

# IBM Research Report

## Clustering Algorithms for Content-Based Publication-Subscription Systems

**Zhen Liu, Joel L. Wolf, Li Zhang, Philip S. Yu, Anton Riabov**  
IBM Research Division  
Thomas J. Watson Research Center  
P.O. Box 218  
Yorktown Heights, NY 10598



Research Division  
Almaden - Austin - Beijing - Delhi - Haifa - India - T. J. Watson - Tokyo - Zurich

# Clustering Algorithms for Content-Based Publication-Subscription Systems

Anton Riabov\*  
IEOR Department  
Columbia University  
New York, NY 10027, USA  
E-mail: anton@ieor.columbia.edu

Zhen Liu, Joel L. Wolf, Philip S. Yu and Li Zhang  
IBM T.J. Watson Research Center  
P.O. Box 704  
Yorktown Heights, NY 10598, USA  
E-mail: {zhenl,jlwolf,psyu,zhangli}@us.ibm.com

## Abstract

We consider efficient communication schemes based on both network-supported and application-level multicast techniques for content-based publication-subscription systems. We show that the communication costs depend heavily on the network configurations, distribution of publications and subscriptions. We propose to use clustering algorithms to form multicast groups. We devise new algorithms and adapt partitional data clustering algorithms for these content-based publication-subscription systems. These algorithms can be used to determine multicast groups with as much commonality as possible, based on the totality of subscribers' interests. These algorithms perform well in the context of highly heterogeneous subscriptions, and they also scale well. An efficiency of 60% to 80% with respect to the ideal solution can be achieved with a small number of multicast groups (less than 100 groups in our experiments). Some of these same concepts can be applied to match publications to subscribers in real-time, and also to determine dynamically whether to unicast, multicast or broadcast information about the events over the network to the matched subscribers. We demonstrate the quality of our algorithms via a set of simulation experiments.

## 1 Introduction

Publication-Subscription systems (*pub-sub* for short) provide information on specific real-time events from publishers to interested subscribers. The subscribers express their interest in the form of multiple subscriptions. The publishers and subscribers may be located at arbitrary nodes in a distributed network.

Pub-sub systems can be characterized into two broad types based on the degree of generality, usability and personalization allowed to the subscribers. We focus in this

paper on the more sophisticated of these, namely *content-based* pub-sub. See, for example, the pioneering work of the Gryphon project [2, 4, 14], as well as NEONet [21] and READY [8]. Such systems differ from the more mature *subject-based* pub-sub systems in their generality, usability and flexibility. Borrowing extensively from the classic stock market example used by Gryphon, a subject-based pub-sub system would allow subscriptions based on broad criteria only. A subscriber might request all events related to IBM, for instance. Such a system is powerful but relatively inflexible. The subscriber might receive many publications involving IBM stock market events of little interest. On the other hand, a content-based pub-sub system would allow subscriptions which were based on the conjunction of potentially multiple predicates related to different attributes.

Expanding on the stock market scenario originally discussed by Gryphon, a content-based pub-sub system would allow subscriptions which are based on the conjunction of potentially multiple predicates concerning different attributes. The motivating Gryphon subscription example was based on three distinct attributes:

- The first was an equality predicate based on a character string referring to the stock name. *name=IBM* would be an example, though one could also imagine categories ("blue chip", for instance) being specified instead.
- The second could be one or two inequality predicates based on a two decimal numerical attribute referring to the stock price, such as  $90.00 < price \leq 110.00$ . Alternatively one could normalize these by the price of the stock at the start of the day, so that a subscription might look for changes in price within 10 percent in a day.
- The third could be one or two inequality predicates based on an integer attribute referring to the volume: *volume > 10,000* might be an example. One could easily imagine translating these volumes into dollars to ensure a common basis amongst differing stocks, so that one might look for volume whose monetary equivalent exceeds one million dollars.

\*The work was carried out while this author was with IBM T.J. Watson Research Center

A client with this (expanded) subscription would receive information about all trades of “blue chip” stocks whose price stays within a 10 percent range during the day and involves more than a million dollars in start-of-day units.

So content-based pub-sub is more general and much more personalizable than subject-based pub-sub. In general, we can assume that content-based pub-sub systems allow each predicate to be *range-based*, composed of *intervals* in the underlying domain of the predicate. (Because computers can handle only inherently finite and discrete attribute values, one can assume without loss of generality that all possible intervals are actually open on the left and closed on the right. This assumption allows the intervals to ‘fit together’ more cleanly.) By decomposing a subscription with multiple such ranges into multiple subscriptions consisting of single ranges we can see that it is sufficient only to consider intervals, albeit at a cost of more subscriptions. And even attributes such as name, not typically thought of as numerical, can be indexed and therefore linearized in some fashion. A predicate involving “blue chip”’s can thus be decomposed into the union of several conjunctions.

This perspective allows us to think of a subscription as a half-open, half-closed aligned rectangle in space, each dimension corresponding to a different attribute. (The term *aligned* means that the rectangle is parallel to the various axes.) A published event becomes a point in the underlying space.

Clearly, the extra function involved with pub-sub systems comes at a price. It is technically challenging for a content-based pub-sub system to efficiently distribute the many publication events to the interested subscribers over the network. And it is technically challenging to do so in a manner which scales with the dimensionality of the underlying event space, the number of publishers and the number of subscriptions.

There are two key dynamic problems that a content-based pub-sub system needs to solve:

1. One must *match* a given real-time event quickly to determine the set of relevant subscribers. This is the so-called *matching* problem.
2. One must decide to unicast, multicast and/or broadcast information about the event over the network efficiently to the matched subscribers (or possibly to a superset of those subscribers, if that is more algorithmically reasonable, to be filtered out as necessary). We shall call this the *distribution method* problem.

Both of these problems must be solved in real time as new events are published. However, there is a related, static “pre-processing” problem that should be solved in order to enable the real-time algorithms to function efficiently. Basically we must precompute a set of high quality multicast groups having as much commonality as possible, based on the totality of subscribers’ interests. We shall call this the *subscription*

*clustering* problem. This paper will be focused on the clustering problem, whereas the matching problem will be analyzed in more detail in a companion paper [16].

The Gryphon papers do tackle these problems in elegant ways. In [2] and [4], for instance, the authors consider the matching problem. The interested reader is referred to [7] for a detailed description of the matching algorithm, a performance analysis and a comparison with existing matching algorithms. In [4] and [14] the authors consider multicasting techniques. It seems fair to say, however, that the authors base their algorithms primarily on subscriptions in which each dimension is based on either equality or wild-card predicates. (A *wild-card* (\*) indicates that the subscriber does not care about the value in that dimension.) While their algorithms will certainly work in full generality, we believe that they are optimized for their motivating predicate types.

However, the above mentioned earlier work of Gryphon project (see, for example, [14]) concluded that multicast mechanism does not provide substantial benefits in many pub-sub systems, and that given the multicast overhead, unicast and broadcast are sufficient for these systems. We believe that the conclusions would be very different depending on the network configurations, distribution of publications and subscriptions. In this paper we consider larger networks with fewer number of subscriptions from each node. We think this setting could be closer to the real world environments. This leads to the potentially large advantage of forming multicast groups. We will evaluate the benefits of employing (both network-supported and application-level) multicast mechanisms. We shall show that the communication costs depends crucially on how to form the multicast groups.

