Efficient Covering for Top-k Filtering in Content-Based Publish/Subscribe Systems

Kaiwen Zhang[1,2,5], Mohammad Sadoghi[3], Vinod Muthusamy[4], Hans-Arno Jacobsen

[1]Middleware Systems Research Group
[2]Technische Universität München
[3]University of California, Davis
[4]IBM T.J. Watson Research Center
[5]University of Toronto

Abstract

We investigate the use of content-based publish/subscribe for data dissemination in large-scale applications with expressive filtering requirements. In particular, we focus on top-k subscription filtering, where a publication is delivered only to the k best ranked subscribers, as ordered using expressive semantics such as relevance, fairness, and diversity. The naive approach to perform filtering early at the publisher edge works only if complete knowledge of the subscriptions is available, which is not compatible with the well-established covering optimization in scalable content-based publish/subscribe systems. We propose an efficient rank-cover technique to reconcile top-k subscription filtering with covering. We extend the covering model to support top-k and describe a novel algorithm for forwarding subscriptions to publishers while maintaining correctness. We also establish a framework for supporting different types of ranking semantics and propose an implementation to support fairness. Finally, we compare our solutions to a baseline covering system and perform sensitivity analysis to demonstrate that our optimized rank-cover algorithm retains both covering and fairness while achieving properties advantageous to our targeted workloads. In a typical setting, our optimized solution is scalable, selects fairly, and provides over 81% of the covering benefit.

CCS Concepts • **Software and its engineering → Publish-subscribe / event-based architectures**;

1 Introduction

The volume of data is increasing exponentially, with some estimating that over 2.5 billion gigabytes of information is generated daily [1]. Given its sheer size, new challenges arise for distributed content dissemination. Therefore, the ability for producers to deliver pertinent data only to selected and interested consumers is an important requirement which exists in many domains including targeted marketing, software maintenance, and job scheduling.

In targeted marketing, the merchants wish to limit the size of their advertising audience based on their social network profile and context such as location. For example, there could be a finite number of promotions to be delivered only to individuals who are most likely to be interested in the offer or redeem the coupon. Selecting the matching individuals may be based on a *relevance* ranking along with other objectives such as *fairness,* where every individual within the targeted demographic must have an equal chance of being informed about new deals.

The enterprise domain also presents data dissemination challenges: cloud providers are distributing software patches to client devices in a selective manner using ranking semantics such as relevance and fairness. The notion of *relevance* is unavoidable since not every software patch should be delivered to every device, such as a situation where a vulnerability only affects Android mobile phones by a particular manufacturer or a service provider. However, solely relying on relevance is insufficient. For instance, the risk of releasing a security patch can be limited by incrementally delivering the updates to a few devices. In this case, *fairness* semantics are necessary in order to select an unbiased sampling of devices.

Another class of problems that requires sophisticated communication patterns is job scheduling. Master nodes submit jobs that need to be processed by a number of workers. The challenge is to select the appropriate set of worker nodes for a given job, based on factors such as their compute capabilities, locality, and cost characteristics [29]. For example, while there may be many workers that satisfy the requirements for a job, some nodes may be faster, closer, or cheaper than others. That is, there is an implicit ranking of the *relevance* among the worker nodes for a given job.

Common to the above scenarios is a network of interconnected consumers (i.e., subscribers) and producers of data (i.e., publishers) with complex filtering requirements. Given the inherently distributed nature of these applications, there is a need to establish a decentralized architecture to cope with the data volume. The content-based pub/sub model, known for its scalability and decoupled nature, is a logical substrate candidate for distributed data dissemination [2, 5, 9, 10, 13, 16, 26, 30–32, 34]. Furthermore, it has been shown that every one of the above scenarios can be built on top of pub/sub primitives: social networks [28], data center maintenance [17], and job scheduling [15].

Although content-based pub/sub systems already offer fine-grained filtering capabilities, there is a need to explicitly limit publication deliveries (*top-k filtering*) in order to restrict network

overhead and information overload. In some applications, a subscriber may wish to control the amount of publications it receives (*consumer-centric* filtering), while in others, a publisher may want to limit the number of subscribers each publication is delivered to (*producer-centric* filtering). The scenarios outlined above call for the latter capability and is the problem addressed in this paper. We also remark that producer-centric top-k filtering is a better fit for pub/sub, since it is processed as a stateless operation which does not store publications, which is what pub/sub systems are innately suited for. On the other hand, consumer-centric top-k filtering is a stateful operation which requires the pub/sub system to be extended to deal with window semantics, which introduces publication buffering and consistency issues, as described in our prior work [34].

More specifically, this paper focuses on delivering each publication to its k highest ranked matching subscriptions over a pub/sub broker overlay. Subscription rankings can be based on a variety of criteria including relevance, fairness, or a combination of these criteria.

There is a requirement for pub/sub systems to employ optimized routing protocols and efficient filtering mechanisms for disseminating information flowing from data producers to data consumers. A well-known routing optimization prevalent in content-based pub/sub systems is covering [20]. This protocol reduces the propagation and management of subscriptions that are forwarded to brokers (i.e., nodes) in the pub/sub network. Covering optimization algorithms are compatible only with simple subscription matching and do not extend trivially to more expressive semantics such as top-k filtering using relevance and fairness.

In general, without the covering protocol, every edge broker to which a publisher is connected to maintains global knowledge of all subscriptions in the system, thereby making top-k subscription filtering a centralized task since the rank of all subscriptions can be computed locally. But once subscription covering is assumed (as is commonly the case in pub/sub systems), the edge broker has only partial knowledge of existing subscriptions. Top-k filtering then becomes a challenging and important problem in that setting.

In this paper, we make the following contributions:

1. Formalize general top-k subscription filtering semantics to express a wide range of ranking objectives including relevance and fairness.
2. Develop a novel rank-cover algorithm (RankCoverK) which leverages a partially ordered set (poset) of subscriptions to enable relevance-based top-k subscription filtering while supporting covering.
3. Introduce an ancestor counting optimization (AncestorCoverK) to further reduce the size of the covering poset produced by RankCoverK and the resulting subscription traffic.
4. Extend the previous algorithm to consider fairness (AncestorFairK) and develop a weighted sort shuffle algorithm for efficient subscription selection.
5. Conduct an extensive sensitivity analysis of our algorithms (on an open-source pub/sub system) with respect to edge broker load reduction (covering effectiveness), end-to-end latency, processing overhead, and fairness semantics. The results demonstrate the scalability of our optimized ancestor counting approach (AncestorCoverK), and the fairness of its extended version (AncestorFairK).

2 Related work

Generally speaking, problems related to retrieving the most relevant answers have been studied in different contexts including database (distributed) top-k querying [3, 6, 14, 18, 22] and publish/subscribe (pub/sub) matching techniques [2, 5, 10, 13, 16, 26, 31, 32, 34].

The most widely adopted database top-k processing model [14, 18] differs from our proposed top-k model in an important respect: our top-k model solves the reverse problem. In the database context, top-k querying means finding the most relevant tuples (events) for a given query (subscription). But in our pub/sub abstraction, matching means finding the relevant subscriptions (queries) for a given event (tuple).

