

# Efficient Query Subscription Processing in a Multicast Environment \*

Arturo Crespo, Orkut Buyukkokten, and Hector Garcia-Molina  
Stanford University

## Abstract

This paper introduces techniques for reducing data dissemination costs of query subscriptions. The reduction is achieved by merging queries with overlapping, but not necessarily equal, answers. The paper formalizes the query-merging problem and introduces a general cost model for it. We prove that the problem is NP-hard and propose exhaustive algorithms and three heuristic algorithms: the Pair Merging Algorithm, the Directed Search Algorithm and the Clustering Algorithm. We develop a simulator for evaluating the different heuristics and show that the performance of our heuristics is close to optimal.

## 1 Introduction

With information dissemination (information push), data is delivered from a set of *producers* to a (typically) larger set of *consumers*. Examples of dissemination-based applications include information feeds (e.g., stock and sports tickers of news wires), traffic information systems, electronic newsletters, and entertainment delivery [16]. We focus on a type of dissemination system where the consumers in advance submit *subscriptions* defining their interests. Each subscription may include one or more queries over the data that the producers hold or generate. The producers run the queries periodically, disseminating information of specific interest to the consumers. Systems such as Pointcast [30], Marimba [28], Backweb [3], and Airmedia [2] are examples of subscription-based dissemination.

Subscription-based dissemination services are well suited to users' needs, but can be very expensive. As a real world example, a 1996 study that monitored Internet traffic found that more than 17% of the HTTP Internet traffic involved PointCast [19]. Additionally, PointCast "pulls" saturated company networks so much that large corporations have limited or even outlawed the use of PointCast on their desktop PCs [29]. In this paper we study a novel technique that can significantly reduce traffic and server loads. The overconsumption of resources in subscription-based dissemination services is the result of three factors. First, the network is point-to-point (i.e., the answers to each query subscription are transmitted separately to each consumer). Second, each query is processed independently. And third, previous work do not make full use of the processing power of clients. Instead, clients are considered "dumb" processes that are unable to perform any post-filtering of data they receive.

---

\*This work was partially supported by DARPA, as part of the BADD Technical Enhancements Program and the NASA Interchange for Joint Research NCC 2-5347.

The overhead of point-to-point dissemination can be reduced by using a multicast network. For example, consider the case where  $n$  clients issued exactly the same query in their subscriptions. A subscription service using a point-to-point network will process and transmit the answers to those queries  $n$  times, while a multicast-based service will establish a “channel” for the answer and will transmit the answer only once [19].

However, in many applications, it is unlikely that a large number of clients will issue exactly the same query, preventing us to fully exploit the advantages of a multicast network. In this paper, we present algorithms for efficient use of a multicast network for such applications. We achieve this by considering merging not only identical queries but also queries with answers that overlap significantly. By merging these queries, the server has to process fewer queries and the amount of information sent may be reduced. (As we will see later, in some cases, merging queries might *increase* the data sent.) On the negative side, the merged answers may contain some data that is irrelevant to a client. As a result, the client needs to make use of their processing power and apply a post-filtering *extraction query* over the received data in order to obtain the answer to its original query. For example, say we merge queries  $q_1 : \sigma_{2 \leq A \leq 40} R(A)$  and  $q_2 : \sigma_{3 \leq A \leq 41} R(A)$  into  $q_3 : \sigma_{2 \leq A \leq 41} R(A)$ . The server can then process this single query and send the result,  $ans(q_3)$ , to the clients that issued  $q_1$  and  $q_2$ . The  $q_1$  client will need to *extract* the  $q_1$  answer from  $ans(q_3)$  by applying the extraction query  $q_1 : \sigma_{A \leq 40}(q_3)$ . Similarly, the  $q_2$  client applies its own extraction query to eliminate all elements less than 3. By merging  $q_1$  and  $q_2$ , we reduce both the amount of work done by the server to process the query and the amount of information sent to the clients; however, this is at the expense of having to post-process the messages at the clients.

