Content-based XML Data Dissemination

Guoli Li   Shuang Hou   Hans-Arno Jacobsen

Middleware Systems Research Group, University of Toronto

{gli,shou}@cs.toronto.edu   jacobsen@eecg.toronto.edu

XML-based data dissemination networks are rapidly gaining momentum. In these networks XML content is routed from data producers to data consumers throughout an overlay network of content-based routers. Routing decisions are based on XPath expressions (XPEs) stored at each router. To enable efficient routing, while keeping the routing state small, we introduce advertisement-based routing algorithms for XML content, present a novel data structure for managing XPEs, especially apt for the hierarchical nature of XPEs and XML, and develop several optimizations for reducing the number of XPEs required to manage the routing state. The experimental evaluation shows that our algorithms and optimizations reduce the routing table size by up to 90%, improve the routing time by roughly 85%, and reduce overall network traffic by about 35%. Experiments running on PlanetLab show the scalability of our approach.

1.1 Introduction

Over the past decade, XML has rapidly evolved as the standard for data representation and exchange. XML marked-up message traffic in intranets and on the Internet ranges from insurance claims, health-care requests, corporate memos, online ads to news items and entertainment information. The standardization of the mark-up language, the wide range of related standards, and the wide-spread adoption of this technology are further amplifying the network externalities created by this technology.

XML-based data dissemination networks are starting to become a reality. In a dissemination network, data messages, marked-up in XML, are routed based on filter expressions stored at intermediate nodes that indicate where the XML message is to be routed to. These routing nodes are often referred to as *content-based routers* or *brokers*, as they route messages to interested recipients by inspecting the message content. Filter expressions, often expressed as XPath expressions (XPEs), are submitted by data consumers who express interest in receiving certain kinds of documents. This architecture is depicted in Figure 1.1. For instance, a globally operating insurance company with many branch offices distributed world-wide is linked by an overlay network of content-based routers that comprise the XML dissemination network. An insurance claim, an insurance bid, or a request for proposal can be submitted anywhere into the overlay network (e.g., by a third party insurance broker or an online client) and be routed toward a currently online, specific expert employee, speaking the same language as the requester. Note,

the latter constraints are expressed as XPE filter expressions against which the XML document is evaluated in transit. This design fully decouples information requesters and information providers, avoids a single point of control and a single point of failure, and increases scalability due to decentralization and distribution.
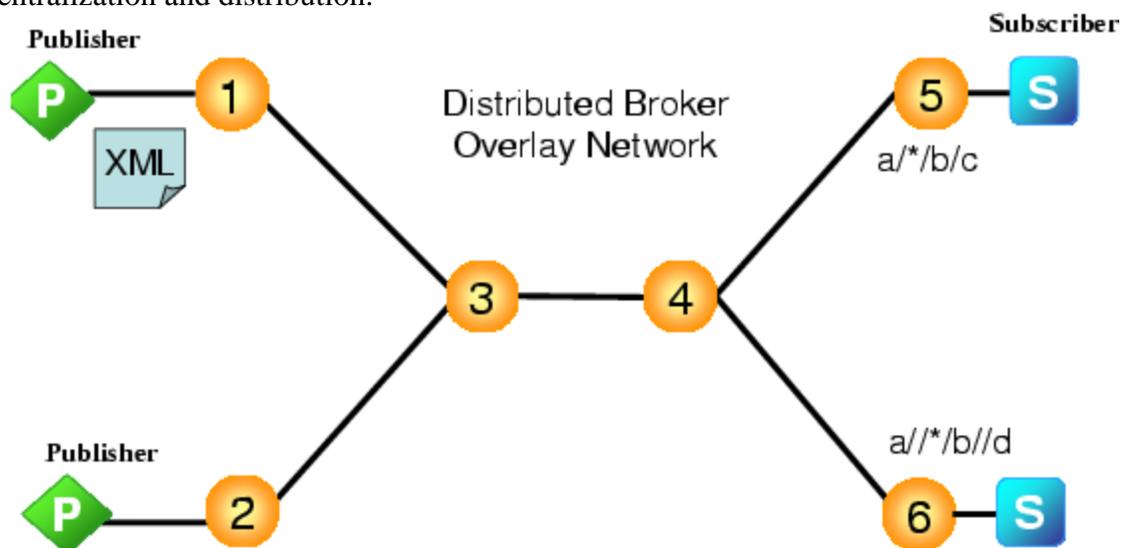

Figure 1.1: Distributed Dissemination Network

The fundamental concepts underlying the content-based dissemination of XML messages are the algorithms and data structures for matching and routing of XML documents against XPEs, the definition of advertisements that efficiently summarize the kind of XML documents that data producers will publish, the interpretation and the development of algorithms for intersecting advertisements with subscriptions, and the definition of covering and merging of XPEs.

This chapter addresses the XML/XPath routing problem. More specifically, this chapter focuses on the problem of efficiently routing an XML document emitted from a data producer at one point in the network to a set of data consumers located anywhere throughout the network. Prior to receiving XML documents, consumers must have expressed interest in receiving XML documents by registering XPEs with the network. This problem statement is akin to the well-known publish/subscribe matching problem. However, the main difference here is that in the case of data dissemination networks, there exists no one single centralized publish/subscribe system, but a network of content-based routers (i.e., a network or federation of publish/subscribe systems.) In the dissemination network, XML documents are routed based on their content and not based on IP address information, which is, due to the completely decoupled design, not available – all routing decisions are exclusively based on content information. Figure 1.2 provides an overview of the dissemination network this chapter assumes. In the overlay network depicted in Figure 1.2 each content-based router, referred to as broker, only knows its neighbors (i.e., in terms of IP network address information.) However, none of the clients – neither data producers nor data consumers – know about each other or about the network topology, except the router they connect to. This design eases system management while offering flexibility to the application design.

In the context of XML-based data dissemination, one of the main challenges is the ability to efficiently deliver relevant XML documents to a large and dynamically changing group of consumers. Not only data sinks (consumers) can change dynamically, but also data sources (producers) can change dynamically. Existing producers and consumers can go away and

re-appear, and new ones can appear at any time. Furthermore, the challenges are to evaluate high rates of XML messages against large numbers of possibly overlapping XPEs to decide where to send an XML message to. Moreover, XPEs can dynamically change, which has to be reflected in the whole dissemination network. A further challenge is that XPEs may operate over XML document structure, XML tag attributes, and XML document content. XPath expressions can define very complex expressions.

Centralized XML filtering (Altinel & Franklin, 2000; Chan et al., 2002b; Diao et al., 2003; Hou & Jacobsen, 2006) and distributed query-based XML retrieval approaches (Chan et al., 2002a; Koloniari & Pitoura, 2004; Koudas et al., 2004) have found wide-spread interest, but do not address the distributed, content-based routing problem articulated above and addressed in this chapter. ONYX is a query-based XML retrieval approach for XML data dissemination in distributed environments (Diao et al., 2004). It is closely related to our approach, but complementary in objectives. ONYX aims at reducing the XML message size. ONYX achieves this through only disseminating parts of an XML message selected by subscriber queries. Several techniques are presented and evaluated to achieve this objective. In the context of our work, subscriber queries are meant to select messages that match the query, however, the subscriber requires the full message and not just parts of it. It is therefore difficult to quantitatively compare both approaches. Content-based routing in distributed publish/subscribe architectures (Carzaniga et al., 2004; Cugola et al., 2001; Li & Jacobsen, 2005; Mühl, 2001), has been studied for non-XML-based data. Their operational model assumes sets of attribute-value pairs joined by Boolean operators. It is not at all obvious how to extend these approaches for semi-structured data, especially due to the hierarchical data model of XML.

Our own prior research on content-based routing and matching develops architectures, algorithms and protocols for content-based routing of non-XML based data (Cheung & Jacobsen, 2006; Fidler et al., 2005; Li et al., 2005; Li & Jacobsen, 2005; Liu & Jacobsen, 2002; Liu & Jacobsen, 2004; Liu et al., 2009). It is a non-trivial problem, as we argue in this chapter, to extend these approaches to the hierarchical structure of XML. Our prior work on developing efficient matching algorithms for XML addresses part of the problem (Hou & Jacobsen, 2006), and another work on efficient routing protocols addresses the important problem of efficiently computing routing decisions, inducing advertisements from XML DTDs, and determining covering and merging relations (Li et al., 2008a; Li et al., 2008b). This chapter is an extended version of our prior work (Li et al., 2008a). This chapter combined with our earlier work on XML matching (Hou & Jacobsen, 2006) and routing (Li et al., 2008a) comprises all components required to build an efficient content-based router for an XML data dissemination network. The research presented in this chapter complements our PADRES content-based publish/subscribe system effort (Fidler et al., 2005; Jacobsen, 2006; Li et al., 2007b). PADRES has to date not investigated the routing of XML content against XPEs. PADRES is based on a proprietary subscription language and publication data model. Subscriptions are conjunctions of predicates, which define filter constraints over attribute value pairs. Publications are sets of attribute value pairs. Unique to PADRES are its capabilities to subscribe to past events (i.e., query the event space for historic data) (Li et al., 2007a; Li et al., 2008c), to correlate past and future events, and to detect composite events by correlating individual atomic and composite events. PADRES achieves this flexibility through a powerful subscription language based on composite subscription expressions that form the basis of its SQL-like subscription language. A composite subscription defines the subscriber's interest in composite events. PADRES proposes novel algorithms for the routing of composite subscriptions and the detection of composite events based on data source intensity and past matching statistics. PADRES also proposes routing protocols

that exploit cyclic overlay topologies to handle failures, react to congestion and peak loads, and to generally offer a robust publish/subscribe service. In PADRES, robustness is achieved by supporting alternate message routing paths (Li et al., 2008b), load balancing techniques to distribute load (Cheung & Jacobsen, 2008), and fault resilience techniques to react to broker failures (Sherafat Kazemzadeh & Jacobsen, 2007; Sherafat Kazemzadeh & Jacobsen, 2008). PADRES also studies client mobility (Muthusamy & Jacobsen, 2005; Muthusamy & Jacobsen, 2007; Petrovic et al., 2005), service selection (Chau et al., 2008; Hu et al., 2008; Muthusamy et al., 2007; Muthusamy & Jacobsen, 2008) and resource and service discovery (Yan et al., 2009). All these concepts complement the work presented in this chapter. However, PADRES does not support the routing of XML messages against XPEs, which is the missing link to briding and combining the work presented in this chapter with PADRES.

In this chapter, we develop algorithms for dissemination of XML data throughout a network of content-based routers towards data consumers who have specified their interests through XPEs. Our contributions are: first, we adapt the use of advertisements to optimize data dissemination. While this idea is common in the publish/subscribe literature (Carzaniga et al., 2004; Cugola et al., 2001; Mühl, 2001), it is not clear how to extend the concepts to the data model of XML. We demonstrate how to use the XML Document Type Definition (DTD) to generate advertisements about the information a data producer is going to publish. We distinguish between a non-recursive and a recursive case depending on the DTD defining the data emitting source. We then develop advertisement-based routing algorithms for both cases. Second, we propose a novel data structure to maintain XPEs by identifying the covering relations among them. We present covering algorithms for XPEs to reduce the routing table size stored at each router and speed up routing computation in the routers. Third, we present an optimization of merging similar XPEs to further reduce routing computation. Finally, we perform a detailed experimental evaluation of our approach on an overlay network comprised of 127 XML routers deployed over a cluster with 20 nodes and deployed on PlanetLab. Our experimental results demonstrate the effectiveness of the approach by reducing the routing table size by up to 90% and improving the routing time by roughly 85%.