Broadly speaking, two classes of matching algorithms have been proposed for pub/sub: counting-based [13, 31, 32] and tree-based [2, 5, 26, 27] approaches. A fresh look at enhancing pub/sub matching algorithms is to leverage top-k processing techniques to improve matching performance. An early top-k model is presented in [21]; however, this model leverages a fixed and predetermined scoring function, i.e., the score for each expression is computed independent of the incoming event. In addition, this approach is an extension of *R*-Tree, the interval tree, or the segment tree structure; as a result, it is ideal for data with few dimensions [21]. In contrast, a scalable top-k model that supports up to thousands of dimensions while incorporating a generic scoring function, i.e., takes the event into consideration, is introduced in [31], which relies on a static and single-layered pruning structure. To alleviate these challenges, a new dynamic and multi-layered pruning top-k structure is developed in [27]. However, our proposed top-k model attempts to solve a different problem, namely, distributed top-k processing; therefore, our model can leverage any of the existing top-k work as building blocks (e.g., [21, 27, 31]).

Another important aspect of pub/sub top-k matching is to explore and identify a set of plausible top-k semantics. Unlike in the database context, formalizing top-k semantics in pub/sub is more involved and not limited to a single interpretation [11, 12, 25]. The most widely used pub/sub top-k semantics is defined with respect to subscribers, i.e., consumer-centric semantics, in which the subscription language is extended (with a scoring function) in order to rank each incoming publication (over a time- or count-based sliding window); thus, delivering only the top-k matched publications to each subscriber [11, 12, 25].

Alternatively, the top-k semantics can be defined with respect to a publisher, i.e., producer-centric semantics, which extends the publication language for ranking subscribers and delivering publications only to the top-k matched subscribers [21, 31]. Producer-centric semantics is suitable for targeted advertisement (e.g., targeting a specific demographic group) and diversified advertisement (e.g., reaching out to most eligible members within each demographic group). Finally, hybrid semantics can be foreseen such that both subscribers and publishers have control on how data is received and disseminated, respectively.

Recent works argue for the importance of top-k matching based on the relevance of subscriptions (i.e., producer-centric) using a non-monotonic and dynamic scoring [10], but without paying attention to the covering routing techniques employed by pub/sub systems and other top-k ranking semantics such as fairness, which is the focus of our work. A window-based top-k matching solution that

also considered covering was studied in [34], but unlike our present work, their focus was on consumer-centric filtering.

Other notable top-k research directions in pub/sub are as follows. To reduce memory footprint, a probabilistic model with a bounded-error was developed in order to determine whether a given publication will eventually fall in top-k results of at least one subscriber [25]. The problems of how to formulate the subscriber's interest and ranking publications were studied in [11]. To further improve the relevance of top-k publication results (e.g., improving the diversity), different top-k measures such as quantitative approaches (based on scoring functions) and qualitative approaches (based on explicit binary relation preferences) were explored in [12]. More recently, the problem of top-k diversification was also studied in the context of finding the most diverse set of text documents for each subscriber over data streams [8]. Furthermore, the problem of computing top-k results under spatial constraints were studied in [9, 30]. Unlike these approaches, our work focuses on efficient subscription covering which supports distributed top-k computation in order to reduce overall network traffic.

In a previous short paper presentation, we introduced the problem of top-k filtering in content-based publish/subscribe [33]. We included this content for the sake of completeness and readability. However, the current full-length paper also contains theorems, associated proofs, support for fairness, which includes a novel weighted shuffle algorithm, and additional evaluation experiments.

Our paper operates under the well-known model of a distributed broker overlay network operating using content-based routing [7]. In particular, this model assumes that broker links are statically defined, and rely on subscription covering to reduce traffic. Alternatives, such as overlays with direct links [19] or elastic operator-based systems [4], are outside the scope of this paper.

3 Semantics and naive solution

In this section, we first revise the conventional definition of subscription covering used in the context of advertisement-based routing with content-based matching. We then formalize a model for top-k subscription filtering in pub/sub. Our model consists of a general framework that supports a variety of ranking semantics used for selecting a top-k of subscriptions. In particular, we describe our relevance scoring and fairness ranking criteria for top-k filtering. Second, we describe how top-k subscription filtering semantics are propagated through the system by piggybacking top-k parameters via advertisements. Finally, we demonstrate how covering, in conjunction with the proposed top-k model, is incompatible with a naive solution and is thus a non-trivial challenge to solve.

3.1 Pub/Sub model with subscription covering

In a content-based matching model, subscriptions consist of a conjunction of predicates of the form $(attribute, operator, value)$. These predicates are then matched against the attribute-value pairs $(attribute, value)$ of incoming publications. A publication p is considered *matching* for a subscription s if all its predicates are satisfied by p. In that case, p must be delivered to s. We employ advertisement-based routing, where publishers must first send advertisements which are flooded through the broker overlay. Subscriptions are routed through the network by matching them against advertisements and forwarding them to *upstream* brokers (towards the corresponding publishers). Subsequently, publishers are allowed to publish publications only within the space they advertised.

i	x	y	area	rank
1	[1,10]	[1,10]	100	1
2	[1,10]	[1,5]	50	2
3	[3,7]	[1,8]	40	3
4	[9,10]	[9,10]	4	7
5	[1,5]	[1,5]	25	4
6	[3,10]	[2,4]	24	5
7	[3,6]	[1,5]	20	6

Table 1. Example subscriptions and scoring at B_2

Subscription covering is a routing optimization which reduces the volume of subscriptions to be propagated through the pub/sub system. A set of incoming subscriptions $\mathbb{S}$ is then processed as shown in Figure 1. First, the subs are matched against all existing advertisements $\mathbb{A}$ to determine which publishers can publish potentially matching publications.

Second, we determine the covering relationships between existing subscriptions and the new ones in $\mathbb{S}$. A subscription s_1 covers subscription s_2 if all publications that match s_2 also match s_1. This can be determined by inspecting the predicates of the two subscriptions. For instance, suppose the predicates of s_1 are $x \in [1, 10], y \in [1, 10]$, and s_2 are $x \in [1, 5], y \in [1, 10]$. In this situation, s_1 covers s_2.

After applying covering, we obtain a subset C of $\mathbb{S}$ with subscriptions which are *not* covered. C is then forwarded towards the sources of advertisements found in A. Note that any subscription in C which has previously been sent is not forwarded again.

Consider the example in Figure 2, where a publisher is connected to broker B_1 and the subscribers connected to brokers B_2 and B_3. Suppose that the subscribers at broker B_2 issue the subscriptions s_1 to s_7 of the form $x \in [a, b], y \in [c, d]$, which are listed in Table 1. We can build a covering partially ordered set (poset) for B_2, as depicted in Figure 3. Higher level subscriptions cover their descendants. In this example, s_1 covers subscriptions s_2 to s_7. Therefore, applying the covering optimization allows broker B_2 to propagate only the subset $C_2 = \{s_1\}$ to upstream broker B_1.

In the experiments in Section 5, we call RegularCover the above technique without top-k filtering employed.