In this paper we address the query merging tradeoffs. We present a framework for studying query merging (sometimes called logical-channel building) and its costs. We present a variety of algorithms for merging, some optimal, and some heuristic. We study the complexity of the algorithms. We use a simulation tool to evaluate their performance (i.e., time required for merging, and dissemination costs saved).

We start by presenting a motivating example (Section 2). We then specify our problem more formally (Section 3) and define our cost model (Section 4). Next, a specific scenario using geographic queries is considered (Section 5). The algorithms are presented in Section 6, and their evaluation in Section 7.

## 2 Motivating Example

To illustrate we use the DARPA Battlefield Awareness and Data Dissemination initiative (BADD), which funded this work. The goal of BADD is to develop an operational system that delivers to combat troops an accurate, timely, and consistent picture of the battlefield and provides access to key transmission mechanisms and worldwide data repositories. Figure 1 outlines the relevant components of the BADD architecture.

In BADD, a database receives new information (e.g., satellite images, intelligence reports) from a set of *information sources*. The database also receives queries from *operation units*, answers the queries (on an on-going basis), and disseminates the answers. The operation units are limited capacity computers that can perform simple operations on the data received.




Figure 1: The BADD Scenario

A common request in the BADD environment is information (troop presence, weather, topography, etc.) about a geographical area. For these requests, sources typically associate a geographical location with each object. For example, if the data source is a relational database, it may have the schema  $R(\textit{longitude}, \textit{latitude}, \textit{attributes})$ , where the pair  $(\textit{longitude}, \textit{latitude})$  identifies the location, and  $\textit{attributes}$  describes the object. This database can be visualized as in Figure 2(a). The dots in the figure represent the objects that have a given longitude and latitude. As stated before, operation units will query this database for objects inside a geographical area. For simplicity, we will assume that such area is a rectangle, defined by two coordinates  $(c_1, c_2)$  and  $(c_3, c_4)$ . The queries over the database will have the form:  $\sigma_{(c_1 \leq \textit{latitude} \leq c_3) \wedge (c_2 \leq \textit{longitude} \leq c_4)} R$ . Figure 2(b) illustrates this query. Although we will use this simple scenario as a running example, we want to stress that our algorithms can handle more complicated queries and database schemas.

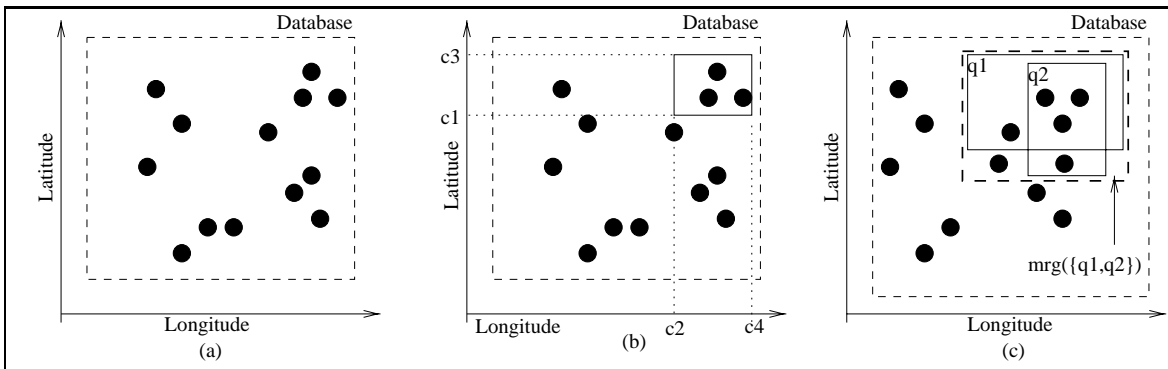


Figure 2: A Sample BADD Database and queries

In Figure 2(c) we illustrate the merging of two BADD queries  $q_1$  and  $q_2$ . Since these queries are very similar, it may be advantageous to merge them in to a single query  $mrg(\{q_1, q_2\})$ . Note that the answer to the new query will contain objects that were not in the answer of  $q_1$ , or in the answer of  $q_2$ , or both. The operation units that receive the answer of  $mrg(\{q_1, q_2\})$  must be able to derive from it the answers to  $q_1$  and  $q_2$ .