Specifically, in this paper:

- We examine the various assumptions in the Gryphon framework and investigate quantitatively different impacts from several aspects of the pub-sub framework.
- We introduce a general framework that allows to adapt partitional data clustering algorithms for pub-sub systems in which subscriber preferences are described more generally than in previous work.
- We devise a number of new clustering algorithms and enhanced some others. Among the algorithms we now consider are *K-means*, an important variant called *Forgy K-means*, a hierarchical clustering algorithm based on *Minimum Spanning Trees (MST)*, the *Pairwise Grouping* algorithm and a variant called *Approximate Pairwise Grouping*, and finally a so-called *No-Loss* algorithm. (The no-loss strategy implies that a publication never needs to be filtered out in the dynamic stage: Every subscriber that receives a publication is indeed interested in that message. The other algorithms described in this paper do *not* have this property.) Our comparisons show new leaders among these algorithms, and our results are more robust and realistic.

- We evaluate our algorithms on a large network model. Our subscriptions are based on three different models of interest, and the same is true for our publication model. We analyze the effects of regional subscriber preferences relative to the network topology. We show that the conclusions in this paper depend dramatically on assumptions about the size and structure of the network.

We will consider two flavors of multicasting in this paper, network supported and application level multicast schemes. The interested reader is referred to [3] for a description of the various tradeoffs, see also [15] for a description of application level multicasting.

The remainder of this paper is organized as follows. We introduce some notation in Section 2. In Section 3, we describe some preliminary investigations which illustrate potential impacts of communication algorithms on the communication costs. We then in Section 4 describe the algorithms for solving common interest clustering problems. In Section 5 we present the results of our experiments and comment on the relative performance of the algorithms. Finally, in Section 6 we summarize the results and discuss about future work. We believe that the many ideas for future work are indicative of the newness and importance of this area of research.

## 2 Problem Notation

In this section we define some key parameters we will use in our problem descriptions.

- Let  $\Omega$  denote the publication event space. Each event being published within the system can be uniquely described with a single value  $\omega$  such that  $\omega \in \Omega$ .
- Let  $N$  denote the number of dimensions (or attributes) in  $\Omega$ , so that  $\Omega \subseteq \mathbb{R}^N$ .
- Let  $p_p(x)$  be the probability that publications events are within set  $x \subseteq \Omega$ .
- Define the underlying network topology as an undirected graph  $G = (V, E)$ .
- Define the communication costs to be  $c_e \geq 0$  for each edge  $e \in E$ .
- Let  $V_P \subseteq V$  be the set of nodes containing publishers.
- Let  $V_S \subseteq V$  be the set of nodes containing subscribers.
- Let  $N_S$  be the number of subscribers.
- In this paper we shall assume that each subscriber  $v_i \in V_S, i = 1, \dots, N_S$ , has a set of subscription preference expressed by (possibly infinite) rectangles  $\mathcal{I}_i = \{b_{ij}\}_{j=1}^{r_i}$  associated with it. Each  $b_{ij} \subseteq \Omega$  is an aligned rectangle in space  $\Omega$ .

- We define  $\mathcal{I} := \bigcup_{v \in V_S} \mathcal{I}_v$  to be the set including all subscription rectangles.
- We define  $k := |\mathcal{I}|$  to be the number of subscriptions.

The size of the clustering problem is essentially determined by the dimension of the event space  $N$  and the number of subscriptions  $k$ . We are interested in algorithms that scale well with respect to these values.

Typically the number  $K$  of multicast groups available to the clustering algorithm is known in advance. In the case of network-supported multicast, this is the number of multicast IP addresses purchased. In the case of application level multicast this number is limited by the amount of memory and processing power of participating computers.

## 3 Preliminary Analyses

Earlier work from the Gryphon project (see, for example, [14]) has demonstrated that multicast mechanism does not provide substantial benefits in many pub-sub systems. Given the well-known structural and performance overhead for applying multicast, unicast and broadcast are sufficient for these systems, having little or no overhead and insignificant increase in the delivery cost. We examine the various assumptions in the Gryphon framework and investigate quantitatively different impacts from several aspects of the pub-sub framework. We believe that the conclusions would be very different depending on the network configurations, distribution of publications and subscriptions.

In our model, events are generated in 4 dimensions, where first dimension corresponds to “regional attribute”. When a publication event occurs, the publication is always set to the identification number of originating subnet (“stub”) for this message. The *degree of regionalism* parameter is the probability that in a subscription this attribute will be required to be equal to corresponding subnet number. Zero degree of regionalism corresponds to no regionalism, and degree 1 to absolute regionalism. Regional subscriptions in this table are generated with this parameter set to 0.4. Non-regional subscriptions have this parameter set to 0 during generation of subscriptions.

The other 3 attributes of events can take on integer values between 0 and 20, according to either uniform or Gaussian distribution. Preference on each parameter can be either a “don’t care” parameter, denoted “\*”, which means that all values of this parameter are of interest, or the interval can be specified. In the uniform case, the probabilities of non having “\*” in position 2 is 0.98, and then decreases with the rate of 0.78 for subsequent parameters, so parameter 3 preference is specified with probability  $0.98 \cdot 0.78$ , and parameter 4 with probability  $0.98 \cdot 0.78^2$ . If parameter preference is specified, then two random numbers between 0 and 20 are generated, sorted if needed, and assigned to the ends of the preference interval. For Gaussian distribution each of the 3 parameters

preference can have value of “\*” with probability  $q_1$ , can be a left-ended interval with probability  $q_2$ , a right-ended interval with probability  $q_3$  and an interval with both ends with probability  $(1 - q_1 - q_2 - q_3)$ . If both ends of the interval are specified, the center of the interval follows a Gaussian distribution with parameters  $(\mu_3, \sigma_3)$  and the length of the interval follows a Pareto-like distribution with a given mean. If the interval is one-ended, then the end of the interval is chosen from Gaussian distribution with parameters  $(\mu_1, \sigma_1)$  and  $(\mu_2, \sigma_2)$  for left-ended and right-ended intervals. The parameters in the experiment were chosen from the following table, to simulate stock name, price and trading volume:

Para	$q_1$	$q_2$	$q_3$	$\mu_1, \sigma_1$	$\mu_2, \sigma_2$	$\mu_3, \sigma_3$	mean
2	0.1	0	0	8,2	10,2	9,6	1
2	0.15	0.1	0.1	8,1	10,1	9,2	4
2	0.35	0.1	0.1	8,1	10,1	9,2	4

The networks were generated by GT-ITM package (see [20] for more details) using transit-stub model with one transit block and the following parameters:

Node	Trans node	Stubs/trans node	Nodes in a stub
100	4	3	8
300	5	3	20
600	4	3	50

Figure 1 illustrates the topology of the network with 100 nodes. More details about the simulation framework can be found in section 5.1.




Figure 1. Network topology with 100 nodes.

Tables 1 and 2 show the communication costs for broadcast, unicast distribution and ideal multicast, where ideal multicast stands for the distribution scheme where a multicast group is formed for each publication event and is composed only of the subscribers interested in this event. The ideal multicast could thus use all the possible  $2^{N_s}$  groups.

Table 1. Potential and Effective Advantage of Multicast for the case of degree 0.4 regionalism

Node	Sub'n	Dist'n	Unicast	Broadcast	Ideal
100	10000	uniform	63758	1430	1397
100	10000	gaussian	97441	1430	1429
100	5000	uniform	31351	1430	1334
100	5000	gaussian	48805	1430	1415
100	2000	uniform	12415	1430	1145
100	1000	uniform	5846	1430	867
100	1000	gaussian	9513	1430	1134
100	80	uniform	750	1430	310
100	80	gaussian	548	1430	287
300	5000	uniform	38612	5079	3453
300	1000	uniform	8181	5079	867
300	1000	gaussian	12692	5079	2205
300	350	uniform	3638	3880	1065
600	10000	uniform	92178	10235	6720
600	10000	gaussian	139020	10235	8212
600	5000	uniform	45320	10235	4820
600	5000	gaussian	69179	10235	6431
600	1000	uniform	5477	10235	1350
600	1000	gaussian	9408	10235	1823

We can see that for the same network, the difference in cost between the broadcast and ideal multicast is small for cases with a large number of subscriptions, and becomes larger as the number of subscriptions decreases. This gap can be as large as 4 times the ideal solution. Therefore, there is need to further investigate more sophisticated content delivery mechanisms.