3.2 Top-k model and ranking criteria

Consider a publication p that matches a set of subscriptions S according to content-based filtering semantics [7]. The problem addressed in this paper is to deliver the publication to the highest ranked k-sized subset of S. Different ranking criteria can be used to determine the subset. For example:

- *Relevance-based (scoring) semantics.* There is a scoring function, score(p, s), that computes a score given a publication p and subscription s. The subscriptions are ranked by descending score, such that the k subscriptions with the highest scores are selected. Table 1 shows the ranking of the example subscriptions using the x-y area as the score.
- *Fairness semantics.* Each rank between $[1, |S|]$ is randomly assigned with equal probability to every subscription in S. Therefore, a fair delivery selects a uniformly random k-sized subset of S. More precisely, the probability of delivering p to any subscription $s_i \in S$ is $\frac{k}{|S|}$.
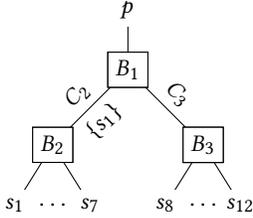
Figure 1. Subscription processing flow with covering
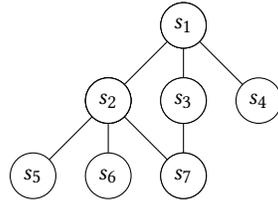

Figure 2. Example with three brokers

Figure 3. Covering poset at B_2

Note that scoring and fairness could be combined together to determine a ranking. For instance, subscriptions could first be sorted by their score. Then, if multiple subscriptions have the same score, we use fairness to break ties such that each subscription holds a unique rank.

Figure 4 shows an overview of the processing flow at a broker using top-k filtering. Each incoming publication p is first matched against known subscriptions $\mathbb{S}$ under the usual pub/sub model (e.g., content-based predicates). The resulting set S of matching subscriptions is sorted according to the provided semantics described above (e.g., relevance, fairness, or some combination of both). This ranked set S of matching subscriptions is then filtered to retain only a k-sized subset of the highest ranked subscriptions. These subscriptions, along with the publication p, are provided to the usual routing component in order to forward p to the appropriate next hops towards the intended destinations (e.g., *downstream* brokers).

3.3 Semantics propagation

The top-k ranking semantics above are from the perspective of the publisher. In particular, the case where a subscriber chooses to receive only a subset of matching publications has been studied and not considered in this paper [34].

As a consequence of this decision, the publisher must specify a value k with each advertisement a it issues, indicating that publications induced by a should be sent to the k highest-ranked matching subscriptions. In addition, a ranking function $rank(p, S)$ is issued either with the advertisement, or there could also be a system-wide ranking function. In this paper, we assume the most general case of a per-advertisement ranking function. Note that each individual publication can request a specific k value, as long as it is lower than the k defined by a.

For a publication p and set of matching subscriptions S, the ranking function $rank(p, S)$ assigns a rank in the range $[1, |S|]$ to each subscription s in S.

We note that our distributed approach is compatible with any ranking function, such as those described in [12] (e.g., see Appendix A on diversity). In this paper, we provide reference implementations for relevance and fairness ranking criteria.

3.4 Naive solution

This paper focuses on the forwarding of advertisement, subscription, and publication messages to achieve top-k dissemination. Notably, the algorithm to compute the top-k set for a publication is assumed to be known. More specifically, given a set of subscriptions $\mathbb{S}$ and a publication p: (i) a matching algorithm can compute the set of subscriptions $S_p \subseteq \mathbb{S}$ that match p, and (ii) a top-k algorithm can use a provided $rank(p, S_p)$ function to compute the k-sized set of top ranked subscriptions in S_p.

If there is no subscription covering, then each broker can compute the top-k ranked subscriptions locally and forward the publications accordingly along with the IDs of the subscriptions to deliver the publication to. We refer to this algorithm, when used without covering, as RegularK in Section 5.

3.5 Problem with naive solution

The naive solution works because each broker knows the complete set of subscriptions in the system, allowing for the ranks of matching subscriptions to be computed. However, routing with covering forwards only a subset of subscriptions to upstream brokers. The subset is chosen to be an appropriate summary of the subscriptions such that the upstream broker can correctly determine which downstream brokers may have matching subscriptions. In other words, the subscription covering algorithm avoids forwarding subscriptions whose interest space is subsumed by another subscription [20].

Given a set of subscriptions, and their pairwise covering relationships, the broker constructs a poset that represents the covering relationships among its subscriptions. It then forwards only those subscriptions that are not covered by any other.

Consider once again the three brokers example in Figure 2, with the covering poset in Figure 3 for broker B_2. Suppose $k = 2$ is employed. When subscription covering is not employed, broker B_1 is aware of all the subscriptions and can select the appropriate top-k set as outlined in Section 3.4. However, with subscription covering, broker B_2 only forwards the roots of the poset, subscription s_1 in this case, to the upstream broker B_1. Since broker B_1 does not have knowledge of all the subscriptions, it cannot accurately rank all the subscriptions (using their area) and therefore determine the k highest ranked subscriptions from B_2 and B_3. In Section 4, we resolve the problem by modifying the covering component of subscription forwarding (see the dotted box of Figure 1) which then allows publication processing to achieve correct top-k filtering while retaining most of the benefits of subscription covering.

4 Top-k forwarding algorithms

This section presents first a distributed top-k algorithm (RankCoverK) that correctly supports covering for relevance-based ranking, as opposed to the RegularK solution shown in Section 3.4. We also present AncestorCoverK, an optimized version of RankCoverK using ancestor counting. Finally, we propose AncestorFairK, an extension of AncestorCoverK which employs a fair ranker in addition to the

Figure 4. Publication processing flow with top-k filtering

relevance scorer. We provide theorems to support our solutions, with the proofs in Appendix B.

4.1 Top-k-aware subscription covering

We now describe a set of algorithms that support subscription covering. The challenge is to decide how many of the k highest ranked subscriptions are reachable from each downstream broker.

Without loss of generality, we consider again the topology in Figure 2. Brokers B_2 and B_3 should employ covering and forward a subset of their subscriptions to B_1 sufficient for supporting top-k filtering. Suppose $k = 2$ for the publisher. Then, when a publication arrives at broker B_1, it needs to distinguish among three cases:

1. **2-0** case: The two highest ranked subscriptions are both reachable through broker B_2.
2. **1-1** case: One of the highest ranked subscriptions is reachable through broker B_2, and the other through B_3.
3. **0-2** case: The two highest ranked subscriptions are both reachable through broker B_3.

The key insight is that broker B_1 only needs to distinguish among the above three cases. In particular, it does not need to know if there are more than $k = 2$ matching subscriptions downstream of a particular broker. Therefore, it suffices for each broker to forward its own k highest ranking subscriptions for all possible publications that may be induced by advertisement a. Note that unlike the RegularK solution, each broker along the dissemination path must perform top-k filtering to ensure that covered subscriptions can be selected.

We will now generalize and formalize this point. Let S_j be the set of subscriptions that match publication p_j at a given broker, and let $s_{j,1}, ..., s_{j,n}$ be the subscriptions in S_j ordered by ranking. Furthermore, let $\hat{S}_j$ be the top-k set in S_j, that is, $\hat{S}_j = s_{j,1}, ..., s_{j,q}$ where $q = min(n, k)$. Finally, let $\hat{S} = \bigcup_j \hat{S}_j$ be the set of highest ranking subscriptions for all possible publications p_j.

Theorem B.1. *If for each downstream broker $C = \hat{S}$, where C is the set of forwarded subscriptions to upstream broker B, then B can compute the top-k subscriptions for any given publication.*

We now outline several algorithms to compute $\hat{S}$, the set of subscriptions each broker must forward to its upstream broker in order to allow top-k filtering with subscription covering.

4.1.1 Rank-cover

Let subscriptions $s_1, \cdots, s_n$ be the subscriptions at a given broker that intersect an advertisement a from an upstream broker.

In this algorithm, the subscription cover is a poset built according to an *extended covering definition* that is based on the subscription predicates and given ranking functions. We call this the **rank-cover** definition. In particular, subscription s_i rank-covers s_j if and only if both these conditions are true:

(i) the predicates of s_i cover those of s_j. (This is the conventional covering definition.)

