Distributed Publish/Subscribe Query Processing on the Spatio-Textual Data Stream

Chen, Zhida; Cong, Gao; Zhang, Zhenjie; Chen, Lisi; Fu, Tom Z. J.

2017

<https://hdl.handle.net/10356/80746>

<https://doi.org/10.1109/ICDE.2017.154>

© 2017 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at: [<https://dx.doi.org/10.1109/ICDE.2017.154>].

Downloaded on 26 Apr 2023 14:58:58 SGT

Distributed Publish/Subscribe Query Processing on the Spatio-Textual Data Stream

Zhida Chen[†], Gao Cong[†], Zhenjie Zhang[‡], Tom Z.J. Fu^{‡#}, Lisi Chen[‡]

[†]Nanyang Technological University, Singapore

[‡]Advanced Digital Sciences Center, Illinois at Singapore Pte. Ltd., Singapore

[#]Guangdong University of Technology, China

[‡]Hong Kong Baptist University, Hong Kong

[†]chen0936@e.ntu.edu.sg, gaocong@ntu.edu.sg [‡]zhenjie@adsc.com.sg, tom.fu@adsc.com.sg

[‡]chenlisi@comp.hkbu.edu.hk

Abstract—Huge amount of data with both space and text information, e.g., geo-tagged tweets, is flooding on the Internet. Such spatio-textual data stream contains valuable information for millions of users with various interests on different keywords and locations. Publish/subscribe systems enable efficient and effective information distribution by allowing users to register continuous queries with both spatial and textual constraints. However, the explosive growth of data scale and user base has posed challenges to the existing centralized publish/subscribe systems for spatio-textual data streams.

In this paper, we propose our distributed publish/subscribe system, called *PS²Stream*, which digests a massive spatio-textual data stream and directs the stream to target users with registered interests. Compared with existing systems, *PS²Stream* achieves a better workload distribution in terms of both minimizing the total amount of workload and balancing the load of workers. To achieve this, we propose a new workload distribution algorithm considering both space and text properties of the data. Additionally, *PS²Stream* supports dynamic load adjustments to adapt to the change of the workload, which makes *PS²Stream* adaptive. Extensive empirical evaluation, on commercial cloud computing platform with real data, validates the superiority of our system design and advantages of our techniques on system performance improvement.

I. INTRODUCTION

Driven by the proliferation of GPS-equipped mobile devices and social media services, data containing both space and text information is being generated at an unprecedented speed. Millions of social-media users, for example, are uploading photos to Instagram with both location and text tags, posting geo-tagged tweets on Twitter, and creating location-aware events in Facebook, using their smart phones. Such user generated content arrives in a streaming manner, and contains valuable information to both individual and business users. On the one hand, individual users may be interested in events in particular regions, and are keen to receive up-to-date messages and photos that originate in the interested regions and are relevant to the events. On the other hand, business users, e.g., Internet advertisers, expect to identify potential customers with certain interest at a particular location, based on their spatio-textual messages, e.g., restaurant diners in a target zone. For both types of users, it is important to find spatio-textual messages satisfying location and textual constraints in real-time, to deliver the service for their business models.

Publish/subscribe systems [1]–[6] provide the basic primitives to support such information processing paradigms, such that *subscribers* register subscription (continuous) queries in the system to catch all messages from *publishers* satisfying the query predicates. In the context of our problem setting, each registered query from the subscribers consists of two components, a space component describing the spatial region of interests and a text component containing a boolean expression of keywords for matching. A message matches a query, if its location lies in the interested region and its textual content satisfies the boolean expression of the query. When a massive spatio-textual data stream floods on the Internet, the publish/subscribe system filters the messages and routes the matching messages to subscribers in real-time.

Existing publish/subscribe systems are capable of handling subscriptions on a spatio-textual data stream at moderate rates, by utilizing indexing techniques tailored for spatio-textual data on a centralized server [7]–[13]. With the growth of spatio-textual messages on social media and registered queries, the computation workload for publish/subscribe systems is quickly increasing, which exceeds the capacity of a single server. This calls for a distributed solution to building a publish/subscribe system over the spatio-textual data stream. Moreover, due to the dynamic nature of the social media, the workload of publish/subscribe systems varies dramatically over time. The system is expected to conduct dynamic load adjustments with small migration cost to fit for the change in the workload.

In this paper, we present our system, called *PS²Stream*, as a scalable distributed system over a spatio-textual data stream, to enable real-time publish/subscribe services on top of a cluster of servers. The system design objective is to accomplish a maximal processing throughput, minimal latency, and ignorable migration cost. To accomplish these goals, we propose a handful of new techniques. Firstly, we propose a new workload partitioning strategy, utilizing both text and space properties of the data, to facilitate distributing the workload with aims of minimizing the total amount of workload and load balancing of workers. Different from existing distributed systems, we need to distribute both spatio-textual data stream and subscription query stream, and different ways of workload distribution will result in different amounts of workload. Secondly, we discuss dynamic load adjustment approaches over our system architecture, in both local and global manners, to

minimize both the overhead of workload reassignment and the total amount of workload in the scenario of workload changing. The main contributions of our work are summarized as follows:

1) We propose a new workload partitioning algorithm that considers both space and text properties of the data for the workload distribution. Our system is the first to consider the optimal workload partitioning problem for distributing both spatio-textual data stream and subscription query stream, which aims at minimizing the total amount of workload, and balancing the load of workers.

2) We propose efficient dynamic load adjustment algorithms to adjust the load of workers in the scenario of workload changing. By considering both space and text properties of the data, our dynamic load adjustment algorithms achieve features of reducing the total amount of workload and invoking small migration cost. To the best of our knowledge, our system is the first to support such dynamic load adjustments.

3) We conduct extensive experiments on Amazon EC2 with a real-life spatio-textual data stream. The results demonstrate that: i) our distribution framework considering both space and text properties performs better than the baselines utilizing either space or text property only, and our system achieves excellent performance on both processing throughput and query response latency; ii) our dynamic load adjustment approaches improve the performance of the system and invoke small migration cost.

II. RELATED WORK

Spatial-Keyword Publish/Subscribe. The problem of building publish/subscribe systems over a spatio-textual data stream on a centralized server has been studied recently [7]–[13]. They focus on developing new indexes to speed up the matching between spatio-textual objects and the spatial-keyword continuous queries. Specifically, Chen *et al* [10] present an indexing structure called IQ-tree. Li *et al* [7] propose an R-tree based indexing structure. Wang *et al* [9] propose to adaptively group the continuous queries using keyword partitions or space partitions based on a cost model. The problem of similarity based spatial-keyword publish/subscribe is also studied [8], [13], which determines the matching between a subscription and a published spatio-textual object based on a similarity score.

In theory, these studies are complementary to our work as these complicated index structures could be employed in workers of our system. However, it would be expensive to jointly maintain and migrate these index structures across multiple workers, which is required in our system.

Distributed Content based Publish/Subscribe. Our work is related to content based subscription queries. Many works [1]–[6] exist for distributed content-based publish/subscribe systems over a wide-area network. They aim at finding an optimal assignment of the subscribers to message brokers with some performance criterion such as the latency [2]. Another line of works [14], [15] consider deploying the publish/subscribe systems on a P2P network, which focus on minimizing the communication cost. E-StreamHub [16] is a distributed content based publish/subscribe system deployed on a cluster of local servers. However, these systems do not provide supports for

the spatio-textual data, which is essential to our problem, and we do not see any sensible way to deploy these systems for our problem.