1.2 Background

1.2.1 Content-based Routing

Content-based publish/subscribe systems provide a flexible and extensible environment for information exchange (Carzaniga et al., 2004; Cugola et al., 2001; Fidler et al., 2005; Jacobsen, 2006; Mühl, 2001). Publishers and subscribers are clients to the publish/subscribe system, are loosely coupled in space and time, and have no knowledge of each other. Messages in content-based publish/subscribe systems are routed based on their content rather than the IP address of their destinations. In order to handle a large amount of dynamic information and reduce the network traffic many optimization techniques, such as advertisements (Carzaniga et al., 2004), covering technique and merging technique (Carzaniga et al., 2004; Cugola et al., 2001; Mühl, 2001) have proven to offer significant benefits for non-XML based publish/subscribe systems. While conceptually, these ideas apply to XML-based data as well, it is not obvious how to apply these concepts to XML due to the structural complexity of XML data. These are the challenges addressed by this chapter.
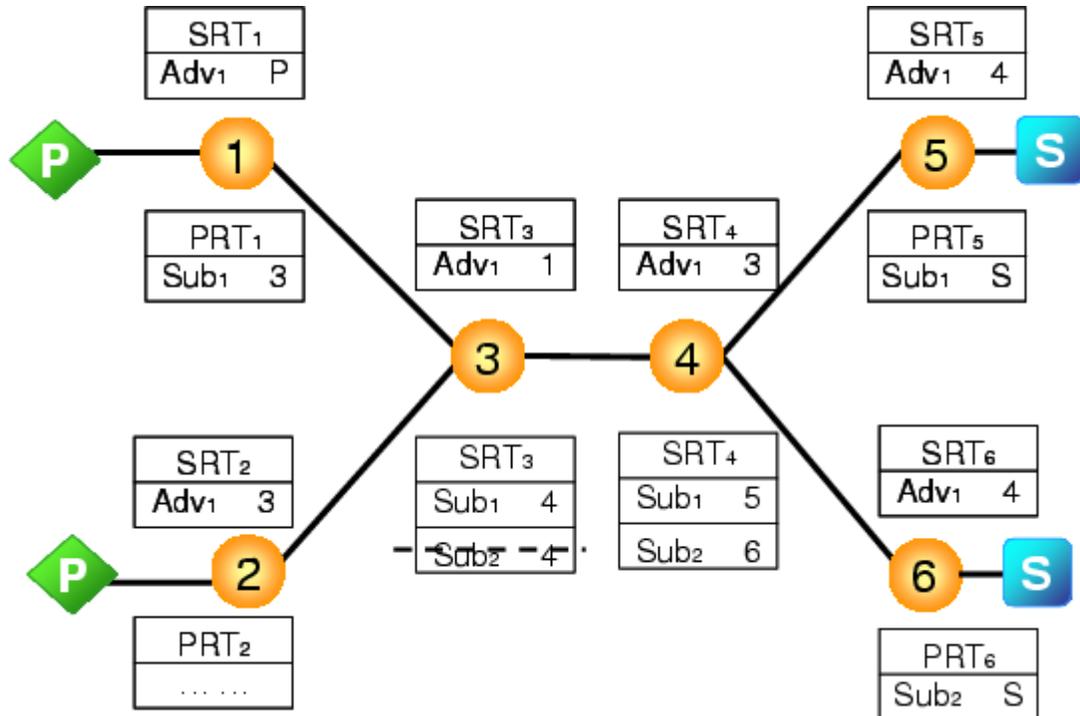
Figure 1.2: Content-based Routing

In advertisement-based publish/subscribe systems, advertisements are specifications of information that the publisher publishes in the future. Advertisements are flooded in the publish/subscribe overlay. The common assumption is that the number of advertisements is much smaller than the number of subscriptions and publications. Advertisements are used to avoid broadcasting subscriptions in the network, so that subscriptions are only routed to the publishers who advertise what the subscribers are interested in. Subscriptions define filters to select publications of interest. Matching publications are delivered to subscribers along the paths built by subscriptions. Figure 1.2 shows a scenario for advertisement-based content-based routing. The subscription routing table (SRT) consisting of <advertisement, last-hop>-tuples stores advertisements in order to route subscriptions. Publications trace back along the path setup by subscriptions to interested subscribers. The publication routing table (PRT) maintains the path information. For example, in Figure 1.2, advertisement adv_1 is broadcast in the network, and is stored on each broker of the network with a different last hop. Consequently, subscriptions that match adv_1 are routed according to these last hops (e.g., sub_1 is routed along the link 5−4−3−1). Note that the subscription sub_1 is not forwarded to Brokers 2 and 6, respectively, since adv_1 indicates that matching publications are from Broker 1. Therefore, publication pub_1 is routed along the reverse path 1−3−4−5 to the subscriber. In the rest of this chapter, we use the notations $P(s)$ and $P(a)$ to refer to the set of publications that match subscription s and advertisement a, respectively.

1.2.2 Covering and Merging

The goal of covering-based routing is to remove redundant subscriptions from the network in order to obtain a compact routing table and reduce the network traffic (Carzaniga et al., 2004; Li

et al., 2005; Mühl, 2001). In Figure 1.2, if subscription sub_1 covers subscription sub_2 at Broker 4, sub_2 is not forwarded to Broker 3. That is, we can safely remove sub_2 at Broker 3 obtaining a compacter routing table while maintaining the same information delivery behavior. All publications matching sub_2 must also match sub_1. A formal definition of the covering relation is as follows: A subscription sub_1 covers sub_2, if and only if, $P(sub_1) \supseteq P(sub_2)$, denoted as $sub_1 \sqsupseteq sub_2$. The covering relation defines a partial order on the set of all subscriptions with respect to $\sqsupseteq$. Since advertisements have the same format as subscriptions, the covering relations among advertisements can be defined in the same manner.

If two subscriptions are not in a covering relation, but their publication sets intersect each other, the two subscriptions can be merged to a more general subscription, which covers the original subscriptions. Suppose subscription sub_m is a merger of sub_1 and sub_2, then we have $P(sub_m) \supseteq P(sub_1) \cup P(sub_2)$. There are two kinds of mergers. If the publication set of the merger is exactly equal to the union of the publication set of the original subscriptions, the merger is a *perfect merger*; otherwise, if $P(sub_m) \supset P(sub_1) \cup P(sub_2)$, it is an *imperfect merger*. After merging, only the merger is forwarded into the network. The merging technique (Li et al., 2005; Mühl, 2001) is used for further minimizing the routing table size, since the merger may introduce new covering relations among subscriptions. Covering and merging are complementary routing optimizations.

1.3   Related Work

A large body of work has focused on developing publish/subscribe-style matching algorithms for evaluating an XML message against a set of XPEs (Altinel & Franklin, 2000; Bruno et al., 2003; Chan et al., 2002b; Diao et al., 2003; Hou & Jacobsen, 2006). However, all these approaches exclusively address centralized matching architectures, not the distributed, content-based XML dissemination networks we address in this work. While matching is an integral step in a content-based router, other routing operations studied in this chapter are equally important in distributed data dissemination architecture. Thus, our work complements matching algorithms for the design of a content-based XML router.

Recent research has focused on XML data dissemination (Altinel & Franklin, 2000; Diao et al., 2003; Diao et al., 2004; Fenner et al., 2005; Koloniari & Pitoura, 2004; Koudas et al., 2004; Snoeren et al., 2001). Among the earlier work on XML dissemination, XFilter (Altinel & Franklin, 2000) is probably among the first approaches. It defines a finite state machine (FSM) for each XPath query, then it proposes an index over these FSMs. These FSMs are executed concurrently for each XML document. When a matching state is reached, the XML document is delivered to the subscriber who issued the query. YFilter (Diao et al., 2003) improved the matching performance by proposing one unique FSM for all queries, which allowed common query paths to be processed only once. The matching performance is improved by sharing common query paths. Other FSM-baed approaches use different techniques for building the FSMs for queries, as well as different type of FSMs. For instance, Deterministic Finite Automaton (DFA) have been built in Green *et al.* from XPath queries (Green et al., 2004). The authors construct a DFA lazily. A lazy DFA is one whose states and transitions are computed from the corresponding NFA at runtime, not at compile time. The number of states in the lazy DFA is small, which guarantees the efficient matching of XML documents.

Holistic information is first used in the matching process by FiST (Kwon et al., 2005), which performs a different, bottom-up approach based on ordered twig patterns using Prüfer sequences. FiST organizes the sequences into a dynamic hash-based index for efficient filtering. Another FSM-based approach BUFF also uses a bottom-up approach (Moro et al., 2007), but it avoids translating documents and queries to Prüfer sequences, but employs a novel pruning technique based on lower and upper bound estimations to reduce the query space.

Another methodology is to aggregate the queries using some index technique. XTrie supports efficient filtering of streaming XML documents based on XPath expressions (Chan et al., 2002b). The authors proposed a novel index structure, termed XTrie, for large-scale filtering with complex XPath expressions, not limited to simple, single-path specifications. It supports both ordered and unordered matching of XML documents. XTrie is space-efficient since the space cost of XTrie is dominated by the number of substrings in each tree pattern.

There are some related research in peer-to-peer networks. Koloniari *el al.* present a decentralized approach for XML dissemination in a peer-to-peer network (Koloniari & Pitoura, 2004). However, in their approach queries are severely restricted in that no wildcards are allowed. Koudas *et al.* propose a flexible routing protocol for XML routers to enable scalable XPath query and update processing in a data-sharing peer-to-peer network (Koudas et al., 2004). Both approaches are solutions to the *location problem*. The location problem states that given a dynamic collection of XML database servers and an XPath query, find the databases that contain data relevant to the query. Our approach evaluates an XML document against a set of XPath queries, and decides where to route the XML document.

Instead of matching XML documents against individual queries, generating a compact set of XPath queries from a given set of XPEs could improve the matching performance as well. This problem is similar to the covering and merging problem discussed in this chapter. Query containment and aggregation address how a compact set could be generated theoretically and practically (Chan et al., 2002a; Dong et al., 2003; Miklau & Suciu, 2004). In Miklau *et al.* it is shown that for a simple fragment of XPath that contains descendant axis (//), wildcards(*), and qualifiers([...]), but without either tag variables or disjunctions (Miklau & Suciu, 2004), query containment is coNP-complete. If any of the above constructs is dropped, (Amer-Yahia et al., 2002) shows that the query containment problem can be reduced to PTIME. Dong *et al.* (Dong et al., 2003) further studies the nested XML queries and shows for queries with fixed nesting depth, the containment problem is coNP-complete.

One of the challenges brought by the query containment and aggregation problem is to guarantee the equivalence between the compact set and the original set. That is, the aggregate query set does not introduce or control false positives (i.e., takes XML documents not originally matched) or false negatives (i.e., misses XML documents originally matched.) The efficient tree-pattern aggregation algorithm proposed in Chan *et al.* makes effective use of document-distribution statistics in order to compute a precise set of aggregate tree patterns within the allocated space budget (Chan et al., 2002a). However they try to minimize the loss in precision due to the aggregation, not to avoid it. While we try to eliminate the false positives and false negatives in the covering approach and quantify the imprecision in the merging approach. These are non-trivial extensions to their work. Furthermore, the tree aggregation approach does not address the generation of advertisements from DTDs, which is central to our approach. Chand *et.al.* discuss a scalable protocol for XML-based data dissemination (Chand & Felber, 2003). They explore subscription aggregation to provide efficient message routing and use a tree structure to maintain the covering relation among subscriptions. They maintain part of the

covering relation in their substitution tree. We use extend the tree structure with *super pointers* to cache a complete covering relationship, which is used for covering-based routing.

Advertisement-based techniques for optimizing content-based routing have been developed in the area of distributed publish/subscribe (Carzaniga et al., 2004; Cugola et al., 2001; Mühl, 2001). It has been demonstrated that the network traffic and routing table size can be reduced by using different routing strategies, including advertising, covering and merging techniques. However, the main differences between these approaches and our approach lie in the subscription language and publication data format. Our approach is based on the hierarchical, tree-based XPath and XML model; while the traditional content-based routing approaches operate with attribute-value pairs and predicate constraints over these pairs. The advertising carried out by XML and XPath data sources is different and more complex than predicate-based languages, for the hierarchical and recursive structure of the model needs to be taken into account. A DTD of an XML document does not have an equivalent in traditional publish/subscribe approaches. Galanis *et al.* explore XML data dissemination based on a distributed hash table (DHT) (Galanis et al., 2003), which is a decentralized distributed system that provides a lookup service similar to a hash table: (key, value) pairs are stored in the DHT, and any participating node can efficiently retrieve the value associated with a given key. The authors use data summaries to ensure that queries are only sent to relevant servers. These data summaries could be taken as a form of advertisements. The data summary is generated from an XML document, so the expressiveness of the data summary is part of the DTD. Our contribution is to generate a complete advertisement set once from a DTD for all related XML documents.