The Gryphon framework has a 100 node network, with an average of 125 subscriptions for each of the 80 nodes. For such types of networks with a small number of nodes, each node having many subscriptions per publication, there is a very high probability that at least one of the subscriptions at this node includes this publication. Therefore, the number of nodes interested in this publication is either very high or very low. This means that the system needs to deliver the publication either to almost every node or to a very small set of nodes. We believe that this fact results in the conclusion by Gryphon that broadcast is sufficient for such systems. For larger networks with relatively fewer subscriptions, on the other hand, it is unlikely that publications need to be delivered to almost every node in the network. Multicast is most beneficial in this kind of setting which messages need to be

**Table 2. Potential and Effective Advantage of Multicast for the case of no regionalism**

Node	Sub'n	Dist'n	Unicast	Broadcast	Ideal
100	5000	uniform	50737	1430	1377
100	5000	gaussian	81779	1430	1428
100	1000	uniform	9409	1430	1039
100	1000	gaussian	16314	1430	1301
100	80	uniform	816	1430	328
100	80	gaussian	1580	1430	545
300	5000	uniform	61513	5079	4019
300	5000	gaussian	98735	5079	4751
300	1000	uniform	13384	5079	2026
300	1000	gaussian	21167	5079	2918
300	80	uniform	61513	5079	1113
300	80	gaussian	6113	5079	1598
600	10000	uniform	151270	10235	7993
600	10000	gaussian	232405	10235	9382
600	5000	uniform	73830	10235	6184
600	5000	gaussian	116952	10235	8000
600	1000	uniform	8276	10235	1791
600	1000	gaussian	14428	10235	2502

delivered to only part of the network.

In Tables 1 and 2 the unicast and ideal multicasts are in general larger for the gaussian distributions than for the uniform. This is due to the fact that more nodes are likely to be interested in the publication events for the gaussian case, hence results in more communication cost. This shows that the publication distributions also affects the multicast benefits.

Furthermore, Tables 1 and 2 show that the communication costs for regional-specific subscriptions are smaller than the costs for non-regional subscriptions. More generally, the dependence of the subscriptions by different nodes can have a big impact on the multicast benefits. Consider independent subscriptions by the nodes for a particular publication. The nodes that are interested in this publication would be scattered around the network evenly with very high probability. The multicast benefit for delivering messages to such a scattered network would not be significant. On the other hand, if the subscriptions have a regional concentration, the interested nodes of a publication would very likely be more localized. It would not be surprising to observe substantial benefits from employing multicasts.

In addition, if the probability that each node subscribes to a given message is independent of other nodes, and there is no concentration of common interest in the event space, then it is difficult to form only a few number of groups and greatly improve communication efficiency – each of  $(2^{N_s} - 1)$  possible multicast groups will be needed with equal probability.

In this paper we consider larger networks with fewer number of subscriptions from each node. We think this setting

could be closer to the real world environments. This leads to the potentially large advantage of forming multicast groups. For such pub-sub systems, using broadcast to deliver messages would not be appropriate due to its large communication overhead. We will evaluate the benefits of employing multicast mechanisms. The communication benefits of multicast depends crucially on how to form the multicast groups. Due to the overhead of managing a large number of multicast groups, we would need to consider forming a limited number of groups. We develop several algorithms for constructing multicast groups and evaluate their performance benefits. Algorithmic complexity is also a key factor for these real-time applications. We further study the cost benefit and running time trade-offs of these algorithm and discover good algorithms for practical applications.

Before going into details on the algorithms, it is important to describe the assumptions of our studies for the aforementioned reasons. First, we assume that the peaks in density of subscriptions follow peaks in density of the messages. It seems likely that multicast will not provide comparable improvements in communication cost without this assumption.

We further assume that the subscriber preferences are regional in the network topology. In our experiments, for example, stock name preference (mean of the distribution) was assigned according to the transit block in the network. In addition we assume subscriptions themselves are unevenly distributed in the network, with higher concentrations of interest in some areas, and lower in others.

Under these assumptions forming a limited number of groups using subscription clustering algorithms can potentially lead to large reduction of communication costs when compared to unicast and broadcast. While relatively restrictive, the assumptions still seem to be practical. Indeed, in many real-life pub-sub systems we would expect that events, in which more people are interested, are typically published more often, than the less interesting events. Also, we can expect the regionalism of subscriptions, with more concentration of users in certain parts of the network, and regionalism of interest, with interest distribution being different in different parts of the network. The model formally presented in section 5.1 follows the above assumptions.

## 4 Algorithms for Subscription Clustering

There are two distinct categories of subscription clustering algorithms that we present in the paper. First category, the *Grid-Based* clustering algorithms, extends earlier work on subscription clustering in content-based pub-sub systems for the case of rectangular preference sets. Work by the Gryphon group [4] and by Katz *et al* [19] employ similar data algorithms for clustering of point subscriptions. In the following subsection we describe the framework that allow us to use data clustering algorithms for clustering subscriptions. In the following subsections we illustrate our approach by

describing how four clustering algorithms: *K-means*, *Forgy K-means*, *MST*, and *Pairwise Grouping* can be used for subscription clustering.

The second category of subscription clustering algorithms presented in this section includes just one algorithm – the *No-Loss* algorithm. While grid-based algorithms sometimes can “lose” messages, sending them to subscribers who are not interested in the particular message, but only happen to have close interests, the No-Loss algorithm guarantees that all subscribers receiving a message are interested in it, thus avoiding redundant communication in the network.

Subscription clustering algorithms form multicast groups, as well as produce data structures that are used for matching events to multicast groups. The last subsection describes matching algorithms for the two categories of clustering algorithms.

#### 4.1 Grid-Based Clustering Framework

Grid-based subscription clustering algorithms (or *cell clustering* algorithms) apply data clustering heuristics to the cells of a regular grid in event space  $\Omega$ . Data clustering algorithms are widely used for grouping “similar” objects. Similarity of objects is determined based on the value of a distance function. In this subsection we define feature vectors and distance function for clustering, and describe the application of several data clustering algorithms to our model in the following subsections.

**Feature Vectors.** Each cell  $a \subset \Omega$  has a subscriber membership vector  $\mathbf{s}(a) \in \{0, 1\}^{N_s}$  associated with it. By definition,

$$\mathbf{s}(a)_i := \begin{cases} 1, & \text{if there exists an index } j \text{ such that } b_{ij} \cap a \neq \emptyset, \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

In this vector non-zero elements correspond to subscribers interested in the cell.

The most commonly used strategy for partitional clustering is square-error minimization criterion, which amounts to minimizing sum of the squared Euclidean distances between feature vectors corresponding to objects over the entire set of objects. We use membership vectors as feature vectors of cells instead of using, for example, the coordinates of the cell center in event space  $\Omega$ . Using coordinates in  $\Omega$  for this purpose would lead to poorer solutions, since our goal is to create groups based on common interest, as opposed to *similar* interest. Grouping based on similar interest is considered, for example, in the work by Katz *et al* [19]. In our model, however, we assume that subscribers are only interested in events for which they have directly expressed interest, and there is a performance penalty to be paid when a message is delivered to a subscriber who not interested in it. On the other hand, when several subscribers are interested in the same events,

possibly scattered over  $\Omega$ , it is still logical to assign these subscribers to the same multicast group. Therefore we conclude that the comparison of sets of interested subscribers is more suitable for identifying common interest than comparison of coordinates in the event space.