Algorithm 1: Rank cover subscription forwarding

1 **foreach** *advertisement a* **do**
  /* Construct rank-cover poset for subs that
     intersect a and are from a different last
     hop.   */
2 | $S_a \leftarrow \{s | s$ matches $a \wedge s.lasthop \neq a.lasthop\}$
3 | $C_a \leftarrow$ rank-covering poset among subscriptions S_a
  /* Construct the forwarding set for a for a
     certain value of k.   */
4 | $F \leftarrow$ subscriptions at depth $\leq k$ in C_a
5 | add F in the subs to forward to $a.lasthop$

(ii) the rank of s_i covers that of s_j for advertisement a. More precisely, $\forall p_b, \{s_i, s_j\} \subseteq S_b \implies (ranking(s_i, S_b) < ranking(s_j, S_b))$, where p_b is any publication induced by advertisement a, and $ranking(s, S_b)$ returns the rank of a subscription s in the matching subscription set S_b, for publication p_b.

Intuitively, rank-covering implies that if a publication matches s_j, then it will also match s_i with a higher rank.

Each broker then forwards all the subscriptions in the first k levels of the poset, where k is the desired number of top-k subscriptions. This will ensure that the upstream broker has enough subscriptions to determine how many of the top-k ranking subscriptions are downstream.

More precisely, this algorithm constructs $\hat{S}$ as $\{s_i | depth(s_i) \leq k\}$ where $depth(s_i)$ is the depth of subscription s_i in the rank-cover poset.

Theorem B.2. *The set $\hat{S} = \{s_i | depth(s_i) \leq k\}$ contains the set of top-k subscriptions for any possible publication by advertisement a.*

For example, consider the topology in Figure 2, and let subscriptions $s_1, ..., s_7$ be the subscriptions at broker B_2 that intersect with advertisement a (with $k = 2$) from broker B_1. Figure 3 depicts the rank-covering poset for the subscriptions at B_2 that intersect advertisement a using area for scoring. Note that in this example, it is identical to the regular covering poset. Since $k = 2$, all nodes with depth ≤ 2 will be forwarded to B_1. So, subs s_1, s_2, s_3, s_4 will be sent to B_1. Broker B_1 now has enough subscriptions from B_2 to determine if 0, 1 or 2 of the top ranked subscriptions are from broker B_2.

Note that the algorithm must be rerun when there is a change in the set of subscriptions or advertisements, as is true of the traditional covering algorithm.

The rank-covering algorithm is defined in Algorithm 1, and is referred to as RankCoverK for the remainder of this paper.

4.1.2 Ancestor counting

The previous rank-cover covering algorithm can be optimized so that certain subscriptions, which have multiple parents, are removed from the forwarding set. This happens if there are enough overlapping ancestors to cover the subscription and satisfy k matches.

Algorithm 2: Ancestor counting forwarding

---

1 **foreach** *advertisement a* **do**

  /* Construct rank-cover poset for subs that intersect *a* and are from a different last hop.   */

2   $S_a \leftarrow \{s | s$ matches $a \wedge s.lasthop \neq a.lasthop\}$

3   $C_a \leftarrow$ rank-covering poset among subscriptions S_a

  /* Construct the forwarding set for *a* for a certain value of *k*.   */

4   $F \leftarrow$ subscriptions at depth $\leq k$ in C_a

  /* Remove subscriptions who have more than *k* ancestors.   */

5   **foreach** *sub* $s \in F$ **do**

6   **if** *s.ancestors* $\geq k$ **then**

7   remove *s* from *F*

8   add *F* in the subs to forward to *a.lasthop*

In this optimization, rather than forwarding all subscriptions in the poset with depth $\leq k$, we only forward those subscriptions with $< k$ ancestors in the poset. This is outlined in Algorithm 2. Precisely, this algorithm constructs $\hat{S}$ as $\{s_i | \text{ancestors}(s_i) < k\}$ where ancestors(s_i) is the number of subscriptions in the rank-cover poset that are in the path from the root to s_i.

Theorem B.3. *The set $\hat{S} = \{s_i | \text{ancestors}(s_i) < k\}$ contains the set of top-k subscriptions for any possible publication from advertisement a.*

If $k = 3$ in the example from Figure 3, every subscription would be forwarded except s_7 which has 3 ancestors. This algorithm is referred to as AncestorCoverK.

4.2 Fairness

This section extends the solution to also take into account fairness when selecting the k subscriptions to forward a publication. This solution, called AncestorFairK, uses the same covering component as AncestorCoverK, but modifies the ranking component for processing publications (see dotted box in Figure 4). First, the fair ranker uses relevance to assign a score to each subscription (when processing an incoming publication). Then, if two subscriptions have the score, our fairness breaks ties by assigning ranks to each tied subscription with equal probability. For instance, suppose a publication p must be delivered to the top 10 highest ranked publications. Based on score, there is a clear order for the top 8, but 3 subscriptions are tied then. The ranks of $9 - 11$ must then be fairly assigned to these subscriptions. Then, the probability that each of these 3 subscriptions will be part of the top-k set is 2/3. In other words, these 3 tied subscriptions are each equally likely to be in the top-k set.

Note that if there is no subscription covering, then the solution is trivial as discussed in Section 3.4: the first broker that receives the publication from the publisher can compute the set of matching subscriptions, uniformly select among them, and forward the ids of these selected subscriptions to each downstream broker.

One attempt to preserve fairness is to break ties fairly only considering subscriptions in the subsets $\hat{S}$ from downstream brokers generated by our above two algorithms in Section 4.1. However, this simple approach could produce incorrect results. For example, consider a broker with matching subscriptions S and corresponding

---

Algorithm 3: Weighted shuffle algorithm

  /* After receiving incoming pub *p*, construct a cumulative weight search tree *wtree*:   */

1 $sum \leftarrow 0$

2 $tail \leftarrow 0$

3 **foreach** *subscription s* **do**

4   $sum \leftarrow sum + s.weight$

5   insert (sum, s) in *wtree*

6   $tail \leftarrow sum$

  /* Shuffle the elements with weighted probability: */

7 $slist \leftarrow \{\}$

8 **repeat**

9   $x \leftarrow random(1, tail)$

10 $e \leftarrow ceiling(x, wtree)$

11 $slist \leftarrow slist \cup \{e\}$

12 replace e with a pointer to the range $[tail - e.weight, tail]$ in *wtree*

13 $tail \leftarrow tail - e.weight$

14 **until** $|S|$ *times*

15 **return** *slist*

---

forwarding set $\hat{S}$ as computed by the aforementioned techniques. Now add a new subscription s' to get $S' = \{s' \cup S\}$. Suppose s' has the same score as some subscription $t \in \hat{S}'$, but $s' \notin \hat{S}'$ (not forwarded upstream). Consequently, $\hat{S} = \hat{S}'$, and hence from the upstream broker's perspective, the case where s' is present is indistinguishable from the case where it is absent. If the upstream broker has to break ties for subscriptions which have the same score as s', it cannot fairly assign a probability to s' since it is not in $\hat{S}'$ and therefore cannot be taken into consideration by the simple approach.