ONYX (Diao et al., 2004) and XTreeNet (Fenner et al., 2005) are systems similar in conception to the one we describe in this chapter. However, their objectives are quite different making a comparison difficult. ONYX (Diao et al., 2004) describes an architecture to deploy XML-based services on an Internet-scale. The approach is based on performing incremental message transformations to reduce message size instead of sending and receiving the entire XML message, published by a data source. ONYX only delivers the parts of the message actually selected by the data sinks' subscribing queries. For many applications, this approach is not feasible, as the entire message published by a source needs to be delivered in its entirety to all subscribing data sinks, which is the message delivery semantic realized by our system and algorithms. To further reduce message size and processing cost, ONYX investigates various representations for XML messages. While a binary message representation will certainly speed up XML message processing, we have found in prior work that XML processing, such as parsing, is not the dominating cost for an XML router (Hou & Jacobsen, 2006) and optimizations of that component are therefore of questionable utility in this context. ONYX uses an NFA-based operator network for representing routing tables. This approach supports the sharing of common prefixes among queries. Our approach goes beyond this by identifying all covering relations among queries processed when constructing the routing tables, so that all covered queries are eliminated completely from the routing computation. Moreover, our work introduces merging and advertising for XML data not addressed in the earlier approach. Our unique contribution is to enable these techniques for the XML data model, which is fundamentally different from the data models underlying non-XML-based content routing approaches, such as (Carzaniga et al., 2004; Cugola et al., 2001; Li et al., 2005; Mühl, 2001). XTreeNet (Fenner et al., 2005) elegantly unifies the publish/subscribe and the query/response model in an XML-aware overlay network. XTreeNet proposes a dissemination protocol to avoid repeatedly matching XML data against queries at intermediate routers. This is an orthogonal optimization that our approach can also employ by attaching the path of overlay hops to the XML query when the subscription tree is

constructed in the network. XML messages only match against the queries at the first router, and are forwarded along the subscription paths to the subscribers.

To the best of our knowledge, our work is the first to address the problem of advertisement for XML data dissemination. However, the advertisement-based routing has been widely studied in the area of content-based publish/subscribe system (Carzaniga et al., 2004; Cugola et al., 2001; Mühl, 2001). They reduce the network traffic and routing table size by using different routing strategies, including advertisement, covering and merging techniques. In contrast to the format of advertisement, conjunction of attribute-value pairs, in the context of content-based publish/subscribe system, our chapter focuses on a more complex format (as the advertisement in our work is defined as recursive and non-recursive with respect to different DTDs), which includes both data contents and structure. The query aggregation scheme given in (Chan et al., 2002a) addresses the problem as part of our work, and they do not address the advertisement-based optimization and support powerful language like recursive advertisement format.

1.4   Advertisement-based Routing

Upon receiving a subscription, a broker matches the subscription against its advertisements. If there is an advertisement whose publication set intersects that of the subscription, it means there is a *match* between the subscription and the advertisement. The broker then routes the subscription to the broker where the advertisement came from.

1.4.1   XML-based Advertisements

In the context of XML/XPath routing, advertisements are generated by exploiting DTD information. The purpose of a DTD is to define the legal building blocks of an XML document. The main building blocks of XML documents are elements surrounded by tags, e.g., *<root>...</root>*, where *root* is the element name in the document. All elements appearing in the XML document must be defined in the corresponding DTD, which determines the structure of elements and their sequence in the document. In this chapter, our discussion focuses on the main building block – elements. Our approach could be easily extended to element attributes and content Hou & Jacobsen (2006), which we omit due to space limitations.

In this chapter, we use the common interpretation of an XML document as a tree of nodes and consider each path from the root node to a leaf node. Thus, we decompose each XML document into a set of XML paths and each path is represented as $e=/t_1/t_2/.../t_n$, where t_i is the XML element name. These paths are extracted from the document before the publisher submits the document to the network. Thus, a publication routed in our system is actually an XML path annotated with a *pathId* and *docId*. This is transparent to publishers and subscribers who handle entire XML documents. Publishers submit entire XML documents, commonly referred to as publications, and subscribers submit XPath expressions (XPEs), commonly referred to as subscriptions. We use the terms XPE and subscription interchangeably in the rest of this chapter.

We use an absolute XPath expression without //-operators as the format of advertisements in the context of XML/XPath data routing. Note that this is not a restriction of our subscription language. Advertisements are a system internal mechanism, which is not exposed to the application or to the user. An advertisement is described as $a=/t_1/t_2/.../t_{n-1}/t_n$, where t_i can be either an element name or a wildcard, and a has the same length as the publication it advertises.

In our approach, advertisements are derived from the DTD, since the DTD allows deriving all possible paths from the root to the leaves appearing in related XML documents.

We call an advertisement a *non-recursive* advertisement if it is extracted from a non-recursive DTD. The above advertisement *a* is an example of non-recursive advertisement. A DTD is recursive if it contains elements that are defined in terms of the elements themselves. The popular NITF DTD, often used for experimentation, is recursive. We call an advertisement a *recursive* advertisement if it is extracted from a recursive DTD. An advertisement may have multiple recursive parts that appear in sequence or are embedded in each other. We classify recursive advertisements into three categories as described below.

Simple-recursive advertisements: A simple-recursive advertisement has only one recursive pattern. The advertisement is described as $a=/t_1/t_2/.../t_{i-1}(/t_i/.../t_j)^+/.../t_n$, where the + operator declares that elements $/t_i/ \dots /t_j/$ must occur one or more times in the advertisement. Note that this is not part of XPath syntax. Advertisements are only used within the system, so the extended XPath syntax has no effect on clients and applications. In the proposed algorithms, we use $a=a_1(a_2)^+a_3$ to simplify the expression, where a_k ($1 \leq k \leq 3$) is a non-recursive advertisement.

Series-recursive advertisements: A series-recursive advertisement includes more than one recursive pattern in sequence. For example, an advertisement containing two recursive patterns in sequence can be described as $a=/t_1/t_2/.../t_{i-1}(/t_i/.../t_j)^+/t_{j+1}.../t_{l-1}(/t_l/.../t_o)^+/.../t_n$, or as simplified expression with non-recursive advertisements, denoted as $a=a_1(a_2)^+a_3(a_4)^+a_5$.

Embedded-recursive advertisements: An embedded-recursive advertisement recursively embeds patterns in other patterns. A possible case is $a=/t_1/t_2/.../t_{i-1}(/t_i/.../t_{l-1}(/t_l/.../t_o)^+/.../t_j)^+/.../t_n$, or $a=a_1(a_2(a_3)^+a_4)^+a_5$. The embedded-recursive advertisement can be more complex.

More types of recursive advertisements can be easily defined based on the above three types of advertisements. We discuss the matching algorithms for non-recursive and recursive advertisements in Sections 1.4.2 and 1.4.3, respectively.

1.4.2 Non-recursive Advertisement

In this section, we discuss the algorithms for subscription and non-recursive advertisement matching in the context of XML/XPath. An advertisement *a* matches a subscription *s* if the publication sets $P(a)$ and $P(s)$ intersect, that is, $P(a) \cap P(s) \neq \varnothing$. Figure 1.3(a) shows all possible relations between the two sets. To forward subscriptions, we need to identify the first two intersecting cases in Figure 1.3 (a). In this chapter, we focus on the subscriptions including parent-child operator (/), wildcard operator (*), and ancestor-descendant operator (//). For other operators appearing in the subscription, such as attribute filters, our approach can be easily extended to support them through value comparison. We discuss the matching algorithm for the following three subscription cases.

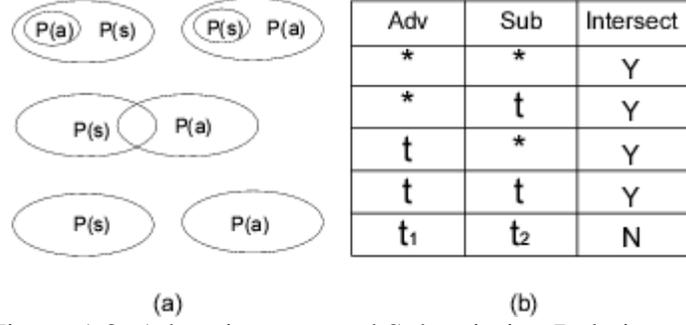| Adv	Sub	Intersect
*	*	Y
*	t	Y
t	*	Y
t	t	Y
t_1	t_2	N

(a) (b)

Figure 1.3: Advertisement and Subscription Relations

Absolute simple XPEs: A simple XPE only contains parent-child and wildcard operators. We describe the matching algorithm for absolute XPEs (without //-operator) and advertisements (`AbsExprAndAdv`) in Figure 1.4. For example, $s=/st_1/st_2/.../st_k$ and $a=/at_1/at_2/.../at_n$, where st_i is the i-th element of s and at_j is the j-th element of a. We use this notation in all algorithms in this chapter. First, the algorithm does not have to be applied, if the given XPE is longer than the advertisement. This observation is exploited because the advertisement has the same length as its publications, and thus, publications in $P(a)$ do not match all the elements in the longer XPE. Next, the algorithm compares each pair of elements or wildcards in the advertisement and the subscription, according to the matching rules shown in Figure 1.3(b). It returns 0, if some pair does not intersect; otherwise it returns 1. For example, given $a=/b/*/*/c/c/d$ and $s=/*/c/*/b/c$, the algorithm returns 0, since the matching rules fail to satisfy for i=4. As shown in Figure 1.3(b), the fifth row indicates that the advertisement includes an element c and the subscription includes an element b at the same position that do not intersect. That is publications matching the advertisement cannot match the subscription.

Input: advertisement a, and subscription s
Output: 1 if $P(a) \cap P(s) \neq \varnothing$, 0 if $P(a) \cap P(s) = \varnothing$
1: **If** $|s|>|a|$ **then return** 0
2: **For** i=1:$|s|$ **do**
3:      **If** matching rules are not satisfied for at_i and st_i
        then return 0
4: **Return** 1

Figure 1.4: AbsExpr. and Adv. Matching

Relative simple XPEs: These expressions are similar to absolute simple XPEs except for the first operator, which cannot be a "/". That is the XPE is relative. The matching could start at any position of the advertisement because the subscription is relative. A naive matching algorithm for this case is repeatedly calling `AbsExprAndAdv`. In iteration i, the algorithm takes the subscription as an absolute one and starts the matching from the i-th position of the advertisement. We skip the details of the naive algorithm as it is straightforward. The complexity of the naive algorithm is $O(n*k)$ where n is the length of the advertisement and k is the length of the subscription. We propose an optimized version of the matching algorithm for relative simple XPEs.

The matching algorithm for relative simple XPEs and advertisements (`RelExprAndAdv`) is a string matching problem (Knuth et al., 1977). We try to find the XPE, *s*, inside the advertisement, *a*, by starting at the first element of *a* that matches st_1 and continue (i.e., comparing to st_2 and so on) until we either complete the match or find a mismatch. In the latter case, we must go back to the place where we started. The difference between the traditional string matching problem and ours is that the wildcard "*" can match any element in our matching rules, as shown in Figure 1.3(b). To improve this algorithm, the KMP algorithm (Knuth et al., 1977) is applied to reduce the number of comparisons to $O(n)$. As shown in Figure 1.5, KMP computes a *next* table, recording all repeating patterns of *s*, to avoid backtracking. Figure 1.6 performs the matching by taking advantage of the *next* table.