Our system differs from these systems in two aspects. Firstly, our system is based on the optimal workload partitioning problem, which aims at both of minimizing the total amount of workload and balancing the load of workers in distributing the workload of both subscription query stream and spatio-textual data stream. Secondly, we support dynamic load adjustments by considering minimizing both total amount of workload and migration cost.

Distributed Systems for Spatial Data. Many distributed systems [17]–[21] have been proposed for large-scale spatial data. Most of them use existing spatial indexes to partition the data to different servers, e.g., SpatialHadoop [18] uses a grid index and a R-tree, and MD-HBase [20] uses a kd-tree and a Quad tree. AQWA [21] uses a kd-tree to partition the spatial data, which aims at minimizing the querying cost based on a query workload. Our system differs from these systems in two ways. Firstly, we consider the workload composed of updating highly dynamic subscription queries and processing spatio-textual objects, while they consider processing disposable queries on a static spatial data set. Secondly, our system design objectives are different. We consider minimizing the total amount of workload and load balancing of workers in distributing workload, as well as dynamic load adjustments. We evaluate for the first time using the data partitioning strategies in these systems for our problem in our experiments.

Distributed Data Stream Processing. The distributed data stream processing systems are also related to our work. A common scheme is to split a continuous query into multiple operators and distribute them to a cluster of servers. NiagaraCQ [22] detects the commonality of different continuous queries and combines their common operators. PSoup [23] improves the processing efficiency of stateful operators such as join and aggregate. Subsequent systems like Aurora* [24] and Borealis [25] further improve the scalability of the data stream processing systems by allowing the same operator being executed by multiple servers. These systems lack support for processing the spatio-textual data stream and their proposed methods do not apply for our system.

A recent demonstration system, called Tornado [26], is presented for indexing spatio-textual data streams to handle querying spatio-textual data streams. However, it does not handle the stream of subscription queries. Furthermore, unlike our system, Tornado does not consider the optimal workload partitioning problem, and it does not support dynamic load adjustments which reduce the total amount of workload with small migration cost. Tornado uses a kd-tree to partition the workload, which has been included in our evaluation of the baselines.

III. PROBLEM STATEMENT AND SYSTEM ARCHITECTURE

A. Problem Statement

Spatio-Textual Object: A spatio-textual object is defined as $o = \langle text, loc \rangle$, where $o.text$ is the textual content of object o and $o.loc$ is the geographical coordinate, i.e., latitude and longitude, of object o .

In this work, we consider a stream of spatio-textual objects, such as geo-tagged tweets, check-ins in Foursquare, geo-tagged posts in Instagram, etc. We aim at building a distributed publish/subscribe system over a stream of spatio-textual objects.

Users may express their interests on the spatio-textual objects with subscription queries. Following the previous works [7], [9], [10], each subscription query contains a Boolean keyword expression and a region. If a new spatio-textual object falls in the specified region and satisfies the Boolean keyword expression specified by a subscription query, the object will be pushed to the user who submits the query. A subscription query is valid until the user drops it. We next present the spatio-textual subscription query.

Spatio-Textual Subscription (STS) Query: A Spatio-Textual Subscription (STS) Query is defined as $q = \langle K, R \rangle$, where $q.K$ is a set of query keywords connected by AND or OR operators, and $q.R$ denotes a rectangle region.

A spatio-textual object o is a result of an STS query q if $o.text$ satisfies the boolean expression of $q.K$ and $o.loc$ locates inside $q.R$.

The large number of STS queries and high arrival rate of spatio-textual objects call for a distributed solution. We build our system on a cluster of servers with several servers playing the role of *dispatchers*, which distribute the workload to other servers. The workload to our system includes the insertions and deletions of STS queries, and the matching operations between STS queries and spatio-textual objects.

Problem Statement We aim at building a distributed publish/subscribe system over a spatio-textual data stream. We expect the system to have the features that (i) the system can achieve a high throughput, (ii) each tuple can be processed in real-time, i.e., low latency, and (iii) the system can dynamically adjust the load of servers when the workload changes.

B. System Architecture

We present a distributed Publish/Subscribe system for handling subscription queries over a Spatio-textual data stream, and we call the system **PS²Stream**. There are three components in **PS²Stream**: dispatcher, worker and merger, which are shown in Figure 1. The dispatcher takes as input a stream of spatio-textual objects, and it receives two types of requests from users, namely submitting new subscriptions and dropping existing subscriptions, which correspond to two operations, i.e., query insertions and deletions, respectively. The dispatcher assumes the responsibility of distributing spatio-textual objects and query insertions/deletions. The worker accepts the workload sent from the dispatcher and conducts the following operations accordingly:

- (1) *Query insertion*: On receiving a new STS query, the worker inserts it into an in-memory index maintained in the worker;
- (2) *Query deletion*: On receiving the indication of deleting an existing STS query, the worker removes the query from the index;
- (3) *Matching a spatio-textual object*: On receiving a spatio-textual object, the worker checks whether the object can be

a match for any STS query stored in the worker. If yes, the matching result is forwarded to the merger.

The merger assumes the responsibilities of removing duplicated matching results produced by workers and sending the results to corresponding users.

Figure 1: Architecture of PS²Stream

IV. HYBRID PARTITIONING

In this section, we present the design of our **PS²Stream** system. We define a new workload distribution problem in Section IV-A, and provide a solution to it in Section IV-B. We introduce the index structure adopted in dispatchers in Section IV-C. Finally, we discuss how workers process their workload in Section IV-D.

A. Definition

We propose **PS²Stream** for processing a stream of subscription queries over a spatio-textual data stream. Different from most existing distributed systems, we need to distribute the workload of both updating highly dynamic STS queries and processing spatio-textual objects. In our problem, different ways of workload distribution will result in different amounts of workload. During the workload distribution, the following two factors are important for the system. One is the total amount of workload distributed to the workers, and the other one is the load balancing of workers.

Next, we first define the load of one worker. Based on it, we propose our Optimal Workload Partitioning problem. To the best of our knowledge, our definition of the workload distribution problem is the first for distributing spatio-textual data that considers minimizing the total amount of workload while balancing the load of workers.

Definition 1: Load of One Worker: Given a time period, the load of a worker w_i for the period is computed by

$$L_i = c_1 \cdot |O_i| \cdot |Q_i^i| + c_2 \cdot |O_i| + c_3 \cdot |Q_i^i| + c_4 \cdot |Q_i^d|,$$

where O_i is the set of spatio-textual objects sent to w_i , Q_i^i is the set of STS query insertion requests sent to w_i , and Q_i^d is the set of STS query deletion requests sent to w_i in the given time period. Here c_1 is the average cost of checking whether a spatio-textual object matches an STS query, c_2 is the average cost of handling one object, c_3 is the average cost of handling one STS query insertion request, and c_4 is the average cost of handling one STS query deletion request. ■

Definition 2: Optimal Workload Partitioning Problem: Given a set of spatio-textual objects O , a set of STS query

Figure 2: Two regions with different data distributions.

insertion requests Q^i and a set of STS query deletion requests Q^d , let S denote the space they reside in and T denote the set of terms appearing in them. We aim to partition S into m subspaces $S_1, S_2, \dots, S_m$ and partition T into m subsets $T_1, T_2, \dots, T_m$, and assign each pair (S_i, T_i) to one worker, where m is the number of workers. An object o is sent to worker w_i only if $o.loc \in S_i$ and $o.text \cap T_i \neq \emptyset$. An insertion/deletion request of an STS query q is sent to worker w_i only if $q.R \cap S_i \neq \emptyset$ and $q.K \cap T_i \neq \emptyset$. The Optimal Workload Partitioning problem is to find a partition such that $\sum_{i=1}^m L_i$ is minimized, subject to the constraint that $\forall i \neq j$ and $L_i \geq L_j, L_i/L_j \leq \sigma$, where σ is a small constant value being larger than 1. ■

Note that the constraint in the problem statement is to achieve load balancing.

