique optimal solution. This statement is supported by two facts: a) from Figure 9, after the 14th minute, all the three curves have the similar average sojourn time, and similar performance trends. Especially for the two curves that experience the re-allocation event, it shows a clear decrease in the average sojourn time; b) the plans kept in the log files further verify this observation.

Another observation we make, according to the four curves experiencing the re-scheduling events shown in Figure 9 - (8:12:2) and (11:9:2) of VLD and (8:12:2) and (7:13:2) of FPD - is that our improved version of re-balancing mechanism led to remarkably low cost, i.e. a neglectable increment in the average sojourn time within the 14th minute only. Besides, the whole re-balancing process of ours only takes a few seconds, comparing to the 1-2 minutes taken by Storm's default version.

Next, we investigate how DRS adjusts resources when it detects the resource shortage/wastage according to the configured parameter  $T_{\max}$ . Two experiments on VLD application are conducted and each one lasts for 27 minutes and the re-balancing function is disabled from the beginning till the end of the 13th minute, and becomes enabled afterwards. The average tuple complete sojourn time of the two experiments in each minute is plotted in Figure 10. In particular, for "ExpA", we set  $T_{\max} = 500$  (ms) and in the initial state, 4 workers with  $K_{\max} = 17$  are allocated; and for "ExpB", we set  $T_{\max} = 1000$  (ms) and initially, 5 workers with  $K_{\max} = 22$  are allocated.

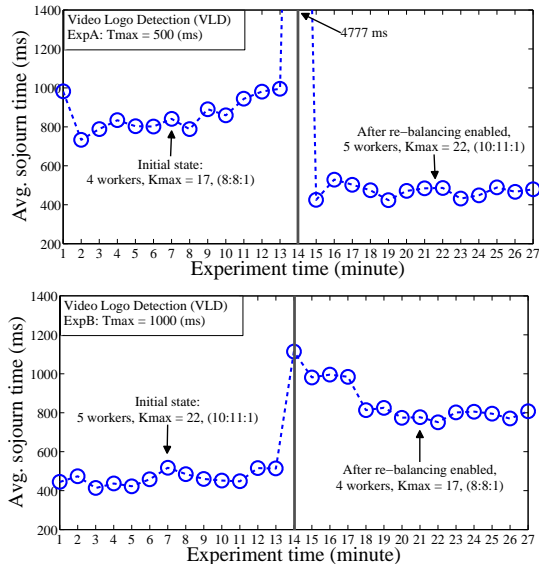


Fig. 10. The average sojourn times under two configurations of the VLD application, where re-balancing function is disabled from the beginning till the end of the 13th minute, and becoming and keeping enabled since the 14th minute. For "ExpA", we set  $T_{\max} = 500$  (ms) and in the initial state, 4 workers with  $K_{\max} = 17$  executors are allocated; and for "ExpB", we set  $T_{\max} = 1000$  (ms) and initially, 5 workers with  $K_{\max} = 22$  executors are allocated.

As shown in Figure 10, the curve of "ExpA" keeps with the allocation configuration (8:8:1) which is actually the suggested allocation when  $K_{\max} = 17$  (by DRS algorithm), for the first 13 minutes. It has experienced the larger average tuple complete sojourn time than the configured  $T_{\max} = 500$  (ms). On the 14th minute right after the re-balancing function is enabled, the DRS quickly triggers the re-scheduling operation including:

(a) initializing and adding an extra machine (thus 5 more executors); and (b) calculating the recommended allocation configuration (10:11:1) when  $K_{\max}$  becomes 22. Since then, the curve of "ExpA" is stably below the target requirement of  $T_{\max} = 500$  (ms). On the other hand, the curve of "ExpB" shows a totally opposite shape to that of "ExpA", which is just as expected: it initially keeps with the configuration of (10:11:1) till the end of the 13th minutes. Afterwards, DRS triggers the re-balancing operation and makes "ExpB" using less resources, i.e., 4 machines,  $K_{\max} = 17$  and (8:8:1), but still satisfying the performance requirement  $T_{\max} = 1000$  (ms).