Input: e subscription *s*
Output: *next*   table for *s* (an array of size |*s*|)
 1:   *next*(1) = -1, *next*(2) = 0
 2:   **For** *i*=3:|*s*| **do**
 3:     j = *next* (*i*−1)+1
 4:     **While** matching rules are not satisfied for st_{i-1} and st_j,
          and *j*>0 **do** j = *next* (*j*)+1
 5:     *next*(*i*) = j
 6:   **Return** *next*
        Figure 1.5: Next Table for Subscription

Input: advertisement *a*, and subscription *s*
Output: *matchPos* if $P(a) \cap P(s) \neq \varnothing$, 0 if $P(a) \cap P(s) = \varnothing$
 1:   **If** |*s*|>|*a*| **then return** 0
 2:   j = 1, I = 1, *matchPos* = 0
 3:   **While** *matchPos* = 0 and *i*≤|*a*| **do**
 4:     **If** matching rules are satisfied for at_i and st_j
          thenj++, i++
 5:     **Else** j = *next* (*j*)+1, **if** *j*=0 **then** j = 1,*i*++
 6:     **If** *j*=*k*+1 **then** *matchPos* = i - k
 7:   **Return** *matchPos*
        Figure 1.6: Optimized Matching

Input: advertisement *a*, and subscription $s=s_1//...//s_p$
        (s_1 is absolute or relative, and s_i, 2≤*i*≤*p*, is relative)
 Output: 1 if $P(a) \cap P(s) \neq \varnothing$, 0 if $P(a) \cap P(s) = \varnothing$
 01:   **For** *i*=1:*p* **do**
 02:             **If** $|s_1|+...+|s_i|>|a|$ **then return** 0
 03:   **If** *s* is absolute **then** *temp* = AbsAdv(a, /s_1)

04:   **Else** *temp* $= \text{RelAdv}(a, s_1)$

05:   **If** *temp* $= 0$ **then return** 0

06:   **Else** $j = k_1 + temp$, *temp* $= 0$

07:   **For** *i*=2:*p* **do**

08:       *temp* $= \text{RelAdv}(/t_j/.../t_n, s_i)$

09:       **If** *temp* $= 0$ **then return** 0

10:       **Else** $j = j + k_i - 1 + temp$

11:       **For** *l*=*i*+1:*p* **do**

12:           **If** $|s_{i+1}|+...+|s_l|>n-j+1$ **then return** 0

13:   **Return** 1

Figure 1.7: DesExpr. and Adv. Matching

Descendant operators in XPEs: Descendant operators indicate that more than one element should appear in the matching advertisement. The matching algorithm for XPEs with descendant operators and advertisements (`DesExprAndAdv`) is based on the above XPE matching algorithms (i.e., `AbsExprAndAdv` and `RelExprAndAdv`). We split the XPE in maximal length sub-XPEs that do not contain any descendant operators, and match each sub-XPE against the advertisement with sequence comparison. As shown in Figure 1.7, lines 1-2 guarantee that the advertisement is longer than the subscription. Next, we match s_1 against the advertisement according to the different types of *s* (lines 3-4), and, recompute the next available matching position in the advertisement (lines 5-6). k_i is the length of s_i. In the rest of this chapter, we use the same notation in other algorithms. The algorithm repeats this process until the end of the subscription is reached (lines 7-10), or returns 0 immediately if it finds the rest of the advertisement is shorter (lines 11- 12). For instance, given *a=/a/*/e/*/d/*/c/b* and *s=*/a//d/*/c//b*, the algorithm matches all sub-XPEs in *s* against *a* in order. It returns 1 because it finds each sub-XPE **/a*, *d/*/c* and *b* matches different parts in *a* (e.g., *a/**, **/d/** and *b*).

1.4.3   Recursive Advertisement

Input: advertisement $a=a_1(a_2)^+a_3$, and subscription *s*

       (a_k, $1 \leq k \leq 3$, is an advertisement)

Output: 1 if $P(a) \cap P(s) \neq \varnothing$, 0 if $P(a) \cap P(s) = \varnothing$

01:   **If** $|s| \leq |a_1 a_2|$ **then return** $\text{AbsAdv}(a_1 a_2, s)$

02:   **Else** *temp* $= \text{AbsAdv}(a_1 a_2, /st_1/.../st_{|a_1 a_2|})$

03:   **If** *temp* $= 0$ **then return** 0

04:   **If** $|s| \leq |a_1 a_2 a_3|$ **then** $q = 0$

05:   **Else** $q = Int((|s| - |a_1 a_2 a_3|)/|a_2|) + 1$

06:   $p = Int((|s| - |a_1 a_2|)/|a_2|) + 1$

07:   **For** *c*=*q*:*p* **do**

08:       *temp* $= \text{AbsAdv}(a_3, /st_{c*|a_2|+|a_1 a_2|+1}/.../st_{|s|})$

09:             **If** *temp* =1 **then return** 1

10:             **If** *c=p* **then** *temp* = AbsAdv(a_2, $/st_{c*|a_2|+|a_1 a_2|+1}/.../st_{|s|}$)

11:             **Else** *temp* = AbsAdv(a_2, $/st_{c*|a_2|+|a_1 a_2|+1}/.../$

$$st_{(c+1)*|a_2|+|a_1 a_2|})$$

12: **If** *temp* =0 **then return** 0

13: **Return** 1

<div align="center">Figure 1.8: AbsExpr. & Simple RecAdv.</div>

We focus on the matching of absolute XPEs and recursive advertisements. The matching of other types of XPEs and recursive advertisements can be implemented based on this algorithm. In Figure 1.8, the matching algorithm for absolute XPEs and simple recursive advertisements (`AbsExprAndSimRecAdv`) calls `AbsExprAndAdv` if the subscription is not longer than the recursive pattern (Line 1). If the subscription is longer, the algorithm estimates the maximum number that the recursive pattern would be repeated in the advertisement according to the length of both subscription and advertisement (Lines 4-6). Next, the algorithm tries all possible advertisements according to the maximum number of repeated recursive patterns (Lines 7-12). For example, given *a*=/*a*/*/*c*(/*e*/*d*)$^+$/*/*c*/*e* and *s*=/*/*a*/*c*/*/*d*/*e*/*d*/*, first, the algorithm compares /*a*/*/*c*/*e*/*d* in *a* with /*/*a*/*c*/*/*d* in *s*, and computes *q*=0 and *p*=1 in Lines 4-6. Second, it supposes that the recursive pattern is repeated only once, compares */*c*/*e* in *a* with *e*/*d*/* in *s* (Line 8) and fails to match. Next, it repeats the recursive part *e*/*d* twice, and continues the comparison (Line 11). Finally, it returns 1 (Line 9) if it finds *a* matches *s* with double recursive patterns in *a*. The complexity of the algorithm is $O(n^2)$, since it actually matches the subscription against each possible advertisement without recursive pattern. From a practical point of view, it is reasonable to limit the maximum nesting depth of items in a document, which would reduce the complexity of processing DTDs.

    The matching algorithm for absolute XPE and series-recursive advertisements (`AbsExprAndSerRecAdv`), where $a=a_1(a_2)^+a_3(a_4)^+a_5$, is implemented by calling the algorithm from Figure 1.8 recursively as shown in Figure 1.9. The matching determines how many times the first recursive pattern could be repeated, and calls the algorithm from Figure 1.8 repeatedly to try all possible advertisement formats. The matching of XPE and embedded recursive advertisements (`AbsExprAnd- EmbRecAdv`) is similar to `AbsExprAndSerRecAdv`. Figure 1.10 describes that, first, it determines how many times the outer recursive pattern could be repeated, and calls `AbsExprAndSerRecAdv` (not restricted to two recursive patterns) repeatedly.

1.5   Covering and Merging

In this section, first, we describe a novel data structure called *subscription tree* for maintaining subscriptions. The data structure captures the covering relations among subscriptions and speeds up the covering detection. Second, we present the covering algorithms for absolute simple XPEs, relative simple XPEs, and XPEs with descendant operators. Last, we explore the merging technique, and discuss the merging rules in the context of XPEs.

Input: advertisement $a=a_1(a_2)^+a_3(a_4)^+a_5$, and subscription *s*

Output: 1 if $P(a) \cap P(s) \neq \varnothing$, 0 if $P(a) \cap P(s) = \varnothing$

1:   **If** $|s| \leq |a_1 a_2|$ **then return** AbsAdv($a_1 a_2$, s)

2:   **Else** $p = Int((|s| - |a_1 a_2|)/|a_2|) + 1$

3:   $temp = 0$

4:   **For** $c = 0{:}p$ **do**

5:         **If** $temp = 1$ **then return** 1

6:         **Else** $temp = \text{SimRecAdv}(a_1(a_2)^{c+1} a_3(a_4)^+ a_5, s)$

7:   **Return** $temp$

<div align="center">Figure 1.9: AbsExpr. & Series RecAdv. Matching</div>

Input: advertisement $a = a_1(a_2(a_3)^+ a_4)^+ a_5$, and subscription s

Output: 1 if $P(a) \cap P(s) \neq \varnothing$, 0 if $P(a) \cap P(s) = \varnothing$

1:   **If** $|s| \leq |a_1 a_2 a_3\ a_4|$ **then return** SimRecAdv($a_1 a_2 (a_3)^+ a_4$, s)

2:   **Else** $p = Int((|s| - |a_1 a_2 a_3 a_4|)/|a_2 a_3 a_4|) + 1$

3:   $temp = 0$

4:   **For** $c = 0{:}p$ **do**

5:         **If** $temp = 1$ **then return** 1

6:         **Else** $temp = \text{SerRecAdv}(a_1(a_2(a_3)^+ a_4)^{c+1} a_5, s)$

7:   **Return** $temp$

<div align="center">Figure 1.10: AbsExpr.& Embedded RecAdv.</div>
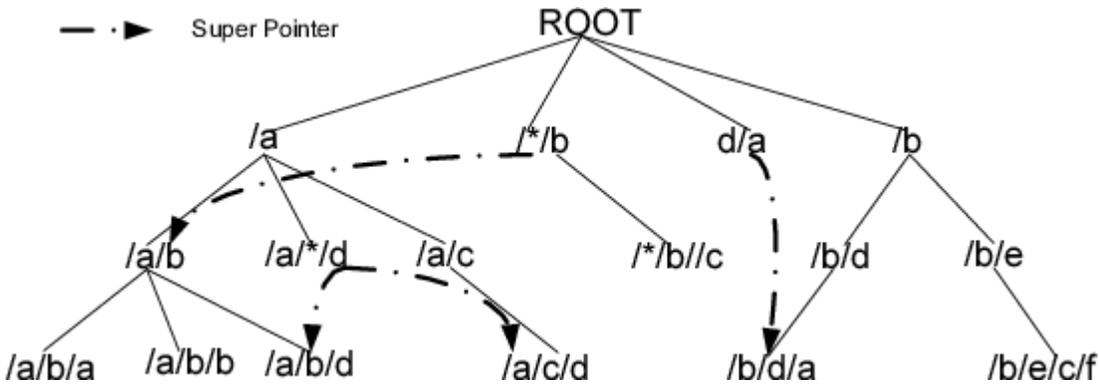
1.5.1   Subscription Tree

Figure 1.11: Subscription Tree

In covering-based routing, if an arriving subscription is covered by an existing subscription in the routing table, the new subscription is not forwarded to the next-hop broker. On the other hand, if the arriving subscription covers existing subscriptions, before it is forwarded, the broker needs to unsubscribe all the subscriptions that are covered by the new subscription. Therefore, the subscription traffic in the network is reduced by removing the redundant subscriptions and the routing table in the next-hop broker is compacted.

At each broker, subscriptions are maintained in a tree data structure. The idea is to store the subscriptions according to the covering relations among them. A subscription at a node in the tree covers all subscriptions in its subtree. Since a covering relation defines a partial order among subscriptions, a tree data structure cannot capture all the covering relations. A subscription node can have only one parent in the tree, but it may be covered by several subscriptions. We allow each node having a set of *super pointers*, which indicate the covering relations with nodes outside its subtree, as shown in Fig 1.11. *Super pointers* are shortcuts to subscriptions that the node covers. The tree and the super pointers form a directed acyclic graph (DAG) capturing the covering relations among subscriptions. With super pointers, a node covers its subtree, the nodes with subtrees pointed to by its super pointers, and the nodes with subtrees pointed to by its offsprings' super pointers.

The tree is maintained as follows. When a new subscription arrives, a breadth first search traverses the tree in order to find a place to insert the subscription. At a given node the following three cases are distinguished.

Case 1: If the new subscription has no covering relation with the node, the node's siblings are searched. If neither sibling has a covering relation with the new subscription, the subscription is inserted as new sibling, After insertion, super pointers maintained by the parent node are updated. If there is a super pointer of the parent pointing to a subscription that is also covered by the new subscription, then the super pointer is moved from the parent node to the new node.

Case 2: If the new subscription covers the current node, the new subscription is inserted between the current node and its parent. As a result, the new subscription becomes the parent of the current node, the old parent becomes the new subscription's parent. The old parent's super pointers are updated and moved to the new node, if there is a covering relation between the inserted subscription and a subscription pointed to by the super pointer.

Case 3: If the new subscription is covered by the current node, its children are searched until the new subscription is inserted. If the current node is a leaf node, the new node is inserted as the current node's child.