Theorem 1: The Optimal Workload Partitioning problem is NP-hard.

Proof: We prove this by a reduction from the Number Partitioning problem. The proof can be found in the longer version of this paper [27]. ■

B. Workload Partitioning

Due to the hardness of the Optimal Workload Partitioning problem, we investigate heuristic algorithms to partition the workload. The text property and the space property of the data inspire us to partition the workload by the text property or the space property. Previous works [28], [29] on constructing a distributed text retrieval system or a distributed information dissemination system use text-partitioning to partition the workload. However, they focus on the communication cost and load balancing. Other works on constructing a distributed system for a static spatial data set [17]–[21] and for a spatio-textual data stream [26], use space-partitioning to partition the workload, and they do not consider minimizing the total amount of workload and the load balancing collectively.

Such partitioning algorithms based on an unanimous scheme perform poorly in minimizing the total amount of workload when the data distributions among different regions are quite different. As shown in Figure 2, space-partitioning performs well in region r_2 where the objects and queries are well spread, but it does not fit for region r_1 as the ranges of

queries in r_1 are large and clustered. Similarly, text-partitioning is good for r_1 but bad for r_2 .

To better solve the Optimal Workload Partitioning problem, we propose a new partitioning algorithm that utilizes both text and space properties of the data. Our hybrid partitioning algorithm decomposes the workload into a set of units by wisely using the text property or the space property in different regions, with the goal of minimizing the total amount of workload after the decomposition. Additionally, we distribute those units to workers, so that the load balancing constraint can be satisfied.

Algorithm overview. The core idea of our algorithm is to first use space-partitioning to identify the subspaces that have large differences in the text distribution between objects and queries, and then for each subspace, we choose between space-partitioning and text-partitioning to minimize the total amount of workload.

The algorithm can be divided into two phases. In the first phase, the algorithm divides the space into subspaces based on the text similarity between objects and queries. The purpose is to identify those subspaces where the text-partitioning would perform better. To achieve this, for a subspace where the text similarity between objects and queries is smaller than a threshold, we check whether partitioning it will result in new subspace(s) where the text similarity becomes smaller. If yes, we partition the subspace recursively and perform the same checking for each new subspace. At the end of this phase, we obtain two types of subspaces represented by N_s and N_t , respectively. The subspaces in N_t have small text similarity between objects and queries, and we partition them using text-partitioning only in the second phase. For the subspaces in N_s , we compare the workloads produced by space-partitioning and text-partitioning, respectively, and select the strategy resulting in a smaller workload. This comparison is necessary as when the query ranges are very large, using space-partitioning will invoke queries being duplicated to multiple workers, which may result in a larger workload than using text-partitioning even though the objects and queries have large text similarity. In the second phase, if the number of partitions is smaller than the number of workers, we further partition the workload in N_s and N_t . To minimize the total amount of workload, we design a dynamic programming function to find the optimal number of partitions for each subspace in N_s and N_t . We then check whether the load balancing constraint can be satisfied. If no, we recursively further partition one subspace until the load balancing constraint can be satisfied.

The output of our workload partitioning procedure is an index structure named kd^t -tree, which is a kd-tree with some leaf nodes being further partitioned by the text property. Figure 3 shows an example of a kd^t -tree, which distributes the workload to 4 workers: the set of leaf nodes is divided into 4 disjoint subsets and each subset is assigned to one worker randomly, e.g., $\{N_{11}\}$ is assigned to worker w_1 and $\{N_{12}, N_{31}\}$ is assigned to worker w_2 . For an object and query locating inside or overlapping with the space range of node N_2 or N_4 , it is sent to worker w_3 or w_4 without checking the textual content. When an object or query locates inside or overlaps with the space range of node N_1 , its textual content is checked and it is sent to worker w_1 (resp. w_2) if it contains

Figure 3: Example of a kd^t -tree

terms in T_1 (resp. T_2). Workers w_2 , w_3 , and w_4 also receive objects and queries that locate inside or overlap with the space range of node N_3 and contain terms in T'_1 , T'_2 and T'_3 , respectively.

Algorithm 1: Partition Workload

Input: The number of workers m , Load balance constraint threshold σ , Text similarity threshold δ , A set of objects O , A set of queries Q ;
Output: The root node of a kd^t -tree with each leaf node assigned a number in $\{1, 2, \dots, m\}$

- 1 Initialize the kd^t -tree structure with a root node $n_r(O, Q)$;
- 2 $N_u \leftarrow \{n_r\}$, $N_t \leftarrow \emptyset$, $N_s \leftarrow \emptyset$;
- 3 **while** $|N_u| > 0$ **do**
- 4 Pop one node n from N_u ;
- 5 **if** $\text{sim}_t(O_n, Q_n) \geq \delta$ **then**
- 6 $N_s \leftarrow N_s \cup \{n\}$;
- 7 **else**
- 8 $\{n_1, n_2\} \leftarrow$ partition n into 2 nodes in the direction that minimizes
 $\alpha = \min\{\text{sim}_t(O_{n_1}, Q_{n_1}), \text{sim}_t(O_{n_2}, Q_{n_2})\}$;
- 9 **if** $|\alpha - \text{sim}_t(O_n, Q_n)| \approx 0$ **then**
- 10 $N_t \leftarrow N_t \cup \{n\}$;
- 11 **else**
- 12 $N_u \leftarrow N_u \cup \{n_1, n_2\}$;
- 13 **if** $|N_t| + |N_s| < m$ **then**
- 14 $A \leftarrow \text{ComputeNumberPartitions}(N_t, N_s, m)$;
- 15 **for each node** n **in** $N_t \cup N_s$ **do**
- 16 $\text{PartitionNode}(n, N_t, N_s, A[n])$;
- 17 **while True do**
- 18 $\text{MergeNodesIntoPartitions}(N_t, N_s, m)$;
- 19 $L_{\max} \leftarrow$ the maximum load value among all partitions;
- 20 $L_{\min} \leftarrow$ the minimum load value among all partitions;
- 21 **if** $L_{\max}/L_{\min} \leq \sigma$ **then**
- 22 **break**;
- 23 **else**
- 24 $n \leftarrow$ the node having the largest load value in $N_t \cup N_s$;
- 25 $\text{PartitionNode}(n, N_t, N_s, 2)$;
- 26 **if** $|N_t| + |N_s| \geq \theta$ **then**
- 27 **break**;
- 28 **return** n_r ;