**Distance Function.** Squared Euclidean distance between cells  $a$  and  $b$  can be computed as  $d_e^2(a, b) = \sum_i (\mathbf{s}(a)_i - \mathbf{s}(b)_i)^2 = \sum_i (\mathbf{s}(a)_i \oplus \mathbf{s}(b)_i)$ , where “ $\oplus$ ” denotes a binary operator of exclusive *or*. In our model, the probability density function of publications can be taken into account in order to better characterize our objective. We define the distance function  $d$  as:

$$d(a, b) := p_p(a) \sum_{i \in V_s} \max\{\mathbf{s}(a)_i - \mathbf{s}(b)_i, 0\} + p_p(b) \sum_{i \in V_s} \max\{\mathbf{s}(b)_i - \mathbf{s}(a)_i, 0\} \quad (2)$$

The value of  $d(a, b)$  is the expected number of messages sent to subscribers, who are not interested in them, if the cells  $a$  and  $b$  are combined in one group. Note that we can similarly define membership vectors and distance functions for sets of cells, simply replacing cells  $a$  and  $b$  with sets of cells in (1) and (2).

The function  $d$  characterizes *expected waste*, the notion first introduced in work by Gryphon group [4]. The objective of clustering in this formulation is to form groups in a way that minimizes expected waste. It is known that the heuristics used for clustering are not guaranteed to achieve a global optimum, but in practice the quality of the solution may be sufficient.

**Implementation Notes.** If for cells  $a$  and  $b$  it is true that  $\mathbf{s}(a) = \mathbf{s}(b)$ , then the cells can be combined in one group inducing zero expected waste. Our implementation during preprocessing stage scans continuous blocks of cells searching for repeating sets of subscribers, and joining the cells into *hyper-cells*, in other words sets of cells having the same membership vector.

Since the number of cells might become too high for the algorithm to handle, it would be helpful if we could somehow select the “most popular” cells for clustering, leaving the rest for unicast. Our implementation sorts hyper-cells (after the grouping step described above) by “popularity rating”  $r(\cdot)$  defined by

$$r(a) := p_p(a) \sum_{i \in V_s} \mathbf{s}(a)_i$$

and keeping only a fixed number of cells, having the largest values of popularity rating.

Our experiments show that after some number of cells, the improvement gained from feeding more cells to the algorithm becomes negligible. In fact, the more cells are given

to clustering algorithm, the worse the quality of solution becomes. This justifies the need for some implementation of outliers removal algorithms for detection of cells that have rather unique combination of subscribers. On the other hand, even without the outlier removal algorithm, clustering a large enough fraction of cells can lead to sufficiently good results.

### 4.2 K-Means and Forgy K-Means Clustering

The use of k-means clustering algorithm in pub-sub systems (with different objective, and therefore different feature vectors and distance functions) was proposed by Katz *et al* in [19]. We have studied the performance of original algorithm by McQueen and as well as the one of its variants by Forgy, see details in [9].

- |  |
|--|
| <ol style="list-style-type: none"> <li>0. Form initial <math>K</math> groups.</li> <li>1. Re-assign each cell to a closest group.</li> <li>2. Repeat step 1 until no cell can be moved.</li> </ol> |
|--|

**Figure 2. K-Means Clustering.**

Figure 2 shows generic pseudo-code for a k-means clustering algorithm for forming  $K$  multicast groups. In step 0, the initial partition is formed. For this purpose,  $K$  hyper-cells with the highest popularity rating  $r(a)$  are chosen to be centroids of groups, and the rest of hyper-cells is assigned to closest groups, based on the expected waste distance function. In step 1 each of the hyper-cells is examined, and assigned to the "closest" cluster. The distance to the cluster is determined using the distance function, measuring the distance between the membership vector of the hyper-cell and the membership vector of the cluster. The K-means algorithm updates the membership vector of the cluster each time a hyper-cell is moved. The algorithm by Forgy updates membership vectors of all clusters after step 1 is finished. A hyper-cell cannot be moved to another cluster if it is the last hyper-cell in its current cluster.

K-means type algorithms are proven to converge to a local optimum in a number of iterations, and in practice they converge quickly. Nevertheless, the processing can be stopped after any iteration, resulting in a feasible partition into  $K$  groups. This also provides an easy way to accommodate changes in cell membership, simply running a number of re-balancing iterations, when new subscribers arrive or subscription rectangles are changed in some other way.

### 4.3 Pairwise Grouping

Pairwise grouping for clustering point interest (or *pairs* for short) was proposed in the work by Gryphon group [4], and we have extended these ideas to the case of interest rectangles. It is a top-down clustering algorithm, which starts with each hyper-cell assigned to its own cluster. If the total

- |  |
|--|
| <ol style="list-style-type: none"> <li>0. Given <math>l</math> cells, form <math>l</math> groups, one cell in each.<br/>Group <math>g_i = \{a_i\}</math>, for each cell <math>a_i</math>, and <math>s(g_i) = s(a_i)</math>.</li> <li>1. Find <math>i</math> and <math>j</math> such that <math>i \neq j</math> and <math>d(g_i, g_j)</math> is minimized.<br/>Reset <math>g_i \leftarrow g_i \cup g_j</math>, update <math>s(g_i)</math>, and remove <math>g_j</math>.</li> <li>2. Repeat step 1 until there are only <math>K</math> groups left.</li> </ol> |
|--|

**Figure 3. Pairwise Grouping.**

number of clusters is larger than the required number  $K$ , the two groups with the minimum distance to each other are chosen, and combined. The membership vector for this group becomes a combination of vectors of the joined groups. The process is repeated until no more than  $K$  groups are left. The algorithm is summarized on Figure 3.

An approximate version of this algorithm inspects a fraction  $1/e$  of the total number of the group combinations during the search for best distance, stores the pair with the shortest distance from this fraction of groups, and terminates the search after that, if a smaller distance is found. This heuristic derives from a well-known solution to secretary problem ([17, p. 114]) for maximizing the chance of choosing the maximum value, if the decision must be done immediately. The algorithm modified in this way works faster, but it may obtain a poorer solution.

### 4.4 Minimum Spanning Tree Clustering

The use of the minimum spanning tree (*MST*) for clustering was proposed by Zahn (1971). Suppose we have a graph  $G$ , in which nodes correspond to hypercells, and there is an edge between each pair of nodes  $i$  and  $j$  of length  $d(i, j)$ . We will say that each node in itself is a component. Processing the edges in non-decreasing order of length, if the edge connects different components, combine the components into one, and proceed to the next edge. Continue until there are  $K$  components left.