Existing super pointers are maintained while inserting. The new subscription may cause new covering relations and new super pointers to be added. Every time a new subscription arrives, we add new super pointers into existing nodes that cover the new subscription while searching the tree. However, this becomes expensive when the subscription tree grows larger. The reason we maintain the updated super pointers is for covering-based routing. When a subscription arrives, if it is not covered by existing subscriptions but it covers a set of subscriptions, we need to unsubscribe the subscriptions it covers and only forward the new subscription to neighbors. In this case, we need the super pointers to tell us what subscriptions should be unsubscribed. That means the updating of super pointers can be postponed to that point. The search space is reduced by super pointers. Note that we only need to unsubscribe subscriptions in the higher level of the tree since nodes in the subtrees are covered and unsubscribed already.

In the worst case, it takes $O(n)$ to identify the covering relations and insert a subscription to the subscription tree. For example, a subscription is covered by all subscriptions which are organized in one path. If the subscription tree created from a subscription workload is balanced, the best case run time is $O(log(n))$, which is the height of the tree.

Optimizations for the subscription searching and insertion can be performed based on the following two properties of the subscription tree.

Property of an Absolute XPE node: For all the absolute simple XPEs which have no wildcard and //-operators, the children's path length is always longer than their parent's path length. The parent is the prefix of its children.

Based on this property, we can perform depth-first search for an XPE to find a start node which has the same length as itself and start breadth-first search at that level. If an absolute XPE has a wildcard or //-operator in the middle of the expression, it is one or more levels higher than other simple XPEs of the same length in the subscription tree. Based on this property we can stop the search earlier.

Property of a Relative XPE node: A relative XPE is a child node of either the root node or another relative node. It will never be inserted in a subtree rooted by an absolute XPE. This property reduces the search space in the subscription tree.

1.5.2   Covering Algorithm

The key problem is how to determine the relationship between two given subscriptions. The covering relation between subscriptions is the containment problem in the context of XPEs. It has been proven that containment of simple path expressions can be tested in PTIME (Milo & Suciu, 1999). It is studied as a part of the problem that checking/finding a prefix replacement for a simple query is in PTIME. In this section, we detect covering relations of XPEs containing wildcard, /- and //-operator in PTIME, and present covering rules and algorithms for determining covering relations between single path XPEs. We say Sub_1 containing an element t_i covers Sub_2 containing an element m_i at the corresponding position, if t_i is a wildcard no matter what m_i is, or $t_i=m_i$, where none of t_i and m_i is a wildcard.

Absolute simple XPEs: The covering relation between two absolute XPEs (without //-operator) is the simplest case. We describe the covering algorithm for two absolute XPEs (AbsSimCov)(e.g., s_1 and s_2) in Figure 1.12. An important observation is that s_1 must be shorter than s_2 if s_1 covers s_2. This is exploited because a shorter XPE s has less constraints on items in an XML document, and refers to a bigger matching set $P(s)$. Next, the algorithm compares each pair of s_1t_i and s_2t_i in s_1 and s_2, respectively, according to the covering rules.

Input: two subscriptions s_1 and s_2
Output: 1 if s_1 covers s_2, 0 if s_1 does not cover s_2
 1:   **If** $|s_2|<|s_1|$ **then return** 0
 2:   **For** $i=1{:}|s_1|$ **do**
 3:      **If** covering rules are not satisfied for s_1t_i and s_2t_i
          then return 0
 4:   **Return** 1

Figure 1.12: Absolute Simple XPEs Covering

Relative simple XPEs: Figure 1.13 shows the covering algorithm for relative simple XPEs (RelSimCov), e.g., s_1 is relative, and s_2 is absolute or relative, calls AbsSimCov repeatedly to determine if s_1 contains subscription s_2 or not. An absolute XPE s_1 can not cover a relative XPE s_2, as the absolute XPE definitely refers to a smaller matching set $P(s_1)$ than $P(s_2)$.

It is important to note that the covering algorithm RelSimCov is also a string matching problem, as we pointed out in the RelExprAndAdv algorithm. The covering algorithm uses

covering rules that are different from subscription and advertisement matching rules used in `RelExprAndAdv`, however, a similar optimization can be applied to reduce the complexity of the covering algorithm from $O(k*n)$ to $O(k)$.

Input: two subscriptions s_1 and s_2 (s_1 is relative, s_2 is absolute or relative)

Output: *matchPos* if s_1 covers s_2, 0 if s_1 does not cover s_2

1: **If** $|s_2|<|s_1|$ **then return** 0
2: **For** $i=1:|s_2|-|s_1|+1$ **do**
3: **If** $AbsCov(/s_1,/s_2t_i/.../s_2t_{|s_2|}) = 1$

   then return *matchPos=i*
4: **Return** 0

Figure 1.13: Relative Simple XPEs Covering

Descendant operators in XPEs: Figure 1.14 describes the covering algorithm for XPEs with descendant operators (`DesCov`), where both s_1 and s_2 can be relative or absolute. It splits the XPE into sub-XPEs without //-operator, and matches each sub-XPE in s_1 against sub-XPEs in s_2 with sequence comparisons. First, it guarantees that s_2 is longer than s_1. Next, it matches the first sub-XPE in s_1 against s_2 according to different types of s_1 and s_2. The algorithm moves to the next sub-XPE in s_1 if it finds a match, and moves to the next sub-XPE in s_2 if it does not find a match. For example, given $s_1=/*/a//*/c$ and $s_2=/a/a/*//c/e/c/d$, first, the algorithm compares the sub-XPE $/*/a$ in s_1 with the sub-XPE $/a/a/*$ in s_2. Second, it moves to the next sub-XPE $*/c$ in s_1 and compares it with $*$ in s_2. Next, it compares $*/c$ in s_1 with the next sub-XPE $c/e/c/d$ in s_2, and finally, it returns *true* since the end of s_1 is reached and a match is found. Generally speaking, a sub-XPE in s_1 could not match a part of s_2 that includes a //-operator. For instance, given $s_1=/*/a//*/c$ and $s_2=/a/a/*//c/b/d$, the sub-XPE $*/c$ in s_1 does not cover $*//c$ in s_2 since $*/c$ refers to a smaller matching set. However, there is a special case that the sub-XPE s_{1i} in s_1 could cover a part of s_2 that includes a //-operator if s_{1i} ended with a wildcard and the matched part in s_2 ended with $//t$, where t can be either a wildcard or an element. For example, given $s_1=/a/*//*/d$ and $s_2=/a//b/c/d$, first, the algorithm compares $/a/*$ in s_1 with $/a//b$ in s_2, where *flag*=1 is used to record the current sub-XPE in s_1 matches a part of s_2 with //-operator. Second, it moves to the next sub-XPE $*/d$ in s_1 and compares it with c/d in s_2. Finally, it returns *true* since a match is found.

It is important to note that the covering detection between non-recursive advertisements is the same with the covering detection for subscriptions, since the non-recursive advertisement has the same format with an absolute simple subscription.

Input: two subscriptions $s_1=s_{11}//...//s_{1q}$, and $s_2=s_{21}//...//$

s_{2p}(both s_1 and s_2 can be relative or absolute subscription)

Output: 1 if s_1 covers s_2, 0 if s_1 does not cover s_2

01: **If** $|s_1|>|s_2|$ **then return** 0

02: **If** both s_1 and s_2 are absolute **then**
$temp = \text{AbsCov}(s_{11},\ s_{21}), flag = 0$

03: **If** $temp = 0$, and $|s_{11}|=|s_{21}|$, and s_{11} ends with $*$, and $p{\geq}2$
then $temp = \text{AbsCov}(s_{11}t_1/.../s_{11}t_{|s_{11}|-1}, s_{21}), flag = 1$

04: **If** $temp = 0$ **then return** 0

05: **Else if** s_1 is relative **then** $temp = \text{RelCov}(s_{11},\ s_{21}), flag = 0$

06: **If** $temp = 0$, and $|s_{21}|^3|s_{11}|$, and s_{11} ends with $*$, and $p{\geq}2$
then $temp = \text{RelCov}(s_{11}t_1/.../s_{11}t_{|s_{11}|-1},$
$s_{21}t_{|s_{21}|-|s_{11}|+2}/.../s_{21}t_{|s_{21}|}), flag \neg 1$

07: **Else return** 0

08: $i = 1, j = 1,\ temp_m = 0,\ temp_1 = 0$

09: **While** $i{\leq}q$ and $j{\leq}p$ **do**

10: **If** $temp = 0$ **then** $j = j +1,\ temp_m = 0$

11: **If** $j{\neq}p+1$ **then**
$temp = \text{RelCov}(s_{1i},\ s_{2j}t_{1+temp_1}/.../s_{2j}t|s_{2j}|), flag = 0$

12: **If** $temp = 0$, and $|s_{2j}|^3|s_{1i}|+temp_1$, and s_{1i} ends with $*$,
and $p{\geq}j$ **then** $temp = \text{RelCov}(s_{1i}t_1/.../s_{1i}t_{|s_{1i}|-1},$
$s_{2j}t_{|s_{2j}|-|s_{1i}|+2}/.../s_{2j}t_{|s_{2j}|}), flag = 1$

13: $temp_1 = 0$

14: **Else return** 0

15: **Else** $temp_m = temp + temp_m + |s_{1i}|-1, i{++}$

16: **If** $i{\neq}q+1$ and $flag{=}0$ **then**
$temp = \text{RelCov}(s_{1i},\ s_{2j}t_{1+temp_m}/.../s_{2j}t|s_{2j}|), flag = 0$

17: **If** $temp = 0$, and $|s_{2j}|^3|s_{1i}|+temp_m$, and s_{1i} ends with $*$,
and $p{\geq}j$ **then** $temp = \text{RelCov}(s_{1i}t_1/.../s_{1i}t_{|s_{1i}|-1},$
$s_{2j}t_{|s_{2j}|-|s_{1i}|+2}/.../s_{2j}t_{|s_{2j}|}), flag = 1$

18: **Else if** $i{\neq}q+1$ and $flag{=}1$ **then** $temp = 0,\ temp_1 = 1$

19: **Else return** 1

<div align="center">Figure 1.14: Descendant XPEs Covering</div>

1.5.3 Merging

Figure 1.15: Subscription Tree

If there is no covering relation among a set of subscriptions, subscriptions can be merged into a new subscription to create a more concise routing table. In this section, we exploit the merging rules for XPEs.

In the subscription tree, child nodes of the same parent have a better chance to be merged. As shown in Fig 1.15, node $/a/b/a$, $/a/b/b$ and $/a/b/d$ can be merged and they are represented by a new node $/a/b/*$ which is a union of the original XPEs. There was a super pointer pointing to node $/a/b/d$ before merging. This pointer should be removed because there is no covering relation between the pointer owner and the merger. If two nodes are merged, their subtrees become siblings of the merger. For example in Fig 1.15, after $/b/d$ and $/b/e$ are merged to $/b/*$, their children are the new node's children. The super pointer at $/b/d/a$ is not changed. To perform the merging in the subscription tree, we define several merging rules.

The subscriptions can be merged if they have only one difference (e.g., different elements). For instance, two subscriptions $s_1=a/*/c/d$ and $s_2=a/*/c/e$ can be merged into $s=a/*/c/*$. Note that if they differ in one operator, they should be in a covering relation to each other. The general form of this rule is:

- $s_1=o_1\ t_1...\ o_i\ t_i\ o_{i+1}\ \mathbf{m}\ o_{i+2}\ t_{i+1}...o_{n+1}\ t_n$
- $s_2=o_1\ t_1...\ o_i\ t_i\ o_{i+1}\ \mathbf{k}\ o_{i+2}\ t_{i+1}...o_{n+1}\ t_n$

are merged to

- $s=o_1\ t_1...\ o_i\ t_i\ o_{i+1}\ *\ o_{i+2}\ t_{i+1}...o_{n+1}\ t_n$

where m and k are different elements, o_i is either a /-operator or a //-operator, and t_i is a wildcard or an element. The number of merging candidates in this rule is not limited to 2.

Another rule is to merge subscriptions with two differences (e.g., different operators or different elements). For example, two subscriptions $s_1=/a/c/*/*$ and $s_2=/a//c/*/c$ that do not cover each other can be merged to $s=/a//c/*/*$. That is, different operators are merged to //-operator, and different elements are merged to *. We represent the general form of this rule as:

- $s_1=o_1\ t_1...\ o_i\ t_i\ o_{i+1}\ m...\ o_{j+1}\ t_j\ /\ t_{j+1}...o_{n+2}\ t_n$

- $s_2 = o_1\ t_1 ...\ o_i\ t_i\ o_{i+1}\ k...\ o_{j+1}\ t_j\ //\ t_{j+1}...o_{n+2}\ t_n$

are merged to

- $s = o_1\ t_1 ...\ o_i\ t_i\ o_{i+1}\ {}^*...\ o_{j+1}\ t_j\ //\ t_{j+1}...o_{n+2}\ t_n$

where m, k and t_i is a wildcard or an element, and o_i is a /-operator or a //-operator.