Algorithm 1 presents the pseudo code of our workload partitioning algorithm. The algorithm first initializes a root node

$n_r(O, Q)$ of the kd^t -tree and puts it into N_u (lines 1–2). The variable N_u denotes the set of nodes that have not been decided to put into N_t or N_s . A while loop is conducted to compute N_t and N_s (lines 3–12). At each iteration, it pops one node n from N_u (line 4) and computes the text similarity between the objects and queries in n . We use cosine similarity in our algorithm. If the text similarity is larger than a threshold δ , we consider node n as being not suitable for text-partitioning, and add it into N_s (line 6). Otherwise, we split node n in either x-dimension or y-dimension as the normal kd -tree does, except that we prefer the direction resulting in a smaller text similarity α between objects and queries in the new subspaces (line 8). If the difference between $\text{sim}_t(O_n, Q_n)$ and α is minor, we consider node n as being consistent in text similarity of the objects and queries in it and add n into N_t (lines 9–10). Otherwise, we add the new nodes n_1 and n_2 into N_u (line 12). In the second phase, the algorithm proceeds to check the number of nodes in N_t and N_s (line 13). If the number of nodes is smaller than m , i.e., the number of workers, a dynamic programming function `ComputeNumberPartitions` is called to compute the number of partitions for each node (lines 13–14). It is for minimizing the total amount of workload by determining an optimal number of partitions for each node. The function `PartitionNode` is then called to partition the nodes (lines 15–16). After that, the algorithm recursively checks whether the load balancing constraint can be satisfied (lines 18–22) and partitions the node having the largest load, until the load balancing constraint is satisfied or the number of nodes reaches a threshold θ , where θ is a threshold of the maximum number of nodes (lines 24–27). In the end, the algorithm returns the root node of the kd^t -tree (line 28).

Computing the number of partitions. For the purpose of minimizing the total amount of workload, a dynamic programming function `ComputeNumberPartitions` is called to compute the number of partitions for each node in N_t and N_s when $|N_t| + |N_s| < m$ (lines 13–14 of Algorithm 1). The core of the function is updating $L[i, j]$, where $L[i, j]$ denotes the minimum sum of loads after partitioning the first i nodes into a total number of j partitions. Let $C[i, k]$ denote the sum of loads after partitioning the i th node into k partitions. The value of $C[i, k]$ can be obtained by calling a function similar to `PartitionNode` except that the node is not really partitioned. Then we can compute $L[i, j]$ by using the equation

$$L[i, j] \leftarrow \min_{1 \leq k \leq j-1} \{L[i-1, j-k] + C[i, k]\}.$$

Based on the arrays L and C , the function can construct a map which maps from each node to the corresponding number of partitions.

Node partitioning. In Algorithm 1, we need to partition a node, if the number of nodes is smaller than the number of workers (lines 13–16), or the load balancing constraint cannot be satisfied (lines 23–27). Function `PartitionNode` takes as input the node n to be split, the set of nodes N_t , the set of nodes N_s , and the number of splits p . It outputs p new nodes by partitioning node n . If node n belongs to N_t , it indicates that the text similarity between objects and queries in n is small, and we partition n into p splits using text-partitioning. Otherwise, we compare the workloads produced by using text-partitioning or space-partitioning to partition n , and select the strategy resulting in a smaller workload.

Figure 4: A grid^t index transformed from a kd^t -tree

Node merging. When the number of leaf nodes in the kd^t -tree exceeds the number of workers m , we call function `MergeNodesIntoPartitions` to divide those leaf nodes into m partitions (line 18 of Algorithm 1). The function first sorts the leaf nodes in descending order of their loads. In that order, for each leaf node n , it finds the partition $part$ such that adding n to $part$ will result in a minimum load increase. If assigning n to $part$ does not increase the load balancing factor (i.e., L_{max}/L_{min}), n is assigned to $part$. Otherwise, n is assigned to the partition that has currently the smallest load.

C. Index Structure on Dispatchers

The kd^t -tree can be used in the dispatcher to distribute the workload. For each spatio-textual object or updating request of an STS queries, the dispatcher obtains the corresponding leaf node(s) of the kd^t -tree by traversing from the root node. The procedure takes $O(\log(m))$ time, which may overload the dispatcher when the arrival speeds of objects and updating requests of queries are very fast.

To alleviate the burden of the dispatcher, instead of maintaining a kd^t -tree, we conduct the workload distribution using a grid^t index with each grid cell containing two hash maps H_1 and H_2 . H_1 maps from terms in the complete term set T to the worker ids, and H_2 maps from terms in STS queries to the worker ids. To distribute a spatio-textual object, the dispatcher first obtains the cell containing the object, and uses H_2 to find the destination worker(s). The object can be discarded if it contains no terms in H_2 . To distribute an updating request of an STS query (note that the request contains complete information of the STS query), if the query keywords are connected by AND operators only, the dispatcher obtains the cells that query overlaps with, and looks up H_1 using the least frequent keyword in each obtained cell to find the destination worker(s). H_2 is updated correspondingly if it does not contain the keyword. For the query containing OR operators, similar operations are conducted except that the set of the least frequent keywords in each conjunctive norm form are used to find the destination worker(s). The grid^t index can be built from the kd^t -tree. The granularity of the cell is decided by the leaf nodes of the kd^t -tree. Figure 4 shows an example of the relation between the grid^t index and the kd^t -tree in Figure 3.

D. Query Processing on Workers

Each worker maintains an in-memory index structure to organize the STS queries for accelerating the matching of incoming spatio-textual objects. The index should be efficient in both matching STS queries and spatio-textual objects and

updating of STS queries. In our system, we adopt an index structure named **Grid-Inverted-Index** (or GI^2 for short) [30]. We choose GI^2 due to its efficiency in construction and maintaining, which is important for processing a dynamic workload like the data stream. Note that our system can be extended to adopt other index structures.

After dividing the STS queries by the cells they overlap with, GI^2 constructs an inverted index for each cell to organize the STS queries. For the query containing AND operators only, it is appended to the inverted list of the least frequent keyword. For the query containing OR operators, it is appended to the inverted lists of the least frequent keywords in each conjunctive norm form. The granularity of the cells can be determined empirically in experiments.

When processing a spatio-textual object o , the worker first looks up which cell contains o and then checks the associated inverted list in the cell for each distinct term in $o.\text{text}$ to see whether o can be matched by any STS query q_i in that list. If there is a match, object o is a result of q_i .

To delete an STS query from GI^2 , we adopt the *lazy-deletion* strategy. Specifically, we do not delete an STS query immediately after receiving the deletion request. Instead, we record the ids of the queries to be deleted into a hash table and remove the queries during the object matching procedure. In particular, while traversing an inverted list, we check whether the id of each query in the inverted list appears in the hash table. If yes, the query can be deleted.

V. DYNAMIC LOAD ADJUSTMENT

The dynamic property of the data stream gives rise to a changing workload. The load of workers may change gradually over time. To handle this, we conduct two types of load adjustments: local load adjustment and global load adjustment.