Given the above point, we present an approximate fairness algorithm which seeks to capture the score of non-forwarded subscriptions to be taken into consideration for scoring tiebreaks. In this algorithm, subscriptions are forwarded as usual, but each forwarded subscription is assigned a weight equal the number of descendants in its rank-cover poset. While forwarding a publication, each broker performs a weighted shuffle using the matching subscriptions, and selects the first k subscriptions in the shuffled sequence. The weighted shuffle algorithm, detailed in Algorithm 3 requires a search tree (e.g., red-black tree), labeled as *wtree*. This tree keeps track of the cumulative sum of the weight of inserted elements for each entry. A randomly generated value in the range up to the total weight of all the entries can be used to retrieve an element whose key is the least greater than the value selected. This algorithm requires n insertions and $O(n)$ searches for a total running time of $O(n \log n)$.

Note that AncestorFairK can support fairness as the sole ranking objective without considering relevance by simply using a scorer which assigns the same score to all matching subscriptions. The fair ranker then decides their rank entirely. More precisely, given a publication p that matches n subscriptions $s_1, ..., s_n$, the probability of delivering p to s_i is k/n. That is, each matching subscription is equally likely to be in the top-k set.

5 Evaluation

This section contains the results of experiments performed on our various algorithm implementations. We first describe scalability experiments which compare the covering performance of three implementations: (1) our baseline top-k solution with covering (RankCoverK), (2) an optimized version using ancestor counting (AncestorCoverK), and (3) our baseline covering algorithm without top-k (RegularCover). These experiments gauge the impact of supporting top-k over covering (i.e., how much covering is lost using RankCoverK and AncestorCoverK compared to RegularCover) and provide grounds for a discussion on the sensitivity of our work with respect to various workload parameters.

We then perform an evaluation of our fair ranker using our optimized top-k solution (AncestorFairK), compared to AncestorCoverK using a fairness metric. The fairness baseline for this experiment is our top-k solution without covering (RegularK). Each subscription has a fair chance to be selected in RegularK since each broker has complete knowledge of all downstream matching subscriptions when covering is not employed (see Section 3.4). Note that AncestorFairK performs equally to AncestorCoverK in terms of covering, and is thus omitted from the scalability experiments mentioned in the previous paragraph.

Finally, we measure the processing overhead of our various solutions using end-to-end publication latency to assess the impact of our top-k filtering and fairness selection algorithms on the publication match-and-forward process. Note that the cost of evaluating the score of each publication and finding the top-k is not considered in our evaluation, since our work does not focus on efficiently performing this local computation, but is rather concerned on the distribution of this computation over a network of brokers.

5.1 Setup

The implementation is performed in Java using the PADRES pub/sub prototype[1]. Experiments are performed on the SciNet testbed using the General Purpose Cluster (GPC)[2]. Up to 60 distinct machines in a LAN are used, each equipped with Intel Xeon quad-core processors and 16GB RAM.

The publication dataset used is extracted from Yahoo! finance stock information, with one publisher emitting 30 publications per minute. A single publisher allows us to better control the distribution of subscribers relative to the publisher which is important for fairness (see Section 5.4). In addition, our measurements of the covering approaches are unaffected by the number of publishers.

The broker overlay topology consists of 10 core brokers connected in a chain. Each core broker is then connected to 5 edge brokers. For our scalability experiments, we connect 200 to 1000 subscribers to each edge broker. We employ between 10000 to 50000 subscribers, each with one subscription. This setup is modeled as a network of data centers connected through gateways (core brokers) and measures the impact of our algorithms on delivery paths with multiple broker hops. For our other experiments, we employ 5 core brokers, each with 5 edge brokers, where each edge broker is connected to 10-20 subscribers.

Subscriptions are randomly and uniformly generated to span between 30% and 60% of the publication space (subscription size). We argue that the large size of the subscriptions is appropriate for

our motivating scenarios where users are receiving an overwhelming amount of data, hence, the need for additional top-k filtering. Furthermore, the narrow size range (30%-60%) creates conservative covering results with shallow and broad covering posets, since neither excessively large subscriptions exist to cover large spaces, nor are small subscriptions present that can easily be covered. We also conduct a sensitivity analysis where we vary the subscribe size range by varying the minimum (10%-60% to 50%-60%) or the maximum (30%-40% to 30%-60%).

The scalability experiments are conducted with two sets of k. First, advertisements are evaluated using low values of k between 1 and 10. These low values are used for micro-benchmarking purposes. Second, advertisements are set with higher values of k ranging between 0.1% to 10% of the total number of subscriptions, which we refer to as the *selectivity*. For instance, in an experiment with 50000 subscriptions, a k set to 5000 has a selectivity of 10%. As top-k is useful for suppressing an overwhelming amount of data, increasing the k value beyond 10% would be semantically ineffective. The other experiments are conducted using k ranging between 0.2% to 2% of the total number of subscriptions. This range of k tested is in line with evaluations from related papers [10, 12].

For all solutions except AncestorFairK, we employ a covering-agnostic random ranking function: each subscription known by the broker for a publication is assigned a uniformly random rank. Since our experiments focus on fairness, this setup allows us to maximize the size of the pool of subscriptions which must be fairly selected from. In practice, our solutions are compatible with any ranking scheme.

The fair ranker used by AncestorFairK extends the random ranker by weighing subscriptions appropriately according to the covered subscriptions (see Section 4.2). Furthermore, we asynchronously update the weights for the ranker by propagating the count of covered subscriptions through propagated subscriptions.

5.2 Metrics

Percentage of forwarded subscriptions: Our main metric for measuring the effectiveness of our covering techniques is the reduction in the number of subscriptions sent upstream (i.e., subscription traffic towards publishers). In our evaluation, we measure the ratio between the number of subscriptions at the publisher edge broker and the total number of subscriptions sent by all subscribers. Since the publisher edge broker is the last hop for every subscription in the system (i.e., every subscription is forwarded to the publisher edge broker if covering is not employed), this number is an overall measure for covering in the system. The lower bound is provided by our baseline covering algorithm RegularCover which does not support top-k, while the upper bound is 100% for top-k without covering (RegularK). Our proposed solutions fall somewhere in between and allow us to assess the impact of top-k dissemination on covering. Covering performance is then defined as the percentage of subscriptions covered and thus not forwarded ($100\% - forwarded\%$).

End-to-end latency: We measure the end-to-end latency from the time, in *ms*, elapsed between sending a publication at the publisher to the receipt of the same publication at a subscriber. This allows us to measure the aggregated processing overhead of our top-k algorithm being run at every broker hop encountered. In particular, this allows us to assess the impact of the weighted fair shuffle algorithm developed in this paper.

Fairness: We measure fairness empirically by comparing the number of publications each subscription receive for each implementation to the baseline top-k algorithm RegularK. We sum up the difference in number of publications received for all the subscriptions during the course of the experiment. This number indicates how much deviation the implementation has to the expected number of publications to be received from each subscription. Note that the subscriptions and publications are generated with the same seed for every experiment.

5.3 Covering performance

Figures 5(a) and 5(b) show the covering performance of our various algorithms relative to the baseline RegularCover with low k values (1-10) and different amount of subscriptions (with size between 30%-60%). A lower number indicates that the covering algorithm has been more effective in reducing the number of subscriptions disseminated.

We first note that the baseline RegularCover propagates only 1.47% of 10000 subscriptions, and 0.69% for 50000 subscribers. Note that the baseline performance is unaffected by the value of k, since it does not support top-k. RegularCover demonstrates scalability through the improvement in subscription reduction with increasing subscription loads.