A more general rule is to replace the different parts in two subscriptions with the //-operator. We generalize this rule to:

- $s_1 = o_1\ t_1 ...\ o_i\ t_i\ XPE_1\ o_{i+1}\ t_{i+1}...o_n\ t_n$
- $s_2 = o_1\ t_1 ...\ o_i\ t_i\ XPE_2\ o_{i+1}\ t_{i+1}...o_n\ t_n$

are merged to

- $s = o_1\ t_1 ...\ o_i\ t_i\ //\ t_{i+1}...o_n\ t_n$

where XPE_1 and XPE_2 are different XPath expressions, t_i is a wildcard or an element, and o_i is a /-operator or a //-operator. This rule is applied if most parts in two subscriptions are equal, otherwise, more false positives will be introduced.

We periodically apply the above merging rules on the subscription tree to aggregate nodes that could be merged. We can compute an *imperfect merging degree* if each broker in the network knows the DTD relative to the XML data producer. An imperfect merger was first introduced in (Li et al., 2005). The imperfect degree of a new merger s, derived from s_1, s_2,..., s_n, is:

$$D_{imperfect} = \frac{|P(s) - \cup_{i=1}^{n} P(S_i)|}{|P(s)|}$$

It measures the imperfectness of an individual new merger. If the publications are distributed uniformly, the bigger the imperfect degree, the more false positive are introduced by the new merger. For example, two subscriptions $s_1 = /a/*/c/d$ and $s_2 = /a/*/c/e$ can be merged into $s = /a/*/c/*$. If the corresponding DTD indicates that the elements a,b,c,d,e are allowed at the fourth position, 60% false positive will be introduced at position 4. We need to consider the distribution of other elements in the subscription, e.g., the probability of each element appearing at other positions, to compute the total number of false positive introduced. Based on the DTD information, if $D_{imperfect}$ is 0, the merger is a perfect merger, and no false positives are introduced in this case. The false positives are not delivered to subscribers. They only occur in the network introduced due to imperfect merging. Clients are not exposed to false positives.

1.6   Evaluation

In this section, we experimentally evaluate the performance of our routing and covering algorithms. All algorithms are implemented in C++. We perform all experiments on a local cluster of 20 nodes and on PlanetLab. Each node in the cluster has an Intel Xeon 2.4GHz processor with 2GB RAM. For generating the XPE workload, we use the XPath generator released by Diao *et al.* (Diao et al., 2003). Queries are distinct, and we set the maximum length of an XPE to 10. We use the IBM XML Generator (Diaz & Lovell, 2003) to create the XML document workload. We use default parameters in this generator except that we set the maximum number of levels of the resulting XML documents to 10, which is consistent with the maximum length of XPEs. We use two different DTDs: the NITF (News Industry Text Format) DTD (NITF, 2005) and the PSD (Protein Sequence Database) DTD (PSD, 2005). The performance metrics we

measure include routing table size, XPE processing time and publication routing time in a single broker. We compare the network traffic (i.e., number of messages) and notification delay (i.e., the time between issuing a publication and receiving a notification) in two broker topologies with 7 brokers and 127 brokers, respectively. We also deployed our system on PlanetLab (PlanetLab, 2006) and measured the notification delay to validate the scalability of our approach.

Routing Table Size (RTS): Our algorithms exploit the covering relations among XPEs. We generate two data sets for NITF which include 100,000 XPEs each. We vary the probability of "*" occurring at a location step (*W*) and the probability of "//" occurring at a location step (*DO*) to generate two data sets *A* and *B* with different covering rates 90% and 50%, respectively. The covering rate indicates the similarity of subscriptions, that is, how many subscriptions are redundant and can be removed from the subscription set to create a compact routing table. For each data set, we evaluate the effect of the covering optimization on routing table size. The routing table size is the number of XPEs in the table. As shown in Figure 1.16, for Set *A*, the routing table size is reduced dramatically by covering. The subscriptions in Set *A* have a higher degree of overlap. The results suggest that the covering algorithm performs better on data sets with higher degree of overlap. That is, the covering technique achieves more benefit when subscribers have similar interests.

Merging can further reduce the routing table size by merging some XPEs according to our merging rules. When $D_{imperfect}$ is 0, the merger is a perfect merger. Figure 1.17 shows that applying perfect merging reduces the routing table size to 87%. When $D_{imperfect}$ increases, more XPEs can be merged. For instance, the routing table is compacted to 67% with $D_{imperfect}$ equal to 0.1.
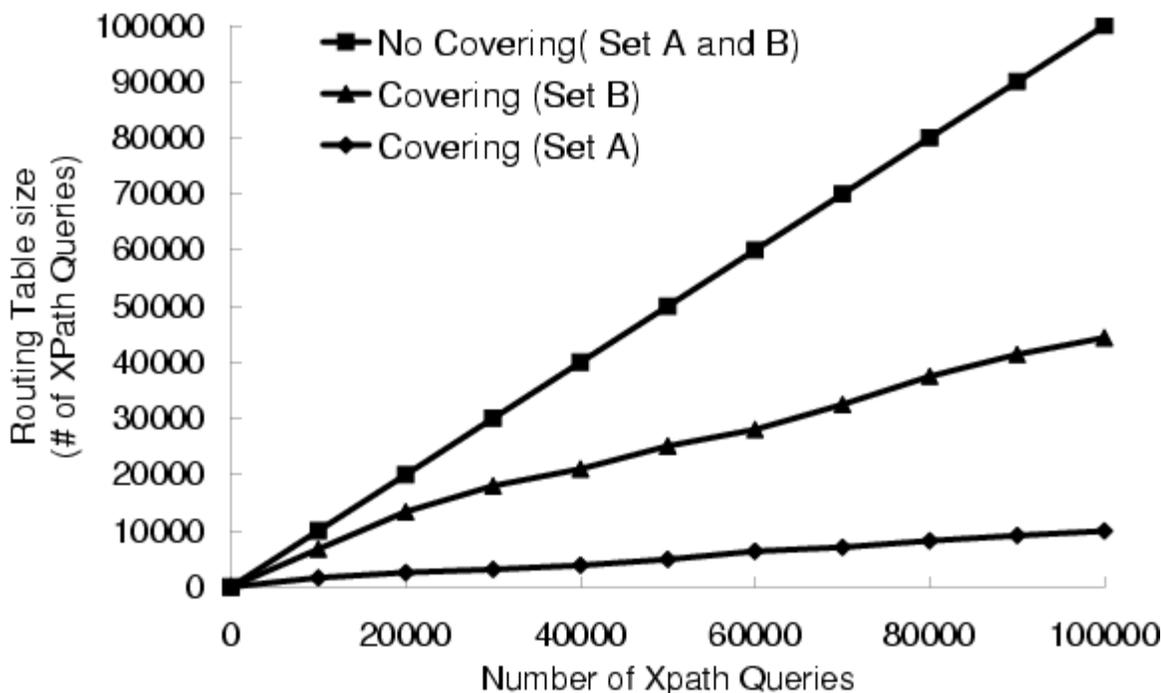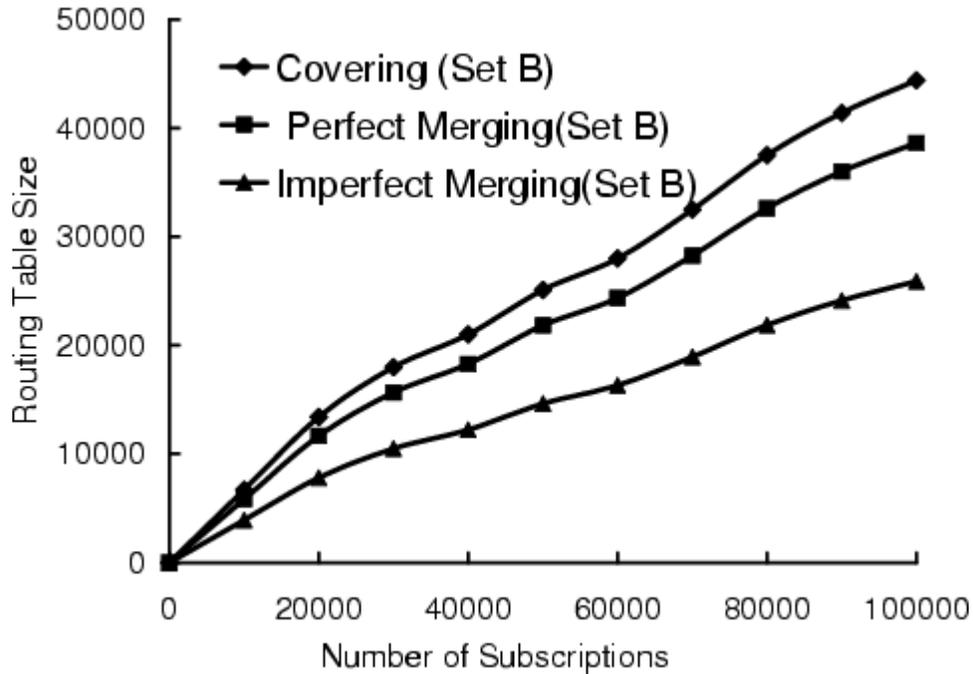

Figure 1.16: RTS

Figure 1.17: RTS

XPE Processing Time: We measure the XPE processing time of the covering algorithm. In covering-based routing, we first check the covering relationship when an XPE arrives at a broker. If the XPE is covered by existing XPEs, it will not be forwarded. Otherwise, we match the XPE against all advertisements and determine where to route it to. Without a covering algorithm, every XPE needs to be matched against all advertisements in order to be forwarded. We issue 5000 XPEs, and Figure 1.18 shows the processing time per XPE. Each data point in Figure 1.18 is the average processing time for 500 XPEs. Although detecting covering relations takes extra time, the experiment shows that the XPE processing time is less in covering-based routing, which avoids matching the covered XPEs against advertisements. For example, among the 5000 PSD XPEs, 90% of the XPEs are covered. The more XPEs are covered, the greater the improvement we achieve. Covering-based routing improves the XPE processing time of NITF XPEs by up to 49.2%, which is more than for the PSD XPEs. The reason is because the number of advertisements generated from the NITF DTD is 35 times larger than that of the PSD DTD. As a result, we benefit more from avoiding advertisement matching, especially when the broker has a large number of advertisements.

 Publication Routing Time: In this experiment, we evaluate the covering-based routing time of each message using data sets *A* and *B*. Note that the performance of non-covering-based routing in the original system has been evaluated against YFilter (Diao et al., 2003) in our previous work (Hou & Jacobsen, 2006). For some scenarios (i.e., the XPE workload with a high percentage of matched expressions, and with many wildcards and descendant operators), our system outperformed YFilter. For a contrasting workload with a very low matching percentage, YFilter outperformed us. We generate 500 XML documents and extract 23,098 publications from these documents. Table 1.1 shows the routing time of the publications against 100,000 XPEs. The measurements are obtained by averaging the time taken to route all publications. Both Set *A* and Set *B* exhibit benefits, derived from subscription covering. After applying the covering algorithm, the routing time for Set *A* and Set *B* are reduced by 84.6% and 47.5%, respectively. The merging

technique generates a more compact routing table, with which we can further improve the publication routing time.

Method	Set A (ms)	Set B (ms)
No Covering	13.96	14.23
Covering	2.15	7.47
Perfect Merging	1.87	6.88
Imperfect Merging	1.27	6.38

Table 1.1: Publication Routing Performance

Network Traffic: The network traffic can be influenced by the broker topology, the distribution of subscribers and publishers, and the routing strategy. In this experiment, we investigate the impact of advertisement-based routing and covering techniques on network traffic, given a tree-like broker topology. The broker overlay network is a tree in which each broker is connected to 2 subordinate brokers. We build two overlays for the experiment. One has three levels, which consists of 7 brokers. The other broker overlay has seven levels with 64 leaf brokers, and 127 brokers in total. Each leaf broker is connected with a subscriber. We extend the size of the broker network to show the scalability of our approach. Publishers randomly connect to the broker overlay.