Similar to the observations we made on Figure 9, the cost incurred by our improved version of the re-balancing mechanism in "ExpA" and "ExpB" are again much lower than that of Storm's default version, as demonstrated by Figure 10. Particularly, "ExpB" just experiences an increase to about 1113 (ms) in average sojourn time in the 14th minute, whereas the overhead of "ExpA" is larger, an increase to around 4777 (ms). This is mainly because of the different actions taken during the re-scheduling, i.e., in "ExpA", new machines are initialized and added to the running topology, in which case, *reusing* JVMs has no effects; in contrast in "ExpB", it only needs to stop and remove some existing working machines. Therefore, there is still room for improvement on our version of the re-balancing mechanism, which we consider as the future work.

**The running overhead of the DRS.** To evaluate the computation overhead of the overall DRS layer, we report the CPU time spent by the whole DRS module, including the processing on measurement results and calculating the optimal allocation. In this experiment, we only test on the video logo detection topology composed by three bolt operators with all the parameters,  $\lambda_0$ ,  $\lambda_i$  and  $\mu_i, i = 1, 2, 3$  fixed. We try different  $K_{\max}$ , i.e. total number of executors for all operators. For each value of  $K_{\max}$ , we run the procedure 100,000 times and report the average running time of the whole DRS layer. The results are listed in Table II, with *Scheduling* as the allocation computation and *Measurement* as the metric processing computation. Generally speaking, the computation done by DRS

TABLE II. COMPUTATION OVERHEADS IN MILLISECONDS UNDER DIFFERENT  $K_{\max}$ .

$K_{\max}$	12	24	48	96	192
Scheduling	0.083	0.158	0.323	0.665	1.250
Measurement	0.100	0.100	0.100	0.100	0.100

is almost neglectable, with overhead less than milliseconds in most of the cases. Moreover, the results are consistent with our intuition that the computation consumption is linear to  $K_{\max}$ , as analyzed over Algorithm 1. The time consumed on processing the measurement results is irrelevant to  $K_{\max}$ . In fact, it is affected by the total number of tasks of the topology, as we discussed in [36] that this number keeps immutable when the topology is continuously running.

## VI. CONCLUSION

This paper proposes DRS, a novel dynamic resource scheduler for real-time streaming analytics in a cloud-based DSMS. DRS overcomes several fundamental challenges, including the estimation of the required resources necessary for satisfying real-time requirements, effective and efficient resource provisioning and scheduling, and the efficient implementation of

such a scheduler in a cloud-based DSMS. The performance model of DRS is based on rigorous queuing theory, and it demonstrates robust performance even when the underlying conditions of the theory are not fully satisfied. In addition, we have integrated DRS into a popular system Storm, and evaluated it by conducting extensive experiments based on real applications and datasets.

Regarding future work, we plan to investigate efficient strategies for migrating the system from the current resource configuration to the new one recommended by DRS. This step should minimize additional overhead and result latency during migration, as well as the migration duration, (e.g., [39]). Another interesting direction for future work is to investigate the possibility of improving performance model accuracy with more sophisticated queuing theory.

#### ACKNOWLEDGMENT

This work is supported by Human-Centered Cyberphysical Systems (HCCS) Programme at Advanced Digital Sciences Center from Singapore's A\*STAR. Ding is partially supported by NSF of China under Grant 61173081, and also supported by GZSI's grant 2013Y2-00046. The authors would like to thank Prof. Hongyang Chao from Sun Yat-sen University for her suggestions to this work.