The graphs show that our rank-cover solution, RankCoverK, is sensitive to increasing values of k: the number of subscriptions forwarded increases from 1.47% to 84.71% when k increases from 1 to 10 with 10000 subscriptions. The performance declines rapidly and is nullified (100% subscriptions forwarded) when k reaches 19. In the 50000 subscriptions case, the number of subscriptions forwarded by RankCoverK increases from 0.69% to 60.53% in the 1-10 range, which is an 28.5% improvement over the 10000 subs case. This indicates that the solution scales with the number of subscriptions in the low k range. However, when k is set to 27, the covering performance is still reduced to 0% of subscriptions covered. This indicates that RankCoverK can only operate at very low k values; there is a fixed limit for k which does not scale with the number of subscribers. This limitation limits the applicability of RankCoverK, as queries typically select a k as a function of the number of subscriptions (i.e., % selectivity). As described in Section 5.1, we selected subscriptions such that the covering posets are shallow and broad. Therefore, traversing a few levels deep is enough to encounter a large number of subscriptions, which explains the significant impact of k on the covering performance.

Our optimized solution, AncestorCoverK, eliminates this limitation by demonstrating covering performance across a wide range of k. In the low range of 1-10, the covering performance stays stable at under 6.03% for 10000 subscriptions and under 2.5% for 50000 subscriptions. AncestorCoverK exhibits the same performance benefits as RankCoverK with respect to the number of subscriptions, while removing the dependency on low values of k. This optimization is well-suited for our workload, since it contains many subscriptions residing at a shallow depth, but covered by multiple parents.

In Figure 5(c), we further test the scalability of AncestorCoverK with regards to k with values between 0.1% to 10%. For example, the experiment with 50000 subscriptions used a k between 50 to 5000. The results show that the covering performance decreases from over 94.24% to 33.3% when in that k range. This indicates that our optimized solution is still sensitive to higher values of k, albeit to a far lesser degree than RankCoverK. The performance is also stable

with regards to increased number of subscriptions, with differences of less than 0.4% across different runs.

In Figure 5(d), we vary the subscription size from a larger range (10%-60%) to a smaller range (30%-40% or 50%-60%) using Ancestor-CoverK at 10000 subscriptions. We first note that the subscription size also affects the baseline RegularCover: the performance varies between 1.22% for the largest range used (10%-60%) to 3.08% for the narrowest range used (30%-40%). When the subscriptions are similarly sized, the covering optimization decreases since it requires large subscriptions to be able to cover smaller ones. For Ancestor-CoverK, the same pattern can be observed. For instance, 40.33% of subscriptions are forwarded for the range 10%-60% at 5% selectivity, compared to 82.05% for 30%-40% at 5% selectivity. Also note that while 30%-40% and 50%-60% both exhibit a range of 10%, 50%-60% performs better (69.59% at 5% selectivity). This is because having larger subscriptions in general increases the chance of covering occurring, since subscriptions are more likely to overlap. We note there is 2.52 times increase in number of subscriptions forwarded for RegularCover from 10%-60% to 30%-40%, whereas there is only a 2.03 times increase for AncestorCoverK. We can therefore conclude that AncestorCoverK is not more sensitive to subscription size than the baseline RegularCover which does not support top-k.

Summary: Both solutions are sensitive to varying degrees to the value of the parameter k. In particular, RankCoverK is bounded by a k of 27, after which no covering is obtained regardless of the number of subscriptions. This is due to the broad and shallow nature of the covering tree, where most subscriptions can be found by traversing at a low k depth. On the other hand, AncestorCoverK retains 81% of the covering benefits of RegularCover at a k set to 1% selectivity, which is 5000 for 50000 subscribers. This range of k is practical for a variety of studied workloads, which typically have 1-2% selectivity [24]. Furthermore, both solutions are scalable to the number of subscriptions, as the covering performance stays stable in that regards. Finally, AncestorCoverK is at least as sensitive than the baseline RegularCover when varying the range of subscription sizes.

5.4 Fairness evaluation

Figure 6(a) evaluates the fairness of our optimized solution AncestorCoverK to AncestorFairK, which is the same optimized solution equipped with our fair ranker. A lower value of fairness is better, as it indicates that the solution sends a number of publications to each subscription closer to the expected value established by RegularK. We note that our fair ranker outperforms the random ranker at all values of k tested by an average of 224%. AncestorFairK also becomes increasingly fair as the value of k increases. This is due to the increase in notification volume which comes with larger values of k. With more subscriptions being selected, the ability of the fair ranker to do a weighted shuffle of those subscriptions is better leveraged. Furthermore, increasing the k value triggers more subscriptions to be sent upstream. Since the metadata used by the fair ranker is propagated via subscriptions, this increased frequency in subscription propagation increases the accuracy of the metadata. On the other hand, AncestorCoverK decreases in fairness as more brokers are encountered as the random ranker used further impacts the fairness of the solution. We also note that the fairness of AncestorCoverK improves by 20.18% between $k = 0.4\%(1)$ to $k = 0.8\%(2)$. As k increases, the number of subscriptions which are uncovered also rises. This larger number of subscriptions found
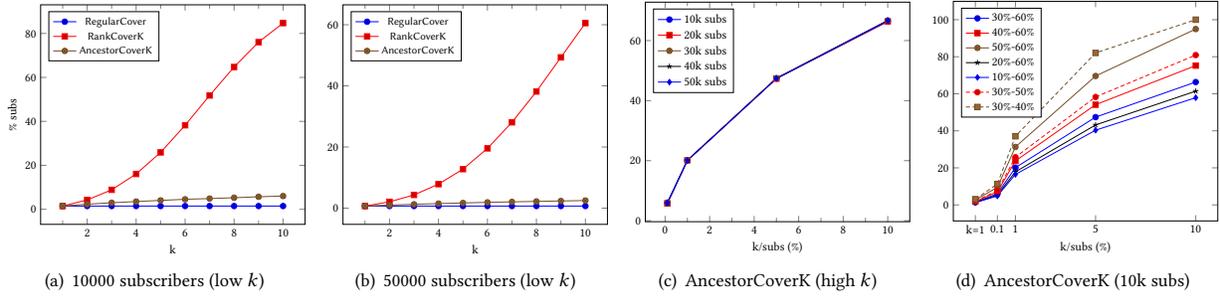
Figure 5. Percentage of forwarded subscriptions (varying k and subscription size)

at the publisher edge broker benefits the random ranker which now has a more accurate representation of the subscriptions in the system, but this benefit is not enough to offset the inaccuracy of larger selections with higher values of k.

In light of these observations, we conclude that AncestorFairK outperforms AncestorCoverK significantly and does not suffer from issues that arise when the k value is at the extreme ends of the intended range for k.

Figure 6(b) shows the sensitivity of fairness for subscriptions at varying distance from the publisher. Distance is measured as the number of core broker hops between the publisher and the subscriber. The figure shows the average fairness of subscriptions across the different values of k, since the same pattern is observed regardless of the k used. We observe that for both solutions, fairness improves when it reaches 3 broker hops, and then decreases when the number of hops increases beyond. This can be explained due to the nature of the inaccuracies. The subscriptions closer are more numerous at the publisher edge broker: because they are propagated through fewer brokers to reach the publisher, there are fewer opportunities for them to be covered. On the other hand, subscriptions farther away are less likely to reach the publisher, since there are more opportunities for them to be covered by other subscriptions at intermediary brokers. As a result, subscriptions closer are overvalued and receive more publications, while subscriptions further away are undervalued and receive fewer publications. In both cases, the fairness metric increases as these subscriptions deviate from the expected outcome. AncestorFairK is subject to this trend to a lesser degree, but for a different reason. The fairness ranker relies on the meta-information which is propagated through subscriptions. In the workload used, staleness in the meta-information tends to underestimate the number of subscriptions. Because these inaccuracies accumulate over multiple hops, subscribers which are closer to the publisher tend to be overvalued, while ones farther away are undervalued. Therefore, we observe the discrepancy in fairness for the closest subscriptions because they receive more publications than expected, while the subscriptions further away loses fairness as they receive fewer publications than expected.