### 3 Problem Specification

Our objective is to reduce the cost of answering a set of query subscriptions made by clients to a server. We attempt to reduce the cost by finding a (possibly) different set of queries, with lower processing and transmission costs, from which the clients can derive the answers to their original queries. In this section, we discuss our conceptual model in detail.

#### 3.1 Conceptual Model

The conceptual model for a query subscription service is shown in Figure 3. In this model, we have a set of clients,  $C = \{c_1, \dots, c_n\}$ , that require information. The information need of  $c_i$  is described by a set of subscriptions. Each subscription consists of a query and its timing requirements (e.g., how often it should be run). For simplicity, we assume that all subscriptions have identical timing requirements. Thus, we can view the subscriptions of client  $c_i$  simply as a set of queries  $Q_i$ . We call  $Q$  the set of all queries received by the server.




Figure 3: Subscription Service Conceptual Model

Clients send their sets of queries to a *server*. The server periodically processes the queries against a database, and sends *answers* to the clients. Before processing queries, the server runs a merge algorithm that combines “similar” queries. The output of the merge process is a collection  $\mathcal{M} = \{M_i\}$  where each  $M_i$  contains the queries that are merged. The queries in each  $M_i$  are merged into a single query,  $mrg(M_i)$ . We use  $ans(q)$  to represent the answer to query  $q$ . Thus, the server generates  $ans(mrg(M_i))$  for each  $M_i$  in  $\mathcal{M}$ . For completeness, we require that  $\cup_i Q_i = \cup_i M_i$ . Similarly, we require that  $ans(q) \subseteq ans(mrg(M_i))$  for every  $q \in M_i$ . We call the difference between the answer to the merged query sent to a client,  $ans(mrg(M_i))$ , and the original query,  $ans(q)$ , the *irrelevant information for q* sent to the client.

To illustrate these concepts, say client  $c_1$  submits queries  $Q_1 = \{x, y\}$ , and  $c_2$  submits  $Q_2 = \{z\}$ . The server may merge them into  $M_1 = \{x, z\}$  and  $M_2 = \{y\}$ . Then the server runs  $mrg(M_1)$  and  $mrg(M_2)$  against the database and generates  $A_1 = ans(mrg(M_1))$  and  $A_2 = ans(mrg(M_2))$ . Note that  $A_1$  needs to be sent to both  $c_1$  and  $c_2$  and that each must apply an *extractor* to obtain the desired answer. For example, the extractor,  $e$ , that  $c_1$  applies to  $A_1$  should yield  $e(A_1) = ans(x)$ . Thus, when the server sends  $A_1$  out, it must include a *header* containing the following information:

- A list of clients that should receive  $A_1$ .
- For each such client  $c$ , one or more pairs,  $(e, q)$ , where  $e$  is an extractor and  $q$  is a query identifier. The extractor  $e$  is what client  $c$  needs to apply to obtain the answer to its original query  $q$ .

Note that more than one  $(e, q)$  pair is needed if multiple  $c$  queries are involved in  $A_1$ . If clients do not need to know what queries generated answers, then the query identifiers are not required. In our example, the information sent with  $A_1$  would be  $c_1 : (e_x, x)$  and  $c_2 : (e_z, z)$ . Client  $c_1$  then applies  $e_x(A_1)$  to obtain its answer to query  $x$  while  $c_2$  applies  $e_z(A_1)$  to obtain its answer.

There are many options for implementing extractors. For example, the server could tag each individual answer object with the identifier of the query that generates the object, or with the identifier of the client that should receive the object. Then each extractor only needs to look for the appropriate tags. In some cases, the extractor for a query is the query itself. In particular, this happens when queries only have selections and projections. A related issue is which component generates an extractor. Above we assumed that extractors were generated by the server and sent with answers. However, if the client can deduce its extractors (e.g., if the extractor is the original query itself), then the server need not send them.