- |  |
|--|
| <ol style="list-style-type: none"> <li>0. Given <math>l</math> cells, form <math>l</math> groups, one cell in each.</li> <li>1. <b>For</b> each pair in the order of increasing distance:</li> <li>2. If the hyper-cells of the pair are not in same group, combine the groups corresponding to cells.<br/>Repeat step 2 (for next pair) until there are only <math>K</math> groups left.</li> </ol> |
|--|

**Figure 4. Minimum Spanning Tree Clustering.**

This algorithm is similar to Kruskal's algorithm [11] for finding MST in graph  $G$ , except that this version stops when exactly  $K$  connected components are formed. Pairwise grouping proceeds in the similar fashion as the MST algorithm, but in the pairwise grouping case the distances are calculated between groups, not between cells. Therefore it is impossible to sort pairs by distance in advance, which



```

0. Set of rectangles:  $S := \mathcal{I}$ ;
   Rectangle weights  $w(b_{ij}) := p_p(b_{ij})$  for each  $b_{ij} \in S$ ;
   Subscriber node lists:  $u(b_{ij}) := \{i\}$  for each  $b_{ij} \in S$ .
1. Sort elements of  $S$  by  $w$ , such that:
   if  $s_l, s_m \in S$  and  $l < m$ , then  $w(s_l) > w(s_m)$ .
2. Retain only first  $T$  elements
   in sets  $S$ ,  $w$  and  $u$ , discarding the rest.
3. For each  $b_{ij} \in \mathcal{I}$  and  $s \in S$ , such that  $t := s \cap b_{ij} \neq \emptyset$ ,
   and  $i \notin u(s)$  do:
   if  $\exists r \in S$  such that  $r \equiv t$ ,
        $u(r) \leftarrow u(r) \cup u(t)$ ;  $w(r) \leftarrow p_p(r)|u(r)|$ .
   else  $S \leftarrow S \cup t$ ;  $u(t) \leftarrow u(t) \cup \{i\}$ ;  $w(t) \leftarrow p_p(t)|u(t)|$ .
4. Repeat from step 2 at most  $k$  times.
5. Sort  $S$  as in step 1.
6. Form  $K$  multicast groups corresponding to
   first  $K$  elements of  $S$ , grouping subscribers
   according to  $u(s_l)$  lists,  $l = 1..K$ .

```

**Figure 5. No-Loss Algorithm**

in effect leads to greater running time of the pairs algorithm, compared to MST on the same data.

#### 4.5 No-Loss Algorithm

The grid-based family of cell clustering algorithms works with cells of a regular grid, and each cell of the grid can be associated with one of the multicast groups. As a result, each subscriber whose interest overlaps with the cell, is assigned to the multicast group. Since interest rectangles are not aligned on cell borders, it is possible that an event sent to a multicast group will reach subscribers that are not interested in this particular event, as well as the ones for which this event was intended. The idea behind the *No-Loss* algorithm is to avoid this kind of wasted communication completely, forming multicast groups corresponding to areas aligned to interest rectangles borders. The algorithm (Figure 5) tries to find the “most popular” intersections of interest rectangles. The popularity (or *weight*) of an area in event space is measured by the number  $|u(s)|$  of subscribers interested in this area multiplied by the probability of publication in the area:  $w(s) = p_p(s)|u(s)|$ .

#### 4.6 Matching Subscriptions To Events

Each time an event arrives in the system, it has to be matched to multicast groups formed by a clustering algorithm in order to find out how to deliver the message. Matching must be done efficiently, since the delay caused by the matching algorithm directly affects the maximum throughput of the system. In this subsection we briefly introduce matching algorithms to illustrate how the data structures produced by subscription clustering algorithms is used in real-

```

1. Given an event  $\omega$ , find the corresponding cell
   of the grid in the event space  $\Omega$ .
2. If the event is matched to a cell:
   2.1. Denote the associated multicast group  $\mathcal{G}$ .
       Find the number (or proportion) of members of  $\mathcal{G}$ ,
       interested in  $\omega$ .
   2.2. If the number is above a predefined threshold:
       2.2.1. Send the message to group  $\mathcal{G}$ .
       Else
       2.2.2. Send it only to interested subscribers.
3. Else (if the event is not matched)
   Send it to the list of interested subscribers.

```

**Figure 6. Matching for Grid-Based Algorithms**

time, and refer the reader to our companion paper [16] for more detailed discussions.

Each group produced by a clustering algorithm can be described as a set of aligned rectangles in the event space  $\Omega$ . Therefore the problem of matching an event  $\omega$  to multicast groups can be reduced to the problem of searching among aligned rectangles in event space  $\Omega$  for the rectangles that contain a given point  $\omega$ . This general problem is most commonly solved using a special pre-built data structure called *R\**-tree (see [5]). In order to achieve better performance a modification of *R\**-tree algorithm, the *S*-tree algorithm described in [1] can be used instead.

**Matching in Grid-Based Algorithms.** Multicast groups formed by a grid-based algorithm are associated with cells of the grid in the event space. Each cell of the grid is either associated with one group or is not associated with any group. Therefore matching algorithm should find which cell the event falls into, and take different actions according to whether the cell is associated with a group or not. If there is no group associated with the cell, the message must be delivered using unicast. If there is an associated multicast group is found, the message is usually delivered via multicasting to this group. However, if we can determine how many of subscribers included in the matched group are actually interested in the message, we may be able to avoid unnecessary communication by sending the message via unicast only to those interested in it. This optimization can help to noticeably reduce communication. We shall further discuss the effects of this optimization, as well as present experiment results, in a companion paper [16]. Figure 6 summarizes the matching algorithm in pseudo-code.

**Matching in No-Loss Algorithm.** No-Loss interest clustering stage forms a list  $A \subseteq S$ , consisting of the first  $n$  elements of  $S$  in the order of decreasing density  $w$ . When an event  $e$  arrives, the matching algorithm in Figure 7 is applied. The algorithm (using an *S*-tree, for example) finds multicast

groups, as well as individual subscribers that are not included in the groups, whose interest rectangle includes the message.

1. If  $e \in s$  where  $s$  is a rectangle,  $s \in A$   
 let  $s$  be such that  $\forall t \in A : e \in t, w(t) \leq w(s)$   
 (i.e. it is the rectangle with greatest density).  
 Send message to multicast group formed of  $u(s)$ .  
 Send message via unicast to subscribers in  
 $(V_S \setminus u(s))$  that are interested in  $e$ .
2. Otherwise send message via unicast  
 to subscribers interested in  $e$ .

Figure 7. Matching for No-Loss Algorithm

## 5 Experiments

### 5.1 Experiment Model

These results have been obtained on 3 different models of subscription interest and message distribution, but on the same network consisting of 600 nodes and the same distribution of subscribers on this network. In all three models 1000 subscription rectangles were generated.




Figure 8. Network topology with 600 nodes.

We adopted the GT-ITM package from Georgia Institute of Technology [20] to generate a network with six hundred nodes according to a hierarchical scheme. This tool generates a hierarchical topology with transit blocks on top, stubs in the middle and nodes at the bottom. In our experiments,

we first generate three transit blocks, with an average of five transit nodes in each block. Each transit node is connected on average to two stubs, and each stub has an average of twenty nodes. Figure 8 illustrates its topology.

For a given network topology, we generate subscriptions for each node. We first generate one thousand subscriptions with a {40%, 30%, 30%} breakdown for the three transit blocks. Within each transit block, there is a Zipf-like distribution for the number of subscriptions between all the stubs connected to this transit block. Subscriptions are distributed according to another (common) Zipf-like distribution within each stub.

The generated interval subscriptions are of the form  $\{bst, name, quote, volume\}$ . The first field *bst*, which could stand for *buy*, *sell*, and *transaction*, takes value  $B, S$  and  $T$  with probabilities 0.4, 0.4, and 0.2, respectively. The center of the interval for the *name* field follows a normal distribution, with mean centered around the points specific to transit block number (3, 10 and 17), and standard deviation of 4. The length of this interval also follows a Zipf distribution. The intervals for the *quote* and *volume* fields are generated according to the same parametric distribution with different parameters. This parametric distribution takes values as follows:

- $(-\infty, +\infty)$ , with probability  $q_0$ ,
- $[n, +\infty)$ , with probability  $q_1$ , and  $n \sim N(\mu_1, \sigma_1)$ ,
- $(-\infty, n]$ , with probability  $q_2$ , and  $n \sim N(\mu_2, \sigma_2)$ ,
- $[n_1, n_2]$ , otherwise, center of interval  $\sim N(\mu_3, \sigma_3)$ ,  
 interval length follows a Pareto distribution.