A. Local Load Adjustment

When the dispatcher detects that the load balancing constraint is violated, it notifies the most loaded worker w_o to transfer part of its workload to the least loaded worker w_l . We expect the load adjustment procedure having the following two features: 1.) The total amount of workload can be reduced after the load adjustment. 2.) The migration cost is small so that the load adjustment can be conducted efficiently. Existing systems [16], [26], [29], [31] do not support dynamic load adjustments that consider both features. To achieve the goal, we propose new load adjustment algorithms.

Algorithm overview. We adjust the workload by migrating STS queries. The queries are migrated in the unit of one cell in the grid^t index. Our load adjustment algorithm is composed of two phases. In the first phase, we check whether some cells in w_o and w_l can be split or merged so that the total amount of workload can be reduced. We conduct related migration operations if such cells exist. In the second phase, if the load balancing constraint is still violated, we continue to select a set of cells in w_o to be migrated to w_l , with the goals of minimizing the migration cost and the load balancing constraint can be satisfied.

Phase I. We first check the p most loaded cells G_p in w_o , where p is a small parameter. For each cell g_s not using text-partitioning in G_p , we do the following checking: after using text-partitioning to partition g_s into two new cells g_1 and g_2 , whether the total amount of workload can be reduced if g_1 or g_2 is migrated to w_l . If yes, we conduct the text-partitioning on g_s and migrate the cell having a smaller size between g_1 and g_2 to w_l . For each cell g_t using text-partitioning in G_p , if there exists a cell g'_t in w_l which shares the same space region as g_t has, we check whether migrating g_t to w_l and merging g_t and g'_t can reduce the total load. If yes, we conduct the migration.

Phase II. If the load balancing constraint is still violated, we conduct the second phase, which solves the Minimum Cost Migration problem. Definition 3 defines the load of a cell, and Definition 4 presents the definition of the Minimum Cost Migration problem.

Definition 3: Given a time period, the load of a cell g for the period is computed by

$$L_g = n_o \cdot n_q,$$

where n_o is the number of spatio-textual objects falling in cell g and n_q is the average number of STS queries stored in cell g in the given time period. ■

Definition 4: Minimum Cost Migration: Let τ denote the amount of load to be migrated. Consider a worker w_o with the set of cells G_o . The minimum cost migration problem is to compute a set of cells G_s to be migrated from worker w_o .

$$G_s = \arg \min_{G_s} \sum_{g \in G_s} S_g \quad \text{s.t.} \\ \sum_{g \in G_s} L_g \geq \tau,$$

where S_g is the total size of the queries in cell g . ■

Theorem 2: The Minimum Cost Migration problem is NP-hard.

Proof: The proof can be found in the longer version of this paper [27]. ■

1) Dynamic Programming Algorithm: We develop a dynamic programming algorithm to solve the Minimum Cost Migration problem. Let $H(i, j)$ denote a subset of cells $\{g_1, g_2, \dots, g_i\}$ such that the total size of its cells is no larger than j and the total load is maximum

$$H(i, j) = \arg \max_{H \subseteq \{g_1, g_2, \dots, g_i\} \wedge \sum_{g \in H} S_g \leq j} \sum_{g \in H} L_g.$$

Let $A(i, j)$ denote the load of $H(i, j)$. We compute $A(i, j)$ for all $i \in \{1, 2, \dots, n\}$ and $j \in (0, P]$, where P is an upper bound of the minimum migration cost. Following shows the computation of $A(i, j)$.

$$A(i, j) = \begin{cases} A(i-1, j) & \text{if } j \leq S_{g_i} \\ \max\{A(i-1, j), A(i-1, j - S_{g_i}) + L_{g_i}\} & \text{otherwise} \end{cases}.$$

The time complexity of the dynamic programming algorithm is $O(nP)$ and it takes long running time when the size of cells is large. Moreover, it needs $O(nP)$ memory space, which is expensive. To address these issues, we next propose Algorithm GR, which is more efficient and requires much less memory space.

2) Greedy Algorithm: We introduce the basic idea of the proposed greedy algorithm, denoted by GR. For each cell g , we compute S_g/L_g , which indicates the relative cost of migrating the cell. It is expected that migrating a cell with a smaller relative cost may have smaller migration cost. We scan cells in G_o in ascending order of their relative costs to find a set of cells to be migrated. For each cell we scan, if the total migrated load after including it in the result is still less than τ , the cell will be included in the result, and we mark it by “GS”; otherwise, the cell becomes a candidate cell to be included in the result, and we mark it by “GL”. The candidate cell will be included in the result, if in the subsequent scan we cannot find cells that can meet the load requirement and incur smaller migration cost.

We illustrate this process with Figure 5, where the cells are sorted by their relative costs, and they are marked either by “GS” or “GL”. For any value of t , we have $\sum_{i=1}^t \sum_{g \in GS_i} L_g < \tau$; and $\forall g' \in GL_t, \sum_{i=1}^t \sum_{g \in GS_i} L_g + L_{g'} \geq \tau$. For any value of t , the set $GS_1 \cup GS_2 \cup \dots \cup GS_t \cup \{g'\}$ is a candidate solution to the Minimum Cost Migration problem.

$$\overbrace{g_1, \dots, g_{k_1}}^{GS_1}, \overbrace{g_{k_1+1}, \dots, g_{k_2-1}}^{GL_1}, \overbrace{g_{k_2}, \dots, g_{k_3}}^{GS_2}, \overbrace{g_{k_3+1}, \dots, g_{k_4-1}}^{GL_2}, \\ \dots, \overbrace{g_{k_{2i-2}}, \dots, g_{k_{2i-1}}}^{GS_i}, \overbrace{g_{k_{2i-1}+1}, \dots, g_{k_{2i}-1}}^{GL_i}, \dots$$

Figure 5: Greedy algorithm for computing the set of cells to be migrated

During the scan, we find the value for t and the cell g' such that $GS_1 \cup GS_2 \cup \dots \cup GS_t \cup \{g'\}$ has the minimum migration cost among all the candidate solutions that we have scanned.

B. Global Load Adjustment

The performance of current workload partitioning strategy will degrade when there exists a significant change in the data distribution. To handle this, we periodically check whether a workload repartitioning is necessary on a recent sample of data. If yes, we conduct a workload repartitioning using the algorithm presented in Section IV-B. To avoid the large migration cost after the workload repartitioning, we make a temporary compromise on the system performance by maintaining two workload distribution strategies: one for the old STS queries before the workload repartitioning, and the other one for the new STS queries. When the amount of old STS queries becomes small, we conduct the migration and stop the old workload distribution strategy. We set a long period for the checking, e.g., once per day, which is reasonable as the data distribution usually changes slowly.

VI. EXPERIMENTAL EVALUATION

A. Experimental Setup

We build **PS²Stream** on top of Storm¹, an open-source stream processing engine. Our system is deployed on the Amazon EC2 platform using a cluster of 32 c3.large instances. Each c3.large instance has 2 vCPUs running Intel Xeon E5-2680 at 2.5GHz and 3.75 GB RAM. The network performance is moderate and no further information about the bandwidth is provided by Amazon. We rent another high performance m4.2xlarge instance (which is equipped with 8 vCPUs and 32 GB RAM, and has high IO performance) to emit streams of spatio-textual objects and STS queries to our system.