Summary: AncestorFairK, which uses our fair ranker, outperforms AncestorCoverK, which selects randomly from known subscriptions, by an average of 224% in the target range for k and becomes increasingly better as the value of k increases. Both solutions tend to send publications to closer subscriptions more frequently,
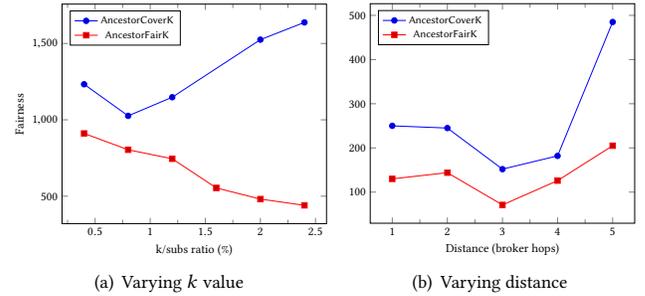

Figure 6. Fairness evaluation (250 subscribers)

while subscriptions further away are less likely to be selected. AncestorFairK is less sensitive in that regard. Additionally, the diameter of pub/sub overlays can be controlled to lessen this effect [23].

5.5 Processing overhead

Figures 7(a) and 7(b) show the end-to-end latency of publication delivery for 250 and 500 subscriptions, respectively. We first note that RegularK is outperformed by AncestorCoverK in lower values of k. This indicates that the processing overhead for top-k is affected by the number of subscriptions stored at a broker. As the value of k grows larger, the covering effectiveness (as explained in Section 5.3) decreases which affects the top-k overhead. Using values extracted from the RegularCover and AncestorCoverK results, we evaluate the overhead of top-k processing at a broker to be 3% of the total latency, if the same number of subscriptions is stored at each broker. The most significant part of the top-k processing overhead is therefore due to the drop in covering performance, which increases the number of subscriptions residing at each broker, which in turn affects the matching and processing times.

The overhead of the fair sorter has a significant impact on latency in higher values of k tested. The overhead follows a superlinear growth with the number of subscriptions, as observed in Figures 8(a) and 8(b). This is in line with the $O(n \log n)$ running time of our weighted shuffle implementation which is invoked at each broker when processing a publication. This result suggests that adaptive solutions which can either limit the number of subscriptions to be shuffled or can avoid the shuffle selection at certain brokers will improve publication latency. It also underlines the importance of our optimization for the weighted shuffle procedure. Nevertheless, the experiments indicate that for values of k less than 1% of the
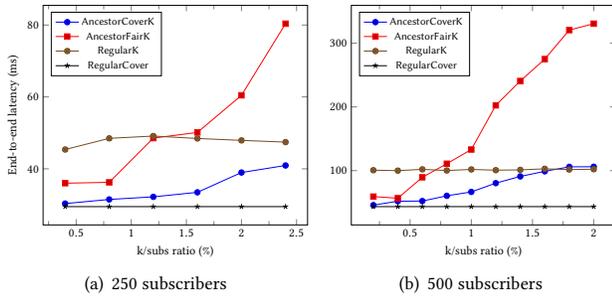
(a) 250 subscribers (b) 500 subscribers

Figure 7. End-to-end latency

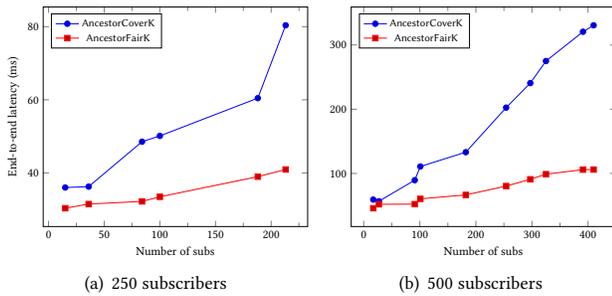(a) 250 subscribers (b) 500 subscribers

Figure 8. Latency sensitivity

total number of subscriptions, it is possible to maintain fairness and still perform better than a top-k implementation with no covering.

Summary: By comparing AncestorCoverK to RegularK, we observe that top-k filtering improves end-to-end latency by an average of 25% in the range of k tested, with the benefit lessening as the covering effectiveness decreases. The fair ranker has a significant impact on the processing overhead at a broker and is sensitive to the number of subscriptions processed. This underscores the importance of optimizing the weighted shuffle procedure, which is performed at every broker while processing publications. In general, in order to support top-k filtering and subscription covering, one should decide between AncestorCoverK or AncestorFairK, based on the trade-off between fairness and processing overhead.

6 Conclusions

Many applications found in domains such as targeted marketing, software maintenance, and job scheduling require a communication middleware that supports selective filtering of data. To this end, this paper extends publish/subscribe subscription filtering with extended top-k semantics that supports general ranking objectives, such as relevance and fairness. In particular, we introduce RankCoverK, a novel rank-cover algorithm to construct an efficient covering partially ordered set (poset) which supports top-k relevance scoring. We further develop AncestorCoverK, an ancestor counting optimization to effectively prune and reduce the cover size, and extend our solution to consider fairness with AncestorFairK. We conclude our work with an extensive sensitivity analysis of our algorithms incorporated in PADRES, an open-source distributed publish/subscribe system, that illustrate the importance of our proposed techniques with respect to subscription traffic reduction, end-to-end latency,

and fairness. AncestorCoverK is scalable and retains over 81% of the covering benefit when using a k set at 1% selectivity. AncestorFairK, which is AncestorCoverK equipped with our fair ranker, provides increased fairness in exchange for higher processing overhead.