Method	Network Traffic	Delay (ms)
no-Adv-no-Cov	58,138	29.02
no-Adv-with-Cov	50,931	7.50
with-Adv-no-Cov	39,849	28.9
with-Adv-with-Cov	38,492	7.45
with-Adv-with-CovPM	25,789	5.15
with-Adv-with-CovIPM	26,146	3.92

Table 1.2: 7 Broker Network

Method	Network Traffic	Delay (ms)
no-Adv-no-Cov	654,871	97.82
no-Adv-with-Cov	572,890	20.74
with-Adv-no-Cov	398,810	98.09
with-Adv-with-Cov	326,796	20.89
ith-Adv-with-CovPM	254,900	16.78
with-Adv-with-CovIPM	257,567	12.24

Table 1.3: 127 Broker Network

In this experiment, we compare routing strategies with different optimization techniques, including the routing with neither advertisement nor covering technique (no-Adv-no-Cov), the routing with covering only (no-Adv-with-Cov), the routing with advertisement only (with-Adv-no-Cov), the routing with both advertisement and covering techniques (with-Adv-with-Cov), the routing with advertisement, covering and perfect merging (with-Adv-with-CovPM), and the routing with advertisement, covering and imperfect merging (with-Adv-with-CovIPM). We generate 1,000 distinct XPEs for each subscriber using the PSD DTD, and 50 XML documents for the publishers. 4,182 publications are extracted from these documents. Table 1.2 and Table 1.3 show the total number of messages in the two broker overlays generated by one publisher. These messages, including advertisements, publications and subscriptions, are received by all brokers in the network under different routing strategies. As can be seen, the two advertisement-based routing methods significantly reduce the network traffic, because in this case a subscription is not flooded, and it is only forwarded to brokers that are on a path from the subscriber to potential publishers. The introduction of advertisements reduces the network traffic to 68.5% and 75.6% for non-covering-based and covering-based routing strategies, respectively. Moreover, applying both advertisement-based routing and covering-based routing techniques can reduce the overall network traffic to 66.2% and 49.9% in the two topologies. The experiment suggests that using advertisements to avoid subscription flooding and removing redundant queries by exploring covering and merging relations among subscriptions can reduce network traffic, save system resources and reduce the publication routing delay. Note that since imperfect merging may introduce false positives, the network traffic due to imperfect merging increases by 1.38% and by 1.04% in the two overlays, respectively. Overall, we achieve more benefit in a larger broker network. The scalability of the system is improved.
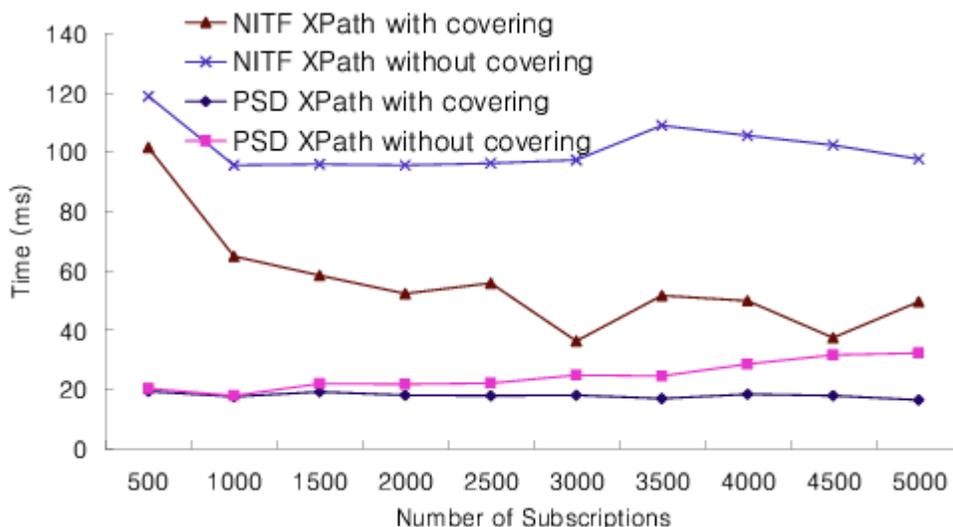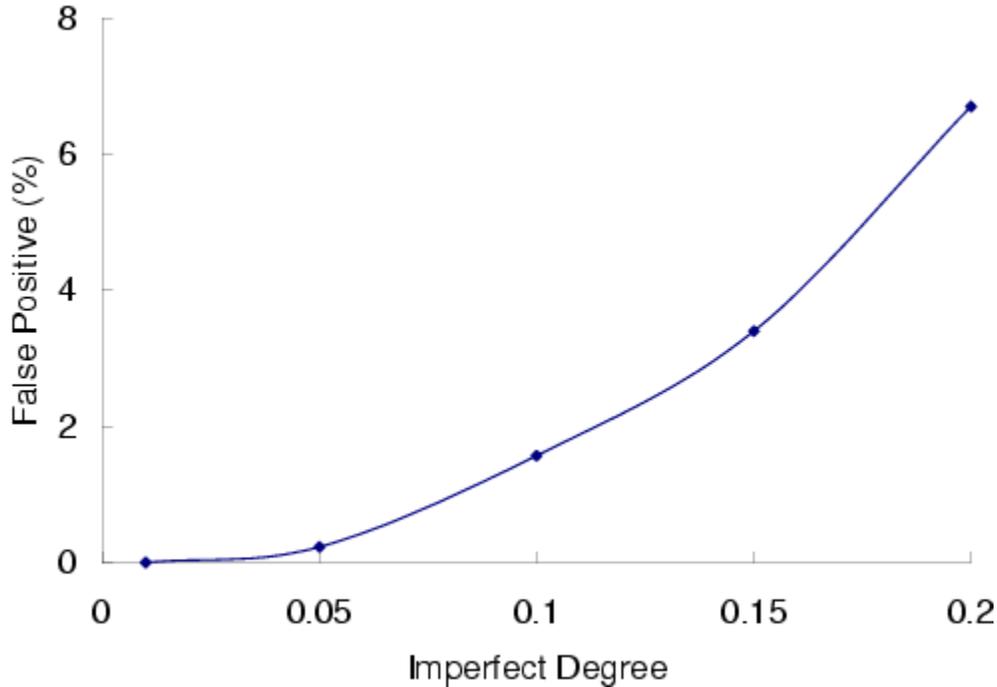

Figure 1.18: XPE Process Time

Figure 1.19: False Positives

False Positives from Imperfect Merging: If the system allows a larger error tolerance, more subscriptions can be merged. An imperfect merger may match more publications than expected. Therefore, a larger $D_{imperfect}$ means that the system allows more false positives. We show the relation between $D_{imperfect}$ and false positives in Figure 1.19. The larger $D_{imperfect}$ is, the greater the number of matched publications, among which some are false positives introduced by imperfect mergers. If the system tolerates up to 2% of false positives, a $D_{imperfect}$ with value less than 0.1 can satisfy the requirement. False positives only occur in the network and are not delivered to clients. Thus, they induce overhead but do not violate the subscription semantic expected by clients.

   Notification Delay on PlanetLab: In this experiment, we measure the notification delay on PlanetLab and demonstrate the scalability of our system. Due to the performance variation on PlanetLab nodes, which in our experiment is up to 15% per data point, we average the results from four experimental runs, as shown in Figure 1.20 and Figure 1.21. We setup a broker network with the maximum end-to-end distance equal to 7 hops. We measure the notification delay from publishers to subscribers for different number of broker hops and different XML document sizes. The experiment shows that the notification delay in covering-based routing is reduced by up to 74% compared with the routing without covering for both NITF and PSD documents. Moreover, the notification delay is linear in the number of hops. In covering-based routing, it increases less with the number of hops than in the content-based routing without covering. The reason is that the routing table size along the routing path has been reduced by the covering technique, as a result, the XML document matching time at each hop decreases, for instance, the routing table size is reduced to 6% for PSD XPEs. The result also indicates that the larger the document the longer the notification delay. A larger document saves more matching time with a condensed routing table. Therefore, the larger the document, the greater the improvement in routing delay we can achieve from the covering technique.
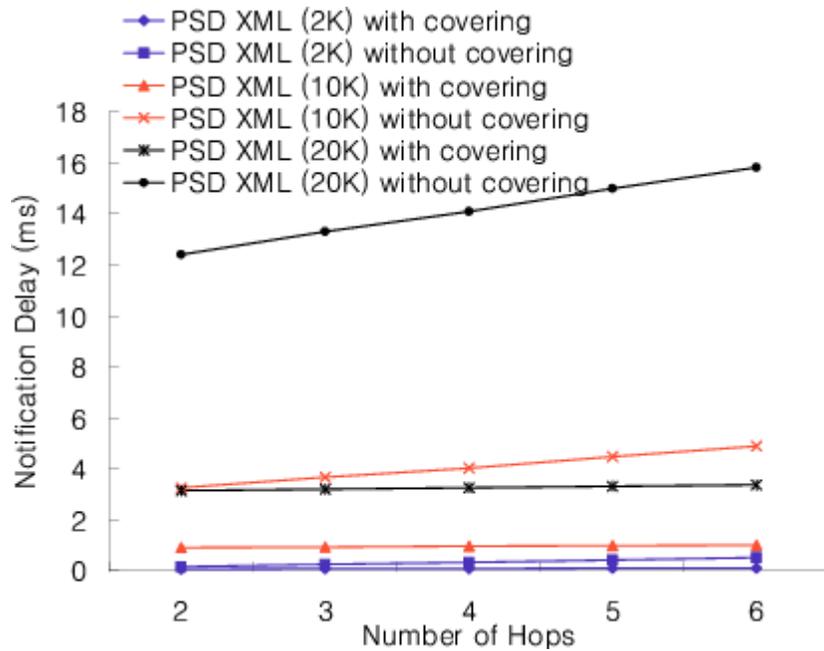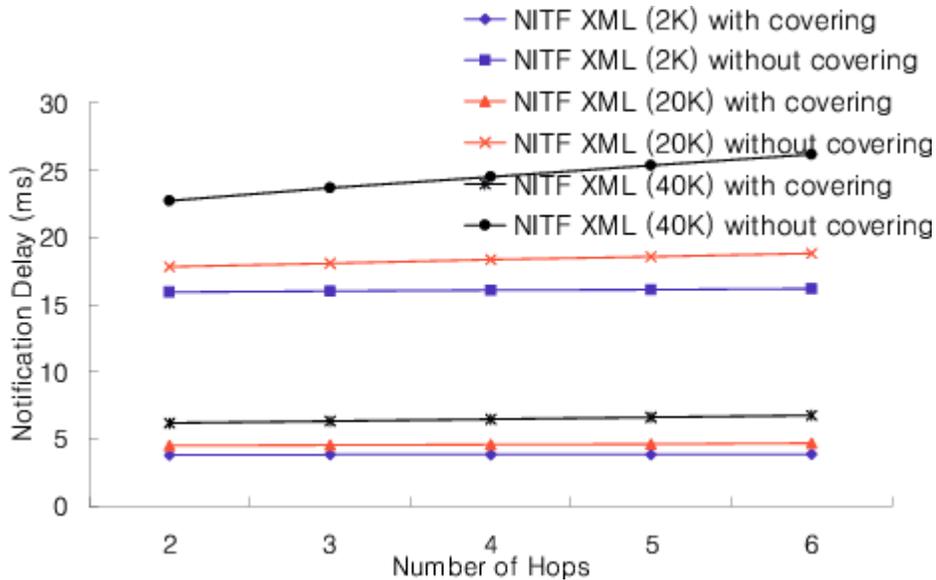
Figure 1.20: PSD XML

Figure 1.21: NITF XML

1.7   Conclusions

In this chapter, we studied the problem of efficiently routing XML data through a data dissemination network comprised of an overlay network of content-based routers. In the dissemination network, publishers' DTD files are transformed into advertisements expressed using XPath-like expressions. An advertisement creates a spanning tree rooted at the publisher. Subscribers specify XPath filters which are forwarded along the reverse paths of this tree for *intersecting* advertisements. XML documents from publishers are forwarded along these routing paths to subscribers with *matching* XPEs. By defining and exploiting covering and merging

relations for XPEs, a compact routing table results. Our techniques improve the routing time at each broker by up to 85% in the most favorable cases. Our experiments demonstrate that the scalability of the system is improved by applying advertisement-based routing, covering, and merging techniques for routing XML documents in a data dissemination network.