Datasets and STS Queries. Our experiments are conducted on two datasets, which are TWEETS-US and TWEETS-UK. TWEETS-US consists of 280 million spatio-textual tweets in America and TWEETS-UK consists of 58 million spatio-textual tweets in Britain. Since we do not have real-life STS queries, we synthesize queries based on tweets. The number of keywords is a random number ranging from 1 to 3 and the keywords are connected by either AND or OR operators. The query range is a square and the center is randomly selected from the locations of tweets in TWEETS-US or TWEETS-UK. We synthesize four groups of STS queries as follows.

(1) STS-US-Q1, STS-US-Q2: The two groups of queries are synthesized based on TWEETS-US. For STS-US-Q1, the side lengths of the rectangle are randomly assigned between 1km and 50km and the keywords share the same distribution as the terms in TWEETS-US, i.e., the keywords in queries satisfy the power-law distribution. For STS-US-Q2, the side lengths of the rectangle are randomly assigned between 1km and 100km and the queries contain at least one keyword that is not in the top 1% most frequent terms in TWEETS-US.

(2) STS-UK-Q1, STS-UK-Q2: The two groups of queries are synthesized based on TWEETS-UK as we do for TWEETS-US.

Workload. The ratio of processing a spatio-textual tweet to inserting or deleting an STS query is approximately 5. The arrival speeds of requests for inserting an STS query and deleting an STS query are equivalent. It indicates that the number of STS queries in the system will be stable eventually. We use a parameter μ to control the number of STS queries. We achieve this by using a Gaussian distribution $\mathcal{N}(\mu, \sigma^2)$ to determine the number of newly arrived STS queries between inserting an STS query and deleting it. We set different values of μ in the experiments with $\sigma = 0.2\mu$.

B. Evaluating Baselines

We evaluate the performance of baseline workload distribution algorithms.

Text-Partitioning. *Text-partitioning* algorithms divide the lexicon into several partitions, each of which is assigned to one worker, and distribute the workload based on the textual content of objects/queries. We implement the following three *text-partitioning* algorithms: (1) Algorithm *frequency-based partitioning* uses the frequency values of the terms to do the

Figure 6: Throughput Comparison (Baselines)

partitioning. (2) Algorithm *hypergraph-based partitioning* [28] constructs a hypergraph based on the co-occurrence of terms and partitions that hypergraph. (3) Algorithm *metric-based partitioning* [29] uses a metric function to do the partitioning.

Space-Partitioning. *Space-partitioning* algorithms divide the space into several partitions, each of which is assigned to one worker, and distribute the workload based on the spatial content of objects/queries. We implement the following three *space-partitioning* algorithms: (1) Algorithm *grid partitioning* [18] represents the space as a set of cells and then partitions that set of cells. We set its granularity as $2^6 \times 2^6$, as it performs best under that granularity. (2) Algorithm *kd-tree partitioning* [21], [26] constructs a kd-tree to do the partitioning, of which each leaf node represents one partition. We transform the kd-tree to a grid index to accelerate the workload distribution in the dispatchers. (3) Algorithm *R-tree partitioning* [18] constructs a R-tree to do the partitioning, and then partitions the set of leaf nodes.

Note that this is the first work on empirically comparing the performance of these algorithms for handling a spatio-textual data stream. These experiments are conducted using 4 dispatchers and 8 workers. The granularity of GI² index of workers in all algorithms is set as $2^6 \times 2^6$ for fairness.

Throughput The throughput is the average number of tuples (i.e., including both query insertion/deletion and object matching) being processed per second when the processing capacity of the system is reached (i.e., the input speed of the data stream is tuned to be closed to the maximum number of tuples that the system can handle in unit time). We set $\mu = 5M$ for Q1 queries, and set $\mu = 10M$ for Q2 queries.

Figure 6 shows the throughputs of the baselines. For Q1 queries, *space-partitioning* algorithms perform better than *text-partitioning* algorithms. For example, for STS-US-Q1, the throughput of *grid partitioning* is 160,659 tuples/second, and the throughput of *frequency-based partitioning* is 76,921 tuples/second. The reason is that *space-partitioning* algorithms impose a smaller amount of workload on the workers: the *text-partitioning* algorithms send each object to multiple workers as the keywords in Q1 queries are frequent in objects.

¹<http://storm.apache.org>

Figure 7: Throughput Comparison

Figure 8: Latency Comparison

For Q2 queries, *text-partitioning* algorithms perform better than *space-partitioning* algorithms. For example, for STS-US-Q2, the throughput of *metric-based partitioning* is 267,415 tuples/second, and the throughput of *R-tree partitioning* is 181,542 tuples/second. The reason is that the Q2 queries have larger query ranges and less frequent keywords, and therefore *text-partitioning* algorithms impose a smaller workload on the workers.

Overall, *metric-based partitioning* performs the best among *text-partitioning* algorithms, and *kd-tree partitioning* performs the best among *space-partitioning* algorithms. Therefore, we select *metric-based partitioning* and *kd-tree partitioning* for further evaluation.

C. Evaluating Hybrid Partitioning

We evaluate the performance of our hybrid partitioning algorithm in this set of experiments. To simulate the situation that users in different regions have different preferences over the spatio-textual objects, we synthesize two new groups of STS queries: STS-US-Q3 and STS-UK-Q3. STS-US-Q3 is synthesized by partitioning the spatial range of the USA into 100 regions of equal size, and for each region we use STS-US-Q1 or STS-US-Q2. Similar procedure is conducted to synthesize STS-UK-Q3.

We first evaluate *hybrid partitioning* using 4 dispatchers and 8 workers. Then we evaluate the scalability of *hybrid partitioning* using 4 dispatchers while varying the number of workers.

Throughput Figure 7 shows the throughputs of *hybrid partitioning*, *metric-based partitioning* and *kd-tree partitioning*. Algorithm *hybrid partitioning* shows the overall best performance regardless of the data distributions. In Figure 7(c), *hybrid partitioning* outperforms *metric-based partitioning* and *kd-tree partitioning* by 30% in terms of throughput. In Figure 7(a),

for STS-US-Q1, *hybrid partitioning* slightly outperforms *kd-tree partitioning*. Both *hybrid partitioning* and *kd-tree partitioning* perform much better than *metric-based partitioning*. It is because that most keywords in STS-US-Q1 are frequent among objects, which results in a larger workload using *text-partitioning*. This can also explain why *hybrid partitioning* and *kd-tree partitioning* have similar performance. Because *hybrid partitioning* wisely choose *space-partitioning* to partition the workload in most regions. In Figure 7(b), *hybrid partitioning* and *metric-based partitioning* have a larger throughput than *kd-tree partitioning* does. This can be explained in two aspects. First, the larger query range results in a larger workload for *kd-tree partitioning*. Second, the keywords in STS-US-Q2 and STS-UK-Q2 are less frequent, which improves the performance of *text-partitioning*.

Latency The latency is the average time of each tuple staying in the system. We evaluate all the algorithms using a moderate input speed of the data stream.

Figure 8 shows the latencies of *hybrid partitioning*, *metric-based partitioning* and *kd-tree partitioning*. Algorithm *hybrid partitioning* has smaller latency. The latency of *kd-tree partitioning* is noticeably larger than *hybrid partitioning* and *metric-based partitioning* in Figure 8(b), e.g., 25ms versus 15ms. This is caused by the larger spatial ranges of queries in STS-US-Q2 and STS-UK-Q2. Algorithm *metric-based partitioning* has similar latency values as *hybrid partitioning* does with the exception on STS-UK-Q1, which takes 407ms. This is caused by the poor performance of *text-partitioning* when query keywords are frequent. Figure 8 demonstrates that *hybrid partitioning* achieves small latency for workload processing.