Note that our basic model does not specify what kind of network is used to send queries to the server and answers to the clients. In the next section, when we discuss cost, we will introduce the multicast network to the model. This model is extended in [10].

## 4 The Cost Model

As described in Section 3.1, the server receives a set of queries  $Q$  and outputs a set  $\mathcal{M}$  where each of its elements is a set of queries to be merged. The query merging problem is to find the set  $\mathcal{M}$  with the minimum cost. The input for the problem is a cost function  $cost()$ , a merge procedure  $mrq()$ , and a set of queries  $Q$ . The output is a collection  $\mathcal{M}$  such that the total cost,  $cost(\mathcal{M})$ , is minimized.

The cost of processing the queries and sending the answers back is represented by the total resources consumed. The total resources consumed are the sum of all the resources used by the server, the network, and the clients. The costs involved in our model can be summarized as follows:

- Server cost to run the merging algorithm and to process the merged queries.
- Cost of transmitting the answers of the merged queries.
- Client cost of applying the extraction procedure.

In order to compute the resources used, we need to estimate the size of the query answers and the cost for computing them. Such estimate can be obtained using well-known database system techniques [27]. We use  $cost(q)$  to denote the estimated cost of retrieving  $q$ 's answer. The estimated total cost of retrieving all the answers (equal to  $cost(mrq(M_1)) + cost(mrq(M_2)) + \dots + cost(mrq(M_m))$ ) will be denoted as  $cost(\mathcal{M})$ . We will denote the estimated size of  $q$ 's answer as  $size(q)$ . The total size of all the answers (equal to  $size(mrq(M_1)) + size(mrq(M_2)) + \dots + size(mrq(M_m))$ ) will be denoted as  $size(\mathcal{M})$  (note that this is the total amount of data that the server needs to transmit to the clients).

As stated before, clients extract the answers to each of their queries *independently* by applying an extraction query to the messages they receive. The independence assumption means that a client may do some redundant work. For instance, consider a client that submits two queries that are then merged by the server. The client receives a single message containing both answers. It

processes the message once to extract the answers for the first query, and then again to get the second set of answer objects. In extracting the first answer, the client will consider some objects as irrelevant, even though they will be later found to be relevant for the second query. We believe this independent processing model is the most realistic since clients are expected to be relatively simple and unable to process different queries concurrently. We will denote the size of the irrelevant information for query  $q_i$  by  $u_i = size(mrg(M_j)) - size(q_i)$ , provided  $q_i \in M_j$ . We will call  $U(Q, \mathcal{M})$  the sum of all  $u_i$  ( $U(Q, \mathcal{M}) = \sum_{q_i \in Q} (u_i)$ ).

The resources used by each component of the system can be computed as follows:

- **Server cost:** If the complexity of the merging algorithm is low, its cost will be insignificant in a subscription service. Therefore we will ignore the cost of executing the merging algorithm.

The other component of the server cost is the time for computing and retrieving the answers from the database. The cost of computing the query answers will be denoted as  $K_A \cdot cost(\mathcal{M})$ , where  $K_A$  is a proportionality constant.

$$Cost_{server} = K_A \cdot cost(\mathcal{M})$$

- **Network cost:** Our network model assumes a multicast medium; namely, one where we can establish channels that allow sending data from one server to many clients. (In the extended version of this paper [10] we also consider a multicast network with a *fixed* number of physical channels.) The network cost will be proportional to two factors. First, the network resources consumed are proportional (by factor  $K_T$ ) to size of the data being transmitted ( $size(\mathcal{M})$ ). Since we expect the size of the header and the size of the queries to be very small compared to the size of the data, we will ignore these when computing the size of an answer message. Second, in some cases we may need to establish network connections or “logical channels” for each  $M_i$  set. Messages then just include a logical channel id, and clients can subscribe to one or more channels. The cost of maintaining logical channels (e.g. table space in the routers, or operating system connection overhead) is proportional (by a factor  $K_M$ ) to the number of merged queries transmitted. Thus,