The parameters are given in the following table:

	$q_0$	$q_1$	$q_2$	$\mu_1, \sigma_1$	$\mu_2, \sigma_2$	$\mu_3, \sigma_3$	$c, \alpha$
price	0.15	0.1	0.1	9, 1	9, 1	9, 2	4, 1
vol	0.35	0.1	0.1	9, 1	9, 1	9, 2	4, 1

The generation of the subscriptions are intended to mimic the real life scenario that people's interests in stocks are centered around the current prices, the popularity of the information for different stocks have a Zipf-like distribution, and the popularity of the participants also have a Zipf-like distribution.

The publications are the points in the subscription space, which are generated according to a mixture of multivariate normal distributions. The different peaks in the multivariate normal distributions represent the multiple hot spots where events are published more frequently. We studied three scenarios, which are mixtures of one, four and nine multivariate normal distributions. The means and standard deviations for the single mode multivariate normal distribution are (1, 1), (10, 6), (9, 2), (9, 6) for each of the four dimensions. The four mode distribution is constructed by sampling independent mixtures of multivariate normal distributions in each dimension. The mean and standard deviations for the first and fourth dimensions are (1, 1) and (9, 6), respectively.

The second dimension is a normal random variable with parameters (12, 3) with probability 0.5, and with parameters (6, 2) with probability 0.5. The third dimension is a normal random variable with parameters (4, 2) with probability 0.5, and with parameters (16, 2) with probability 0.5. Similarly, for the nine mode distribution, the parameters for the first and fourth dimensions remain the same. The third dimension is  $N(4,3)$  with probability 0.3,  $N(11,3)$  with probability 0.4 and  $N(18,3)$  with probability 0.3. The fourth dimension is  $N(4,3)$  with probability 0.3,  $N(9,3)$  with probability 0.4 and  $N(16,3)$  with probability 0.3.

It should be noted that this experimental framework is flexible enough to accommodate other probability distributions for the subscriptions and publications. In the following study, we constructed 1000 subscriptions for the network with 600 nodes generated by the GT-ITM package.

We performed experiments on the generated testbed to evaluate the performance of the different schemes for forming multicast groups under two multicast frameworks: network supported and application-level multicast schemes. Network supported multicast requires the network routers to have the multicast capabilities, to be able to recognize multicast groups, and forward the information to the proper members of the group. There are two types of multicast algorithms currently used in routers: *dense mode* and *sparse mode* multicast. The implementations differ in the amount of state information and in the structure of the routing tree. We assume the *dense mode* multicast where the routing tree is a shortest path tree rooted at publisher. The amount of state information is proportional to both the number of publishers and the number of groups. In recent years, the study of multicast mechanisms has focused on the application level multicast, which does not require full support at the network routers. The members of a multicast group communicates through unicasts. They form a minimum spanning tree and forward the messages from one member to another through the minimum spanning tree. We also evaluated the impacts of the application level multicasts for the different algorithms.

## 5.2 Experiment Results

The following plots summarize simulation results, in which cost of communication was computed by summing up the edge costs (generated by GT-ITM package) on links on which communication takes place. Since absolute values of costs differ for different networks, we normalize the costs to make comparison easier. Thus the vertical axis in most plots shows “improvement percentage” over unicast. In other words, 0% communication cost improvement is achieved by using unicast to deliver each message. 100% cost improvement corresponds to the cost of delivering each message to a multicast group specially formed only of clients interested in this particular message, which is the best possible, and in the worst case requires as many as  $O(2^{N_s})$  multicast groups.

The goal of clustering algorithms is to get as close to this performance bound as possible, while using no more than  $K$  groups.

The absolute communication costs depend on different parameters. For the case of one-mode Gaussian subscription distribution, the unicast is 7139, the broadcast is 8536, and the ideal solution of network supported multicast is 1763.

Figure 9 shows how the communication cost changes as more groups become available for different algorithms. We are interested in algorithms that demonstrate monotone improvement, since intuitively when more groups can be formed we expect the algorithms to do a better job. Each plot is shown for network-level multicast and application level multicast. While application-level multicast results in slightly higher costs, the trend remains the same, and the algorithms that perform better under network multicast maintain their leadership under application-level multicast.

Performance of k-means is almost the same as one of Forgy k-means. Approximate pairs curve closely follows the curve of the pairs algorithm, as one can see in figure 12, so the latter is not shown in order to make plots readable.

The algorithms were run with the following parameters. K-means and Forgy used 6000 rectangles and maximum of 100 iterations (usually the number of actual iterations was less than 20). Approximate pairs algorithm was using only 2000 rectangles. The time complexity graph in figure 12 shows that in this case time complexity is almost the same as the one of K-means. No-Loss algorithm was run with 5000 rectangles kept after intersection and 8 iterations. Figure 10 shows how the algorithm depends on these parameters. MST was run with 6000 rectangles.

Figure 11 shows that the trend of the algorithm performance does not depend greatly on the network topology. The left plot is taken from figure 9, and the right plot is obtained from simulation on a different network and subscription assignment generated according to the same parameters, but with different random seeds. While the grid-based clustering algorithms achieve local optimal solutions at best, in practice the solutions usually are good enough, reaching to 60% in this case.

Figure 13 combines the left and right plots of figure 12. Notice that Forgy and K-Means performance plots go down, when given more time. This only means that when more data is given to the algorithms. It consumes time without improving performance, It can even worsen the quality of solution, because the number of wasted messages becomes large. Data clustering algorithms usually make use of outlier removal techniques to avoid these problems. We leave the study of outlier removal effects for future work, simply noting here that the parameter that regulates the number of cells given to the clustering algorithm regulates its performance, as well as the time it takes to produce a solution, as shown in figure 12.

The results in figures 13 and 12 indicate that Forgy clus-




Figure 9. Algorithms Compared.

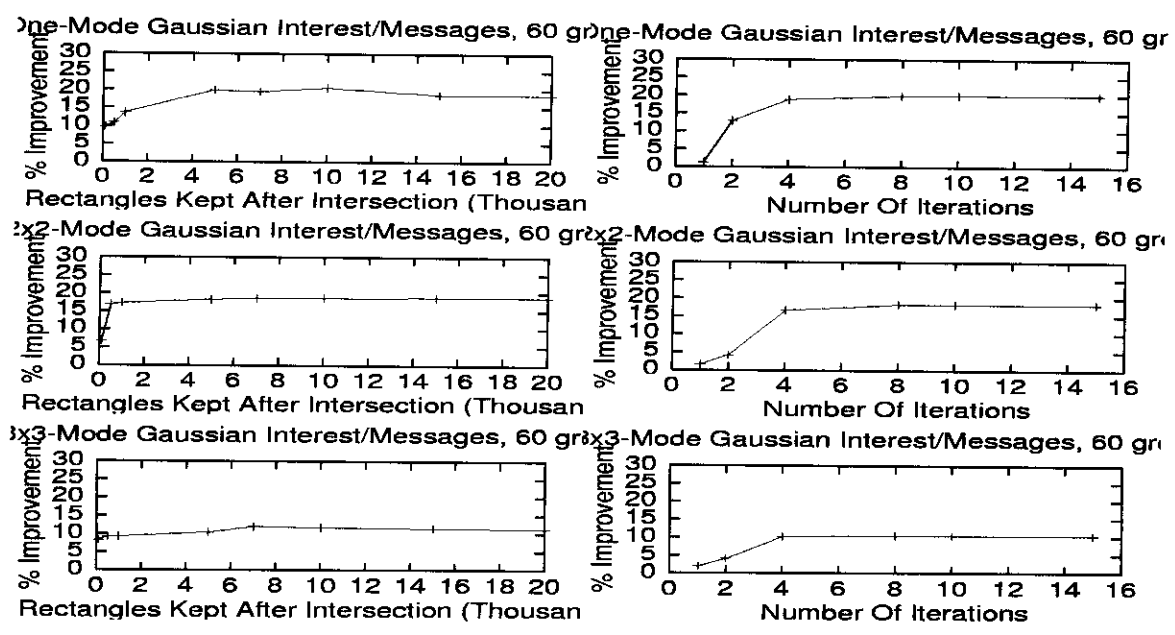


Figure 10. Effect of the number of rectangles kept after intersecting and of the number of iterations on no-loss algorithm.