Memory Figure 9 shows the average memory usages of dispatchers of *hybrid partitioning*, *metric-based partitioning* and *kd-tree partitioning*. Overall, the memory usage of all the methods is not large, e.g., less than 1000MB, which is

Figure 9: Memory Comparison (Dispatchers)

Figure 10: Memory Comparison (Workers)

Figure 11: Scalability

acceptable in practice. Algorithm *kd-tree* partitioning uses less memory than *metric-based* partitioning and *hybrid* partitioning. The memory usage of *hybrid* partitioning is the highest for Q2 queries than for Q1 and Q3. This is because that, for Q2 queries, more cells in the grid^t index need to maintain additional text partitioning information.

Figure 10 shows the average memory usages of workers of the three methods. Algorithm *hybrid* partitioning has an overall best performance, which invokes the smallest memory usage in most cases. This is due to the reason that *hybrid* partitioning distributes STS queries to workers considering the data distributions in different regions, which reduces the cases of one STS query being stored in multiple workers. We also observe that all the three methods do not impose large memory requirements on workers.

Scalability Figure 11 shows the throughputs of the evaluated algorithms as we vary the number of workers. As we can see, *hybrid* partitioning exhibits the best performance in most cases, and scales well with the number of workers. Algorithm *metric-based* partitioning has the worst scalability in Figure 11(a) as the keywords in STS-UK-Q1 are frequent, which results in a larger workload on the workers. Figure 11(b) shows that *kd-tree* partitioning has the worst scalability.

D. Evaluating Dynamic Load Adjustment

In this set of experiments, we evaluate the performance of our dynamic load adjustment algorithms. The experiments are conducted on 4 dispatchers and 8 workers. In addition to our proposed dynamic programming (DP) algorithm and greedy (GR) algorithm, we implement two other algorithms for comparison.

SI: It is another greedy algorithm which adds cells into G_s (the set of cells to be migrated), in descending order of their sizes.

RA: It selects the cells to be migrated randomly.

Comparing the running time of selecting cells for migration We run the four algorithms on the same worker (having the same set of cells) and measure their running time for selecting cells for migration, and the results are shown in Figure 12(a) and Figure 13. Note that workers run out of memory when running DP on queries used for the experiments in Figure 13. As shown in Figure 12(a), DP runs significantly longer than the other algorithms. This is due to the high time complexity of DP. As expected, RA is the fastest among all the algorithms. However, both GR and SI are very efficient too. By comparing Figures 13(a) and 13(b), we can see that the running time of

Figure 12: Migration Experiments (#Q=1M, STS-US-Q1)

Figure 13: Average Time of Selecting Cells (STS-US-Q1)

GR, SI and RA does not change with the number of queries. The reason is that their running time is only determined by the number of cells.

Comparing the migration cost We conduct the evaluation by running a workload consisting of processing a sample of spatio-textual tweets in 60 days and the insertions/deletions of STS queries. Since our tweets are only a small sample of real-life tweets, we scale out the workload by reading 4 hours of tweets in every 10 seconds. Note that we utilize the timestamps of tweets for the scale-out. In each workload migration, we measure the size of migrated data and the time in doing migration.

We report the average size of the migrated data and the average time in doing migration in Figure 12(b) and Figure 14. In each figure, the left y-axis shows migration cost and the right shows time. In Figure 12(b), DP and GR incur the smallest migration cost and require the least time for migration. DP requires slightly longer time than GR, since DP needs longer time in selecting the cells to be migrated. GR performs the best as shown in Figure 14. In Figure 14(a), it incurs 30%–40% less migration cost than SI and RA. It takes the least amount of time in doing migration, which is nearly half of the time required by RA and 70% of the time required by SI. Note that the migration cost (resp. migration time) in Figure 14(b) is larger than the migration cost (resp. migration time) in Figures 14(a). This is because that a larger number of STS queries impose heavier load on the system and the size of each cell to be migrated becomes larger too.

Comparing latency To further evaluate the effect of different algorithms on the system, we also measure the processing latency. The results are presented in Figure 12(c) and Figure 15. In Figure 12(c), GR has the smallest side effect on the system, where 80% of tuples are not affected by the migration operations (less than 100ms) and 4% of tuples are significantly delayed (larger than 1 second). DP is second to GR, where 72% of tuples are not affected by the migration operations. It results in a larger fraction of tuples being delayed between

Figure 14: Migration Cost and Time (STS-US-Q1)

Figure 15: Percentage of Latency Values (STS-US-Q1)

Figure 16: The Effect of the Dynamic Load Adjustments

100ms and 1 second. This is due to its long running time in selecting cells to be migrated. SI and RA perform worse: SI results in 10% more tuples being delayed than GR. RA results in 20% more tuples being delayed than GR.

Evaluating the effect of dynamic load adjustments We evaluate the benefits of our dynamic load adjustment algorithms by comparing the throughputs of the system conducting dynamic load adjustments, and the system without conducting dynamic load adjustments. We use our proposed GR algorithm in the dynamic load adjustments. To simulate a workload with varying data distributions over time, we use query set STS-US-Q3 for this experiment, and every interval of 10M queries, the types of queries in 10% of the regions switch between STS-US-Q1 and STS-US-Q2. We set $\mu = 10M$ in this experiment.

Figure 16 shows the experimental results. The system with

dynamic load adjustments outperforms the system without dynamic load adjustments by 26% in terms of throughput. The results demonstrate the effectiveness of our dynamic load adjustment algorithms.

VII. CONCLUSION

In this work, we develop an efficient distributed publish/subscribe system, PS²Stream, over a spatio-textual data stream. We consider the optimal workload partitioning problem and propose a new hybrid partitioning algorithm. We also propose effective dynamic load adjustment approaches to adjust the load of workers in the scenario of workload changing. Our experimental results show that our workload distribution framework performs better than the baselines in both throughput and latency, and our dynamic load adjustment approaches improve the performance of the system with small migration cost.

ACKNOWLEDGEMENT

This work was carried out at the Rapid-Rich Object Search(ROSE) Lab at the Nanyang Technological University, Singapore. The ROSE Lab is supported by the National Research Foundation, Singapore, under its Interactive Digital Media(IDM) Strategic Research Programme. This work is also supported in part by a Tier-1 grant (RG22/15) and a Tier-2 grant (MOE-2016-T2-1-137) awarded by Ministry of Education Singapore, and a grant awarded by Microsoft. Tom Fu and Zhenjie Zhang are supported by the research grant for the Human-Centered Cyber-physical Systems Programme at the Advanced Digital Sciences Center from Singapore's Agency for Science, Technology and Research (A*STAR). They are also partially supported by Science and Technology Planning Project of Guangdong under grant (No. 2015B010131015).