#### REFERENCES

- [1] W. Lam, L. Liu, S. Prasad, A. Rajaraman, Z. Vacheri, and A. Doan, "Muppet: Mapreduce-style processing of fast data," *Proc. of the VLDB Endowment*, vol. 5, no. 12, pp. 1814–1825, 2012.
- [2] Z. Qian, Y. He, C. Su, Z. Wu, H. Zhu, T. Zhang, L. Zhou, Y. Yu, and Z. Zhang, "Timestream: Reliable stream computation in the cloud," in *Proc. of ACM European Conference on Computer Systems*, 2013.
- [3] J. Dean and S. Ghemawat, "Mapreduce: simplified data processing on large clusters," *Comm. of ACM*, vol. 51, no. 1, pp. 107–113, 2008.
- [4] F. Li, B. C. Ooi, M. T. Özsu, and S. Wu, "Distributed data management using mapreduce," *ACM Computing Surveys*, vol. 46, no. 3, p. 31, 2014.
- [5] [http://hadoop.apache.org/docs/r1.2.1/fair\\_scheduler.html](http://hadoop.apache.org/docs/r1.2.1/fair_scheduler.html).
- [6] [http://hadoop.apache.org/docs/r1.2.1/capacity\\_scheduler.html](http://hadoop.apache.org/docs/r1.2.1/capacity_scheduler.html).
- [7] M. Zaharia, D. Borthakur, J. Sen Sarma, K. Elmeleegy, S. Shenker, and I. Stoica, "Delay scheduling: a simple technique for achieving locality and fairness in cluster scheduling," in *Proc. of the European conference on Computer systems*, 2010, pp. 265–278.
- [8] S. Melnik, A. Gubarev, J. J. Long, G. Romer, S. Shivakumar, M. Tolton, and T. Vassilakis, "Dremel: interactive analysis of web-scale datasets," *Proc. of the VLDB Endowment*, vol. 3, no. 1-2, pp. 330–339, 2010.
- [9] M. Traverso, "Presto: Interacting with petabytes of data at facebook."
- [10] S. Zhang, Y. Yang, W. Fan, L. Lan, and M. Yuan, "Oceanrt: Real-time analytics over large temporal data," in *ACM SIGMOD, Demo*, 2014.
- [11] S. Zhang, Y. Yang, W. Fan, and M. Winslett, "Design and implementation of a real-time interactive analytics system for large spatio-temporal data," *Proc. of the VLDB Endowment*, vol. 7(13), pp. 1754–1759, 2014.
- [12] Z. Zhang, H. Shu, Z. Chong, H. Lu, and Y. Yang, "C-cube: Elastic continuous clustering in the cloud," in *Proc. of IEEE ICDE*, 2013.
- [13] R. Cai, Z. Zhang, and Z. Hao, "Sada: A general framework to support robust causation discovery," in *Proc. of ICML*, 2013, pp. 208–216.
- [14] A. Thusoo, J. S. Sarma, N. Jain, Z. Shao, P. Chakka, N. Zhang, S. Antony, H. Liu, and R. Murthy, "Hive-a petabyte scale data warehouse using hadoop," in *Proc. of IEEE ICDE*, 2010, pp. 996–1005.
- [15] 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. of USENIX NSDI*, 2011.
- [16] <http://hadoop.apache.org/docs/current/hadoop-yarn/hadoop-yarn-site/>.
- [17] Z. Zhang, R. T. B. Ma, J. Ding, and Y. Yang, "Abacus: An auction-based approach to cloud service differentiation," in *Proc. of IEEE International Conference on Cloud Engineering*, 2013, pp. 292–301.
- [18] A. Arasu, B. Babcock, S. Babu, M. Datar, K. Ito, I. Nishizawa, J. Rosenstein, and J. Widom, "Stream: the stanford stream data manager (demo)," in *Proc. of the ACM SIGMOD*, 2003, pp. 665–665.
- [19] A. Arasu, S. Babu, and J. Widom, "The cql continuous query language: semantic foundations and query execution," *The International Journal on Very Large Data Bases*, vol. 15, no. 2, pp. 121–142, 2006.
- [20] S. Babu, K. Munagala, J. Widom, and R. Motwani, "Adaptive caching for continuous queries," in *Proc. of IEEE ICDE*, 2005, pp. 118–129.
- [21] D. J. Abadi, D. Carney, U. Çetintemel, M. Cherniack, C. Convey, S. Lee, M. Stonebraker, N. Tatbul, and S. Zdonik, "Aurora: a new model and architecture for data stream management," *The International Journal on Very Large Data Bases*, vol. 12, no. 2, pp. 120–139, 2003.