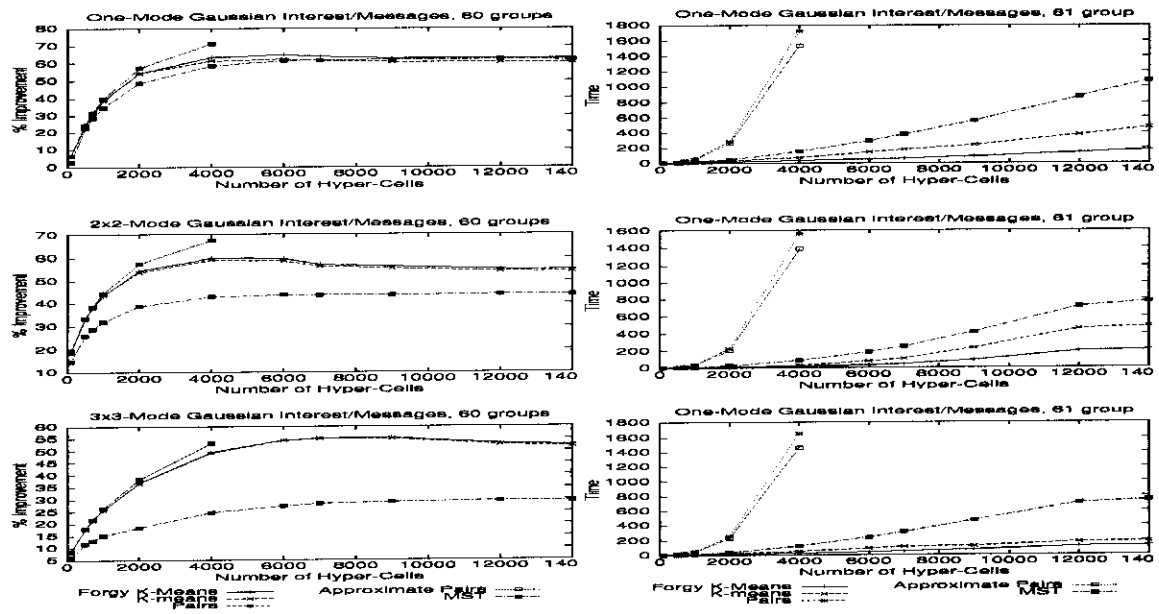


Figure 12. Effect of the size of input data on the performance and running time of cell clustering algorithms.

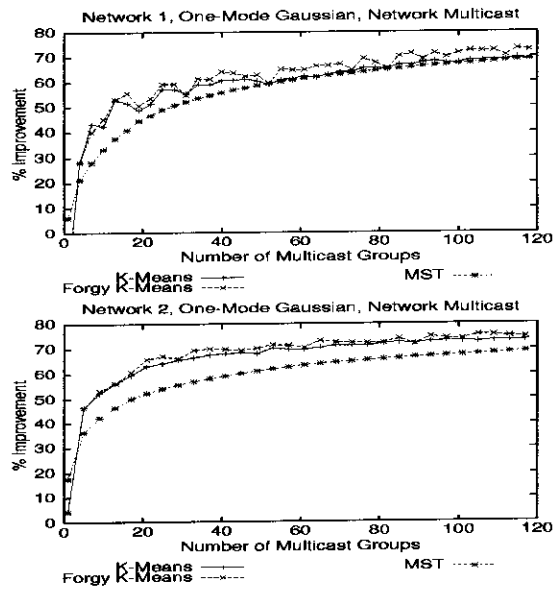


Figure 11. Performance of K-Means, Forgy and MST on 2 different networks generated according to the first model (one-mode Gaussian distribution.)

tering should be preferred over k-means, since it gives better or comparable results faster. Cell-based clustering works well when the dimensionality of the event space is not too high, and the granularity of subscription interest is not too high. In this case these clustering methods should be preferred over No-Loss algorithm. We leave high-dimensional case for future study.

## 6 Conclusions and Discussions

We have considered the issue of efficient communication schemes based on multicast techniques for content-based publication-subscription systems. We have proposed and adapted clustering algorithms to form multicast groups for these content-based publication-subscription systems. These algorithms perform well in the context of highly heterogeneous subscriptions, and they also scale well. An efficiency of 60% to 80% with respect to the ideal solution can be achieved with a small number of multicast groups (less than 100 groups in our experiments).

Our experiments indicate that under several assumptions, which include high degree of regionalism of interest, and distribution of the interest close to the distribution of message parameters, the Forgy algorithm should be preferred for most purposes. Iterative clustering algorithms (K-means and Forgy) seem to be better suited for subscription dynamics, although other algorithms can be adapted as well. Hierarchical clustering algorithms (MST and Pairs) has worse performance than iterative clustering, but has the advantage of




Figure 13. Dependence of quality of solution on time.

monotone improvement: when more multicast groups become available for the system to form, the new groups are formed by sub-division of existing ones. The analysis of influence of algorithm parameters on algorithm performance presented in this paper must be helpful in practical implementation of the algorithms.

There are still many open issues to be addressed in the future. In what follows, we discuss a few of them.

**Non-Rectangular Subscriptions.** It is obvious that rectangular interest sets represent only a special case of real-world structure of subscription interests. Work has been done on the subject of more general formulation of subscription interest. Many proposed solutions employ a dialect of the SQL query language, for example the work on continual queries at Georgia Tech or the Gryphon project papers. Proposed algorithms can be adapted to make use of non-rectangular subscription interest sets by rounding the sets to appropriate shapes. While the no-loss algorithm relies on the rectangular interest set assumption, it is not very important to the other (grid-based) algorithms. The same grid data structures can be created without requiring the sets to be rectangles. Of course this requires rounding to cell borders, but this is the same kind of tradeoff that exists in rectangular case.

**Network Topology Assumptions.** The transit-stub model of network topology [20] is widely used for modeling Internet networks, but does not always reflect the real-world topology correctly.

One simple variant of the topology that may be evaluated in the future is the one in which for each subscription a special node is added, connected only to the corresponding "subscriber" node, where the links between this node correspond to the last-mile connection of a client to ISP. The intuition behind this is that the last-mile links are the slowest and the most congested ones, and we may want to reduce the traffic going through these links as much as possible.

**Real-World Evaluation of Heuristics.** Evaluation of the algorithms on the real-world data would be very helpful for making decisions about implementation of the algorithms. However although real-time stock trading data, for example, can be used to simulate a stream of events coming into the system, we do not see a straightforward way of obtaining information about the structure of real-world subscriptions. This information includes the distribution of preference rectangles in the event space, as well as the distribution of the subscriptions in the network. We have made assumptions on these distributions, but we do not have the evaluation of how well they reflect the real-world systems, and in which systems these assumptions can be employed.

**Size of Messages.** In our experiments we did not simulate real network packets, implicitly assuming that there are no

delays caused by congestion of network links. This is a reasonable assumption to make when the message size is relatively small – of the order of 1K or less, like a stock quote, for example. This allows the sending of 3-5 messages per second even through a modem dial-up connection, which seems to be enough for many real-life applications. Weather updates can be a good example of small and relatively infrequent messages.