There are several directions for future work that will contribute to the more wide-spread adoption of XML dissemination networks. While the text-based nature of XML has clear advantages, it comes at a cost, which may counteract its continued proliferation in many application domains. For example, in many finance applications, processing speeds are critically important. This opens up several venues for future work. A binary XML representation standard may help reduce processing, if the matching and routing algorithms can adequately take advantage of this representation without repeated pre-processing. Also, a compressed representation combined with matching and routing that can operate on the compressed data could help reduce processing. In addition, hardware-software co-designs that allocate compute-intensive loops to hardware blocks can help speed up processing. The adaptation of existing algorithms and the development of new algorithms in this operation context constitute important avenues for future work. Questions that are simple to answer in software, such as how to support dynamic subscription expression updates, become much more challenging under the constraints imposed by hardware-based solutions. Moreover, the future development of applications that exploit the content-based publish/subscribe model described in this chapter will help to understand how to best use the model described in this chapter.

Acknowledgments

Bibliography

Altinel, M., & Franklin, M. J. (2000). Efficient filtering of XML documents for selective dissemination of information. In *Proceedings of the 26th International Conference on Very Large Data Bases (VLDB'00)*, (pp. 53–64), San Francisco, CA.

Amer-Yahia, S., Cho, S., Lakshmanan, L. V. S., & Srivastava, D. (2002). Tree pattern query minimization. *The VLDB Journal*, *11*(4), 315–331.

Bruno, N., Gravano, L., & Doudas, N. (2003). Navigation-vs. index-based XML multi-query processing. In *Proceedings of the 19nd International Conference on Data Engineering (ICDE'03)*, (pp. 139–150), Washington, DC.

Carzaniga, A., Rutherford, M. J., & Wolf, A. L. (2004). A routing scheme for content-based networking. In *Proceedings of IEEE Conference on Computer Communications (INFOCOM '04),* Hong Kong, China.

Chan, C.-Y., Fan, W., Felber, P., Garofalakis, M., & Rastogi, R. (2002a). Tree pattern aggregation for scalable XML data dissemination. In *Proceedings of the 28th international conference on Very Large Data Bases (VLDB'02)*, (pp. 826–837), Hong Kong, China.

Chan, C.-Y., Felber, P., Garofalakis, M., & Rastogi, R. (2002b). Efficient filtering of xml documents with xpath expressions. *The VLDB Journal*, *11*(4), 354–379.

Chand, R., & Felber, P. (2003). A scalable protocol for contentbased routing in overlay networks. In *Technical Report RR-03-074, Institut EURECOM, Feb.*

Chau, T., Muthusamy, V., Jacobsen, H.A., Litani, E., Chan, A., & Coulthard, P. (2008). Automating SLA modeling. In *Proceedings of the 2008 conference of the Centre for Advanced Studies on Collaborative Research*, Richmond Hill, Ontario, Canada.

Cheung, A. K., & Jacobsen, H.-A. (2006). Dynamic load balancing in distributed content-based publish/subscribe. In *Proceedings of the ACM/IFIP/USENIX 2006 International Conference on Middleware (Middleware'06)*, (pp. 141–161), New York.

Cheung, A., & Jacobsen, H.- A. (2008). *Efficient Load Distribution in Publish/Subscribe.* Technical report, Middleware Systems Research Group, University of Toronto.

Cugola, G., Nitto, E. D., & Fuggetta, A. (2001). The JEDI event-based infrastructure and its application to the development of the OPSS WFMS. *IEEE Transactions on Software Engineering*, *27*(9), 827–850.

Diao, Y., Altinel, M., & Franklin, M. J. (2003). Path sharing and predicate evaluation for high-performance XML filtering. *ACM Transactions on Database Systems*, *28*(4), 467–516.

Diao, Y., Rizvi, S., & Franklin, M. J. (2004). Towards an internet-scale XML dissemination service. In *Proceedings of the Thirtieth international conference on Very large data bases (VLDB'04)*, (pp. 612–623), Toronto, Canada.

Diaz, A. L., & Lovell, D. (2003). *XML generator*. Retrieved from `http://www.alphaworks.ibm.com/tech/xmlgenerator`

Dong, X., Halevy, A., & Tatarinov, I. (2003). *Containment of nested XML queries*. Tech. Rep. UW-CSE-03-12-05, University of Washington.

Fenner, W., Rabinovich, M., Ramakrishnan, K. K., Srivastava, D., & Zhang, Y. (2005). XTreeNet: Scalable overlay networks for XML content dissemination and querying (synopsis). In *Proceedings of the 10th International Workshop on Web Content Catching and Distribution*, French Riviera, France.

Fidler, E., Jacobsen, H.-A., Li, G., & Mankovski, S. (2005). The PADRES distributed publish/subscribe system. In *International Conference on Feature Interactions in Telecommunications and Software Systems (ICFI'05)*, (pp. 12–30), Leicester, UK.

Galanis, L., Wang, Y., Jeffery, S. R., & DeWitt, D. J. (2003). Locating data sources in large distributed systems. In *Proceedings of the 29th international conference on Very Large Data Bases (VLDB'03)*, (pp. 874–885), Berlin, Germany.

Green, T. J., Gupta, A., Miklau, G., Onizuka, M., & Suciu, D. (2004). Processing XML streams with deterministic automata and stream indexes. *ACM Transactions on Database Systems*, *29*(4), 752–788.

Hou, S., & Jacobsen, H.-A. (2006). Predicate-based filtering of XPath expressions. In *Proceedings of the 22nd International Conference on Data Engineering (ICDE'06)*, (p. 53), Washington, DC.

Hu, S., Muthusamy, V., Li, G. & Jacobsen, H.- A., (2008). Distributed Automatic Service Composition in Large-Scale Systems. In *Proceedings of the second international conference on Distributed event-based systems* (pp 233-244), ACM: New York.

Jacobsen, H.-A. (2006). *The PADRES content-based publish/subscribe system web site*. Retrieved `http://padres.msrg.toronto.edu/Padres/`

Knuth, D. E., Morris, J. H., & Pratt, V. R. (1977). Fast pattern matching in strings. *SIAM Journal on Computing*, *6*(2), 323–350.

Koloniari, G., & Pitoura, E. (2004). Content-based routing of path queries in peer-to-peer systems. In *Proceedings of the 7th International Conference on Extending Database Technology (EDBT'04)*, (pp. 29–47).

Koudas, N., Rabinovich, M., Srivastava, D., & Yu, T. (2004). Routing xml queries. In *Proceedings of the 20th International Conference on Data Engineering (ICDE'04)*, (p. 844), Washington, DC.

Kwon, J., Rao, P., Moon, B., & Lee, S. (2005). Fist: Scalable xml document filtering by sequencing twig patterns. In *Proceedings of the 31rd international conference on Very large data bases (VLDB'05)*, (pp. 217–228), Trondheim, Norway.

Li, G., Cheung, A., Hou, S., Hu, S., Muthusamy, V., Sherafat, R., Wun, A., Jacobsen, H.- A. & Manovski, S. (2007a). Historic data access in publish/subscribe. In *Proceedings of the 2007 inaugural international conference on Distributed event-based systems* (pp 80-84), Toronto, Ontario, Canada.

Li, G., Hou, S., & Jacobsen, H.-A. (2005). A unified approach to routing, covering and merging in publish/subscribe systems based on modified binary decision diagrams. In *Proceedings of the 25th International Conference on Distributed Computing Systems (ICDCS'05)*, (pp. 447–457), Columbus, OH.

Li, G., Hou, S., & Jacobsen, H.-A. (2008a). Routing of XML and XPath queries in data dissemination networks. In *Proceedings of The 28th International Conference on Distributed Computing Systems (ICDCS'08)*, (pp. 627–638), Washington, DC.

Li, G., & Jacobsen, H.-A. (2005). Composite subscriptions in content-based publish/subscribe systems. In *ACM/IFIP/USENIX, 6th International Middleware Conference (Middleware'05)*, (pp. 249–269), Grenoble, France.

Li, G., Muthusamy, V. & Jacobsen, H- A. (2007b). *NIÑOS: A Distributed Service Oriented Architecture for Business Process Execution*. The Journal of ACM Transactions on Computational Logics, 4(1), Article 2, Jan 2010.

Li, G., Muthusamy, V. & Jacobsen, H- A. (2008b). Adpative content-based routing in general overlay topologies. In *Proceedings of the 9th ACM/IFIP/USENIX International Middleware Conference* (pp 249-269). Berlin: Springer.

Li, G., Muthusamy, V., & Jacobsen, H- A. (2008c). *Subscribing to the past in content-based publish/subscribe*. Technical report, Middleware Systems Research Group, University of Toronto.

Liu, H., & Jacobsen, H.-A. (2002). A-ToPSS: A Publish/Subscribe System Supporting Approximate Matching. In *Proceedings of 28th International Conference on Very Large Data Bases* (pp 1107-1110), Hong Kong, China.

Liu, H., & Jacobsen, H.-A. (2004). Modeling uncertainties in publish/subscribe systems. In *Proceedings of the 20th International conference on Data Engineering* (pp 510-522), Boston.

Liu, H., Muthusamy, V., & Jacobsen, H.-A. (2009). *Predictive Publish/Subscribe Matching*. Technical report, Middleware Systems Research Group, University of Toronto.

Miklau, G., & Suciu, D. (2004). Containment and equivalence for a fragment of XPath. *Journal of the ACM*, *51*(1), 2–45.

Milo, T., & Suciu, D. (1999). Index structures for path expressions. In *Proceedings of the 7th International Conference on Database Theory (ICDT'99)*, (pp. 277–295), London, UK.

Moro, M. M., Bakalov, P., & Tsotras, V. J. (2007). Early profile pruning on xml-aware publish-subscribe systems. In *Proceedings of the 33rd international conference on Very large data bases (VLDB'07)*, (pp. 866–877), Vienna, Austria.

Muthusamy, V., & Jacobsen, H.- A. (2005). Small-scale Peer-to-peer Publish/Subscribe. In *Proceedings of the MobiQuitous Conference* (pp 109-119). New York: ACM.

Muthusamy, V., & Jacobsen, H.- A. (2007). *Infrastructure-less Content-Based Pub*. Technical report, Middleware Systems Research Group, University of Toronto.

Muthusamy, V., Jacobsen, H.- A., Coulthard, P., Chan, A., Waterhouse, J. & Litani, E. (2007). SLA-Driven Business Process Management in SOA. In *Proceedings of the 2007 conference of the center for advanced studies on Collaborative research* $(pp\ 264-267)$, Richmond Hill, Ontario, Canada.

Muthusamy, V., & Jacobsen, H.- A. (2008). SLA-driven distributed application development. In *Proceedings of the 3rd Workshop on Middleare for Service Oriented Computing* (pp. 31-36), Leuven, Belgium.

Mühl, G. (2001). Generic constraints for content-based publish/subscribe systems. In *Proceedings of the 6th International Conference on Cooperative Information Systems (CoopIS'01)*, (LNCS vol. 2172, pp. 211–225), Trento, Italy.

NITF (2005). NITF DTD. Retrieved from http://www.nitf.org/IPTC/NITF/3.3/documentation/nitf.html

Petrovic, M., Muthusamy, V., & Jacobsen, H.- A. (2005). Content-based routing in mobile ad hoc networks. In *Proceedings of the Second Annual International Conference on Mobile and Ubiquitous Systems: Networking and Services* (pp 45-55), San Diego, CA.

PlanetLab (2006). *PlanetLab*. Retrieved from https://www.planet-lab.org/

PSD (2005). *PSD DTD*. Retrieved from http://matra.sourceforge.net/dtdtree/bio/psdml_dtdtree.php

Sherafat Kazemzadeh, R. & Jacobsen, H.-A., (2007). δ*-Fault-Tolerant Publish/Subscribe systems*. Technical report, Middleware Systems Research Group, University of Toronto.

Sherafat Kazemzadeh, R. & Jacobsen, H.-A., (2008). *Highly Available Distributed Publish/Subscribe Systems*. Technical report, Middleware Systems Research Group, University of Toronto.

Snoeren, A. C., Conley, K., & Gifford, D. K. (2001). Mesh-based content routing using XML. *ACM SIGOPS Operating Systems Review*, *35*(5), 160–173.

Yan, W., Hu, S., Muthusamy, V., Jacobsen, H.-A. & Zha, L, (2009). Efficient event-based resource discovery. In *Proceedings of the 2009 inaugural international conference on Distributed event-based systems*, Nashville, TN.