- [22] C. Cranor, T. Johnson, O. Spatschek, and V. Shkapenyuk, "Gigascop: a stream database for network applications," in *Proc. of the ACM SIGMOD*, 2003, pp. 647–651.
- [23] S. Chandrasekaran, O. Cooper, A. Deshpande, M. J. Franklin, J. M. Hellerstein, W. Hong, S. Krishnamurthy, S. R. Madden, F. Reiss, and M. A. Shah, "Telegraphcq: continuous dataflow processing," in *Proc. of the ACM SIGMOD*, 2003, pp. 668–668.
- [24] H. Andrade, B. Gedik, K.-L. Wu, and P. Yu, "Processing high data rate streams in system s," *Journal of Parallel and Distributed Computing*, vol. 71, no. 2, pp. 145–156, 2011.
- [25] B. Babcock, S. Babu, M. Datar, R. Motwani, and D. Thomas, "Operator scheduling in data stream systems," *The International Journal on Very Large Data Bases*, vol. 13, no. 4, pp. 333–353, 2004.
- [26] D. J. Abadi, Y. Ahmad, M. Balazinska, U. Çetintemel, M. Cherniack, J.-H. Hwang, W. Lindner, A. Maskey, A. Rasin, E. Ryvkina et al., "The design of the borealis stream processing engine," in *Proc. of Conference on Innovative Data Systems Research*, vol. 5, 2005, pp. 277–289.
- [27] M. Zaharia, T. Das, H. Li, S. Shenker, and I. Stoica, "Discretized streams: an efficient and fault-tolerant model for stream processing on large clusters," in *Proc. of the USENIX conference on Hot Topics in Cloud Computing*. USENIX Association, 2012, pp. 10–10.
- [28] A. Toshniwal, S. Taneja, A. Shukla, K. Ramasamy, J. M. Patel, S. Kulkarni, J. Jackson, K. Gade, M. Fu, J. Donham et al., "Storm@ twitter," in *Proc. of ACM SIGMOD*, 2014, pp. 147–156.
- [29] L. Neumeyer, B. Robbins, A. Nair, and A. Kesari, "S4: Distributed stream computing platform," in *Proc. of IEEE International Conference on Data Mining Workshops (ICDMW)*, 2010, pp. 170–177.
- [30] SAMZA, <http://samza.incubator.apache.org/>.
- [31] Y. Xing, S. Zdonik, and J.-H. Hwang, "Dynamic load distribution in the borealis stream processor," in *Proc. of IEEE ICDE*, 2005, pp. 791–802.
- [32] G. R. Bitran and R. Morabito, "State-of-the-art survey: Open queueing networks: Optimization and performance evaluation models for discrete manufacturing systems," *Production and Operations Management*, vol. 5, no. 2, pp. 163–193, 1996.
- [33] H. C. Tijms, *Stochastic modelling and analysis: a computational approach*. John Wiley & Sons, Inc., 1986.
- [34] J. R. Jackson, "Jobshop-like queueing systems," *Management science*, vol. 10, no. 1, pp. 131–142, 1963.
- [35] O. Boxma, A. Rinnooy Kan, and M. Van Vliet, "Machine allocation problems in manufacturing networks," *European Journal of Operational Research*, vol. 45, no. 1, pp. 47–54, 1990.
- [36] T. Z. J. Fu, J. Ding, R. T. B. Ma, M. Winslett, Y. Yang, and Z. Zhang, "DRS: dynamic resource scheduling for real-time analytics over fast streams," Tech. Rep., April 2015. [Online]. Available: <http://arxiv.org/abs/1501.03610>
- [37] T. Lindeberg, "Scale invariant feature transform," *Scholarpedia*, vol. 7, no. 5, p. 10491, 2012.
- [38] D. Burdick, M. Calimlim, and J. Gehrke, "Mafia: A maximal frequent itemset algorithm for transactional databases," in *Proc. of IEEE ICDE*, 2001, pp. 443–452.
- [39] J. Ding, T. Z. J. Fu, R. T. B. Ma, M. Winslett, Y. Yang, Z. Zhang, and H. Chao, "Optimal operator state migration for elastic data stream processing," Tech. Rep., Feb. 2015. [Online]. Available: <http://arxiv.org/abs/1501.03619>