References

[1] 2013. The IBM Strategy. http://www.ibm.com/annualreport/2013/. (2013).
[2] Marcos K. Aguilera, Robert E. Strom, Daniel C. Sturman, Mark Astley, and Tushar D. Chandra. Matching events in a content-based subscription system. In *PODC'99*.
[3] Brian Babcock and Chris Olston. Distributed top-k monitoring. In *SIGMOD'03*.
[4] Raphaël Barazzutti, Thomas Heinze, Andre Martin, Emanuel Onica, Pascal Felber, Christof Fetzer, Zbigniew Jerzak, Marcelo Pasin, and Etienne Riviere. Elastic Scaling of a High-Throughput Content-Based Publish/Subscribe Engine. In *ICDCS'14*.
[5] Alexis Campailla, Sagar Chaki, Edmund Clarke, Somesh Jha, and Helmut Veith. Efficient filtering in publish-subscribe systems using binary decision diagrams. In *ICSE'01*.
[6] Pei Cao and Zhe Wang. Efficient top-K query calculation in distributed networks. In *PODC'04*.
[7] Antonio Carzaniga, David S. Rosenblum, and Alexander L Wolf. 2001. Design and Evaluation of a Wide-Area Event Notification Service. *TOCS* (2001).
[8] Lisi Chen and Gao Cong. Diversity-Aware Top-k Publish/Subscribe for Text Stream. In *SIGMOD '15*.
[9] L. Chen, G. Cong, X. Cao, and K. L. Tan. Temporal Spatial-Keyword Top-k publish/subscribe. In *ICDE'15*.
[10] William Culhane, K. R. Jayaram, and Patrick Eugster. Fast, Expressive Top-k Matching. In *Middleware '14*.
[11] Marina Drosou, Evaggelia Pitoura, and Kostas Stefanidis. Preferential Publish/Subscribe. In *PersDB'08*.
[12] Marina Drosou, Kostas Stefanidis, and Evaggelia Pitoura. Preference-aware publish/subscribe delivery with diversity. In *DEBS'09*.
[13] Françoise Fabret, Hans-Arno Jacobsen, François Llirbat, João Pereira, Kenneth A. Ross, and Dennis Shasha. Filtering algorithms and implementation for fast pub/sub systems. In *SIGMOD'01*.
[14] Ronald Fagin, Amnon Lotem, and Moni Naor. Optimal aggregation algorithms for middleware. In *PODS'01*.
[15] E. Fidler, Hans-Arno Jacobsen, Guoli Li, and Serge Mankovski. The PADRES Distributed Publish/Subscribe System.. In *ICFI'05*.
[16] Marcus Fontoura, Suhas Sadanandan, Jayavel Shanmugasundaram, Sergei Vassilvitski, Erik Vee, Srihari Venkatesan, and Jason Zien. Efficiently evaluating complex Boolean expressions. In *SIGMOD'10*.
[17] Albert Greenberg, James Hamilton, David A. Maltz, and Parveen Patel. The Cost of a Cloud: Research Problems in Data Center Networks. In *SIGCOMM'08*.
[18] Ihab F. Ilyas, George Beskales, and Mohamed A. Soliman. 2008. A survey of top-k query processing techniques in relational database systems. *Comput. Surv.* (2008).
[19] Chamikara Jayalath, Julian James Stephen, and Patrick Eugster. Atmosphere: A universal cross-cloud communication infrastructure. In *Middleware'13*.
[20] Guoli Li, Shuang Huo, and Hans-Arno Jacobsen. A Unified Approach to Routing, Covering and Merging in Publish/Subscribe Systems Based on Modified Binary Decision Diagrams. In *ICDCS'05*.
[21] Ashwin Machanavajjhala, Erik Vee, Minos Garofalakis, and Jayavel Shanmugasundaram. Scalable ranked publish/subscribe. In *VLDB'08*.
[22] Sebastian Michel, Peter Triantafillou, and Gerhard Weikum. KLEE: a framework for distributed top-k query algorithms. In *VLDB'05*.
[23] Melih Onus and Andréa W Richa. 2011. Minimum maximum-degree publish-subscribe overlay network design. *TON* (2011).
[24] Navneet Kumar Pandey, Kaiwen Zhang, Stephane Weiss, Hans-Arno Jacobsen, and Roman Vitenberg. Minimizing the Communication Cost of Aggregation in Publish/Subscribe Systems. In *ICDCS'15*.
[25] Krešimir Pripužić, Ivana Podnar Žarko, and Karl Aberer. Top-k/w publish/subscribe: finding k most relevant publications in sliding time window w. In *DEBS'08*.
[26] Mohammad Sadoghi and Hans-Arno Jacobsen. BE-Tree: An Index Structure to Efficiently Match Boolean Expressions over High-dimensional Discrete Space. In *SIGMOD'11*.
[27] Mohammad Sadoghi and Hans-Arno Jacobsen. Relevance Matters: Capitalizing on Less (Top-k Matching in Publish/Subscribe). In *ICDE'12*.
[28] Vinay Setty, Gunnar Kreitz, Roman Vitenberg, Maarten van Steen, Guido Urdaneta, and Staffan Gimåker. The Hidden Pub/Sub of Spotify: (Industry Article). In *DEBS'13*.
[29] Abhishek Verma, Luis Pedrosa, Madhukar Korupolu, David Oppenheimer, Eric Tune, and John Wilkes. Large-scale Cluster Management at Google with Borg. In *EuroSys'15*.
[30] Xiang Wang, Ying Zhang, Wenjie Zhang, Xuemin Lin, and Zengfeng Huang. Skype: Top-k Spatial-keyword Publish/Subscribe over Sliding Window. In *VLDB'16*.

[31] Steven Whang, Chad Brower, Jayavel Shanmugasundaram, Sergei Vassilvitskii, Erik Vee, Ramana Yerneni, and Hector Garcia-Molina. Indexing Boolean Expressions. In *VLDB'09*.
[32] Tak Yan and Hector Garcia-molina. 1994. Index Structures for Selective Dissemination of Information Under the Boolean Model. *TODS* (1994).
[33] Kaiwen Zhang, Vinod Muthusamy, Mohammad Sadoghi, and Hans-Arno Jacobsen. Subscription Covering for Relevance-based Filtering in Content-Based Publish/Subscribe Systems. In *ICDCS'17*.
[34] Kaiwen Zhang, Mohammad Sadoghi, Vinod Muthusamy, and Hans-Arno Jacobsen. Distributed Ranked Data Dissemination in Social Networks. In *ICDCS'13*.

A Diversity as a ranking objective

Diversity is another possible ranking criteria. Suppose there is a distance metric that quantifies the pairwise similarity of subscriptions: $distance(s_i, s_j)$. Each subscription s in S is ranked in descending order using $\sum_{j=1}^{|N|} distance(s, s_j)$, which is the cumulative sum of pairwise distances between s and all subscriptions in S. The resulting k-sized subset of S is *diverse*, the intention being to avoid delivering the publication to subscriptions with similar interests.

The framework in this paper can support diversity, but it remains to devise a reasonable and efficient pairwise distance metric among subscriptions. The definition and justification of such a metric is out of the scope of this paper, and is treated as an interesting avenue for future work.

B Proofs of correctness

Below are the proofs for the theorems found in the paper.

Theorem B.1. *If for each downstream broker $C = \hat{S}$, where C is the set of forwarded subscriptions to upstream broker B, then B can compute the top-k subscriptions for any given publication.*

Proof. Suppose by way of contradiction that there exists a subscription s' is in the top-k set for the set of matching subscriptions S_j for a given publication p_j, and $s_i \notin C$ for all downstream brokers. s' must come a certain downstream broker B'. Note that the set of matching subscriptions S'_j at broker B' is a subset of the set of matching subscriptions S_j at B, since B is an upstream broker from B'. Since s' is in the top-k set of S_j, then s' must also be in the top-k set in S'_j. If this is the case, however, $s' \in \hat{S}'$ and thus $s' \in C'$. This contradicts the original supposition. □

Theorem B.2. *The set $\hat{S} = \{s_i | depth(s_i) \le k\}$ contains the set of top-k subscriptions for any possible publication by advertisement a.*

Proof. Suppose by way of contradiction that there exists a subscription $s' \notin \hat{S}$ that is among the top-k highest ranking subscriptions for some publication. Consider the set of subscriptions $\tilde{S}$ in the path from the root of the poset to s'. By our supposition, there are at least k such subscriptions in $\tilde{S}$. Also by the definition of the poset construction, for any publication that matches s', each subscription in $\tilde{S}$ will also match p and have a higher score than s'. Therefore s' cannot be among the top-k most relevant subscriptions for p. □

Theorem B.3. *The set $\hat{S} = \{s_i | ancestors(s_i) < k\}$ contains the set of top-k subscriptions for any possible publication from advertisement a.*

Proof. The proof is almost identical to that of Theorem B.2. □