REFERENCES

- [1] F. Cao and J. P. Singh, "Efficient event routing in content-based publish-subscribe service networks," in *INFOCOM*. IEEE, 2004, pp. 929–940.
- [2] M. A. Tariq, B. Koldehofe, G. G. Koch, and K. Rothermel, "Providing probabilistic latency bounds for dynamic publish/subscribe systems," in *Kommunikation in Verteilten Systemen (KiVS)*. Springer, 2009, pp. 155–166.
- [3] O. Papaemmanouil, Y. Ahmad, U. Cetintemel, J. Jannotti, and Y. Yildirim, "Extensible optimization in overlay dissemination trees," in *SIGMOD*, 2006, pp. 611–622.
- [4] R. Baldoni, C. Marchetti, A. Virgillito, and R. Vitenberg, "Content-based publish-subscribe over structured overlay networks," in *ICDCS*. IEEE, 2005, pp. 437–446.
- [5] A. Carzaniga, D. S. Rosenblum, and A. L. Wolf, "Design and evaluation of a wide-area event notification service," *TOCS*, vol. 19, no. 3, pp. 332–383, 2001.
- [6] A. Yu, P. K. Agarwal, and J. Yang, "Subscriber assignment for wide-area content-based publish/subscribe," *TKDE*, vol. 24, no. 10, pp. 1833–1847, 2012.
- [7] G. Li, Y. Wang, T. Wang, and J. Feng, "Location-aware publish/subscribe," in *KDD*. ACM, 2013, pp. 802–810.
- [8] H. Hu, Y. Liu, G. Li, J. Feng, and K. Tan, "A location-aware publish/subscribe framework for parameterized spatio-textual subscriptions," in *ICDE*, 2015, pp. 711–722.
- [9] X. Wang, Y. Zhang, W. Zhang, X. Lin, and W. Wang, "Ap-tree: Efficiently support continuous spatial-keyword queries over stream," in *ICDE*, pp. 1107–1118.
- [10] L. Chen, G. Cong, and X. Cao, "An efficient query indexing mechanism for filtering geo-textual data," in *SIGMOD*. ACM, 2013, pp. 749–760.
- [11] X. Wang, Y. Zhang, W. Zhang, X. Lin, and Z. Huang, "Skype: top-k spatial-keyword publish/subscribe over sliding window," *Proceedings of the VLDB Endowment*, vol. 9, no. 7, pp. 588–599, 2016.
- [12] M. Yu, G. Li, and J. Feng, "A cost-based method for location-aware publish/subscribe services," in *Proceedings of the 24th ACM International Conference on Information and Knowledge Management*. ACM, 2015, pp. 693–702.
- [13] L. Chen, G. Cong, X. Cao, and K. Tan, "Temporal spatial-keyword top-k publish/subscribe," in *ICDE*, 2015, pp. 255–266.
- [14] W. W. Terpstra, S. Behnel, L. Fiege, A. Zeidler, and A. P. Buchmann, "A peer-to-peer approach to content-based publish/subscribe," in *DEBS*, 2003.
- [15] A. Gupta, O. D. Sahin, D. Agrawal, and A. El Abbadi, "Meghdoot: Content-based publish/subscribe over P2P networks," in *Middleware*, 2004, pp. 254–273.
- [16] R. Barazzutti, T. Heinze, A. Martin, E. Onica, P. Felber, C. Fetzer, Z. Jerzak, M. Pasin, and E. Riviere, "Elastic scaling of a high-throughput content-based publish/subscribe engine," in *Distributed Computing Systems (ICDCS)*, 2014 IEEE 34th International Conference on. IEEE, 2014, pp. 567–576.
- [17] A. Aji, F. Wang, H. Vo, R. Lee, Q. Liu, X. Zhang, and J. Saltz, "Hadoop gis: a high performance spatial data warehousing system over mapreduce," *Proceedings of the VLDB Endowment*, vol. 6, no. 11, pp. 1009–1020, 2013.
- [18] A. Eldawy and M. F. Mokbel, "Spatialhadoop: A mapreduce framework for spatial data," in *2015 IEEE 31st International Conference on Data Engineering*. IEEE, 2015, pp. 1352–1363.
- [19] A. Akdogan, U. Demiryurek, F. Banaei-Kashani, and C. Shahabi, "Voronoi-based geospatial query processing with mapreduce," in *Cloud Computing Technology and Science (CloudCom)*, 2010 IEEE Second International Conference on. IEEE, 2010, pp. 9–16.
- [20] S. Nishimura, S. Das, D. Agrawal, and A. E. Abbadi, "Md-hbase: a scalable multi-dimensional data infrastructure for location aware services," in *MDM*, vol. 1. IEEE, 2011, pp. 7–16.
- [21] A. M. Aly, A. R. Mahmood, M. S. Hassan, W. G. Aref, M. Ouzzani, H. Elmeleegy, and T. Qadah, "Aqwa: adaptive query workload aware partitioning of big spatial data," *Proceedings of the VLDB Endowment*, vol. 8, no. 13, pp. 2062–2073, 2015.
- [22] J. Chen, D. J. DeWitt, F. Tian, and Y. Wang, "Niagaracq: A scalable continuous query system for internet databases," in *ACM SIGMOD Record*, vol. 29, no. 2. ACM, 2000, pp. 379–390.
- [23] S. Chandrasekaran and M. J. Franklin, "Psoup: a system for streaming queries over streaming data," *The VLDB Journal*, vol. 12, no. 2, pp. 140–156, 2003.
- [24] M. Cherniack, H. Balakrishnan, M. Balazinska, D. Carney, U. Cetintemel, Y. Xing, and S. B. Zdonik, "Scalable distributed stream processing," in *CIDR*, vol. 3, 2003, pp. 257–268.
- [25] D. J. Abadi, Y. Ahmad, M. Balazinska, U. Cetintemel, M. Cherniack, J.-H. Hwang, W. Lindner, A. Maskey, A. Rasin, E. Ryvkina *et al.*, "The design of the borealis stream processing engine," in *CIDR*, 2005, pp. 277–289.
- [26] A. R. Mahmood, A. M. Aly, T. Qadah, E. K. Rezig, A. Daghistani, A. Madkour, A. S. Abdelhamid, M. S. Hassan, W. G. Aref, and S. Basalamah, "Tornado: A distributed spatio-textual stream processing system," *PVLDB*, vol. 8, no. 12, pp. 2020–2023, 2015.
- [27] Z. Chen, G. Cong, Z. Zhang, T. Z. Fu, and L. Chen, "Distributed publish/subscribe query processing on the spatio-textual data stream," *arXiv preprint arXiv:1612.02564*, 2016.
- [28] B. B. Cambazoglu, E. Kayaaslan, S. Jonassen, and C. Aykanat, "A term-based inverted index partitioning model for efficient distributed query processing," *TWEB*, vol. 7, no. 3, p. 15, 2013.
- [29] F. Basik, B. Gedik, H. Ferhatosmanoğlu, and M. E. Kalender, "S³-tm: scalable streaming short text matching," *The VLDB Journal*, vol. 24, no. 6, pp. 849–866, 2015.
- [30] S. Vaid, C. B. Jones, H. Joho, and M. Sanderson, "Spatio-textual indexing for geographical search on the web," in *SSTD*, 2005, pp. 218–235.
- [31] V. Gulisano, R. Jimenez-Peris, M. Patino-Martinez, and P. Valduriez, "Streamcloud: A large scale data streaming system," in *ICDCS*. IEEE, 2010, pp. 126–137.