If the messages are of large sizes, or the variability of message sizes is big, with some messages being of significantly large size, a different type of communication cost evaluation must be used. In particular, such parameters as network bandwidth and message arrival process have to be taken into account.

The systems that intend to deliver content of larger size, for example digital movies or software updates, need to employ different approaches and algorithms. We were designing the algorithms assuming that setting up a multicast group is much more time consuming than sending a message, which is true for small messages. This assumption leads to the two-stage design of the algorithms: multicast groups are fixed in the beginning and then the messages are sent to the existing set of groups. However for big messages the assumption no longer holds; it becomes cheaper to set up a multicast group including only the interested subscribers, and do this each time new content becomes available.

**Subscription Dynamics.** In our model we have assumed that the set of subscriptions can be fixed, and then the clustering algorithms can be applied. In reality subscribers change their subscriber preferences, join and leave the network. Therefore the clustering groups need to be constantly updated.

Katz *et al* propose to use iterative clustering to iteratively improve the solution when subscription structure changes [19]. Although a different type of distance measurement is used in that paper, the same kind of technique can be used to accommodate changes in cell clustering algorithms, especially in the k-means and Forgy modifications, that are well very suited for dynamic changes in feature vectors due to their iterative improvement nature.

When these algorithms are used, each modification of subscriptions results in changes in membership vectors in some of the cells. different multicast groups iteratively, possibly resulting in changes in multicast group membership. This re-balancing process does not necessarily have to take place at once, several iterations can be stretched over time. We only need to make sure that all messages are delivered to all interested subscribers.

Future work should include evaluation of different methods of re-balancing cell membership in clusters. It is likely that this iterative improvement feature can be counted as an additional advantage of iterative cell partitioning (clustering) algorithms, including Forgy and classic k-means. Also note

that performance evaluation shows that these algorithms provide the best results in shortest time, which supports the suggestion that these algorithms should be preferred in pub-sub systems that satisfy the constraints of our model.

**Decentralized Matching and Message Lifetime.** Since the message delivery is distributed, an average message travels through many nodes in the network. And, especially in application multicast case, these nodes are able make decisions about distributing the messages downstream. In our model we have assumed that the matching of messages to groups or individual subscribers is done once: the first “intelligent” node that receives the message decides how to route it.

In an alternative approach, which we will refer to as *decentralized matching*, and which is a part of approach described in several Gryphon project papers [2, 14], each intermediate node knows about the preferences of its neighbors, and matches each event against its specific data structures to find the neighbors to which the event must be forwarded next. This approach may save communication and matching time. However in practice dynamics of subscriptions require subscription changes to propagate quickly in the network, which makes this approach difficult to implement. The following example illustrates this difficulty. Suppose that a subscriber announces a new subscription. The interests of its neighbors change as a result, and this can lead to a change in interest of their neighbors, and so on. If a message that matches subscriber’s interests is posted after the subscriber announced her interests, but before all of interest updates in the network are completed, the message may never reach the subscriber.

Current work on combining “push” content delivery with databases is also related to this subject. These systems make sure that the subscribers receive messages even if they are not functioning in always-on mode, and that they can retrieve past events even if they were not subscribed at the time the event is published. Clients can connect at any time and get information on topics of interest, and receive updates corresponding to their expressed interest while they are online. See IBM Research report [6] for an example of this approach.

## References

- [1] C. Aggarwal, J. Wolf, P. Yu and M. Epelman, “Using Unbalanced Trees for Indexing Multidimensional Objects”. *Knowledge and Information Systems*, Vol. 1, pp. 309-336, 1999.
- [2] M. Aguilera, R. Strom, D. Sturman, M. Astley and T. Chandra. “Matching Events in a Content-Based Subscription System”, *Proceedings of the 18th Annual ACM Symposium on Principles of Distributed Computing (PODC '99)*, Atlanta, USA, May 1999.
- [3] K. Almeroth, “The Evolution of Multicast: From the Mbone to Inter-Domain Multicast to Internet2 Deployment”, *IEEE Network*, January/February 2000.

- [4] G. Banavar, T. Chandra, B. Mukherjee, J. Nagarajarao, R. Strom and D. Sturman, "An Efficient Multicast Protocol for Content-Based Publish-Subscribe Systems" *Proceedings of the 19th International Conference on Distributed Computing Systems*, May 1999.
- [5] N. Beckman, H. Kriegel, R. Schneider and B. Seeger, "The R\*-Tree: An Efficient and Robust Method for Points and Rectangles," *Proceedings of the ACM SIGMOD Conference*, pages 322-331, May 1990.
- [6] C. Binding, S. G. Hild, R. Hermann and A. Schade, "Intelligent Messaging for the Enterprise." IBM Research Report, RZ 3353 (# 93399), June 2001.
- [7] F. Fabret, F. Llirbat, J. Pereira, and D. Shasha. "Efficient matching for content-based publish/subscribe systems." Technical report, INRIA, 2000. <http://www.caravel.inria.fr/pereira/matching.ps>.
- [8] R. E. Gruber, B. Krishnamurthy, and Panagos, "The Architecture of the Ready Event Notification Service", *Proceedings of the 19th IEEE International Conference on Distributed Computing Systems Middleware Workshop*, 1999.
- [9] A. K. Jain and R. C. Dubes. 1988. *Algorithms for Clustering Data*. Englewood Cliffs, New Jersey: Prentice Hall.
- [10] I. Kamel and C. Faloutsos, "On Packing R-Trees," *Proceedings of the 2nd International Conference on Information and Knowledge Management*, 1993, pages 490-499.
- [11] J. B. Kruskal, "On the shortest spanning subtree of a graph and the traveling salesman problem.", *Proceedings of the American Mathematical Society*, volume 7, pages 48-50, 1956.
- [12] B. Levine, J. Crowcroft, C. Diot, J. Garcia-Luna-Aceves and J. Kurose, "Consideration of Receiver Interest for IP Multicast Delivery", *Proceedings of Infocom 2000*, Tel Aviv, March 2000.
- [13] M. Oliveira, J. Crowcroft and C. Diot, "Router Level Filtering for Receiver Interest Delivery", *Proceeding of NGC*, 2000. published by ACM. Stanford (USA). November 8-10, 2000.
- [14] L. Opyrchal, M. Astley, J. Auerbach, G. Banavar, R. Strom and D. Sturman, "Exploiting IP Multicast in Content-Based Publish-Subscribe Systems", *IFIP/ACM International Conference on Distributed Systems Platforms*, New York, April, 2000.
- [15] D. Pendarakis, S. Shi, D. Verma and M. Waldvogel, "ALMI: An Application Level Multicast Infrastructure", Symposium on Internet Technologies, March 2001.
- [16] A. Riabov, Z. Liu, J. Wolf, P. Yu and L. Zhang, "On the Matching Problem in Content-Based Pub-Sub Systems", in preparation.
- [17] S.M. Ross, "Introduction to Probability Models", Sixth Edition, Academic Press, San Diego, 1997.
- [18] T. Sellis, N. Roussopoulos and C. Faloutsos, "The R+ tree: A Dynamic Index Structure for Multi-Dimensional Objects", *Proceedings of the 13th International Conference on VLDB*, pages 507-518, September 1987.
- [19] T. Wong, R. Katz and S. McCanne, "An Evaluation of Preference Clustering in Large-Scale Multicast Applications", *Proceedings of IEEE Infocom 2000*, Tel Aviv, March 2000.
- [20] E. Zegura, K. Calvert and S. Bhattacharjee, "How to Model an Internetwork", *Proceedings of IEEE Infocom 1996*, San Francisco, 1996.
- [21] New Era of Networks Inc. NEONet <http://www.neonsoft.com/products/NEONet.html>.