$$Cost_{network} = K_M \cdot |\mathcal{M}| + K_M \cdot |\mathcal{M}|.$$

- **Client cost:** As answers are multicast, clients need to spend resources receiving the information they want plus the irrelevant information added by the merging algorithm. These resources are proportional to the amount of data received. The total amount of relevant data received by the clients is  $\sum_{q_i \in Q} size(q_i)$ , while the total amount of irrelevant data is  $U(Q, \mathcal{M})$ . (Recall that queries are processed independently at each client, as we discussed earlier.) Therefore, the cost at the clients is equal to  $K_U \cdot (\sum_{q_i \in Q} size(q_i) + U(Q, \mathcal{M}))$ , where  $K_U$  is a proportionality constant. Note that when comparing merging alternatives, we can ignore the cost of listening to the relevant data since this cost does not depend on the merging algorithm and it will cancel out in the comparison. (Of course, when computing actual costs we need to consider both costs.) In the following expression we focus only on the differential cost between merging strategies.

$$Cost_{clients} = K_U \cdot U(Q, \mathcal{M}).$$

Using the three cost components, we can compute the total cost as:

$$Cost_{total} = Cost_{server} + Cost_{network} + Cost_{clients}$$

$$Cost_{total} = K_A \cdot cost(\mathcal{M}) + K_M \cdot |\mathcal{M}| + K_T \cdot size(\mathcal{M}) + K_U \cdot U(Q, \mathcal{M}).$$

## 5 Geographic Queries

The framework and cost model presented so far is very general. The parameters in the cost model allow us to handle a wide range of capabilities in the servers and in the clients. We can model scenarios with clients ranging from very simple palm devices to sophisticated field computers. Similarly, we can handle servers having the functionality of a simple file system, to servers supported by a full fledged database. However, due to the limited space available in this paper, we will focus on a particular scenario. To illustrate how the merge procedure may operate and how the cost model can be used, we will use geographic query example presented in Section 2.

### 5.1 The Merging Procedure for Geographic Queries

As before, we consider the database to be a single relation  $R$ , that has position attributes (e.g., “latitude” and “longitude”), as well as other attributes describing that position. A geographic query has the form  $\sigma_{(c_1 \leq latitude \leq c_3) \wedge (c_2 \leq longitude \leq c_4)} R$ .

In Figure 4 we illustrate three different merge procedures that can be used in this geographic query scenario. In the figure, the solid lines represent the queries, and the dotted line represents the result of the merge procedure. Figure 4(a) shows the *bounding rectangle merging procedure*, the merging procedure introduced in Section 2. This procedure merges a set of 2-dimensional selection queries into a single 2-dimensional selection query. We can visualize this merged query as the smallest rectangle that bounds the original queries. The bounding rectangle merging procedure is very simple (and therefore fast to execute). Additionally, it is easy to extract the answers to the original queries from the answers to the merged query, as we just need to re-apply the original geographical query on the received answers. However, a disadvantage is that the answer includes objects that will be irrelevant to some or all of the input queries.

There are other possible merge procedures for the geographic query scenario. Figure 4(b) shows the *bounding polygon merging procedure*. This procedure also generates a single merged query, but, the query may have disjunctions. Although, the merge query contains less irrelevant information than the bounding rectangle merging procedure, irrelevant information is still present (the area of the polygon outside each query is irrelevant to the query). We can again use the original query as the extraction function for this merging procedure. Figure 4(c) shows a merge procedure that completely eliminates irrelevant information. However, five “merged” queries are generated. A client implementing the extraction function for this merging procedure needs to combine the answers to the five merged queries in order to find the answer to the original query.

In summary, there are many choices for merge procedures that trade off complexity of the merged query, complexity of the extractor and amount of irrelevant information added.

In this paper, we have assumed that  $|mrg(M)| = 1$ . However, our model can be extended to the case when  $|mrg(M)| > 1$  by taking the union of the answers in  $mrg(M)$  to create a single answer.

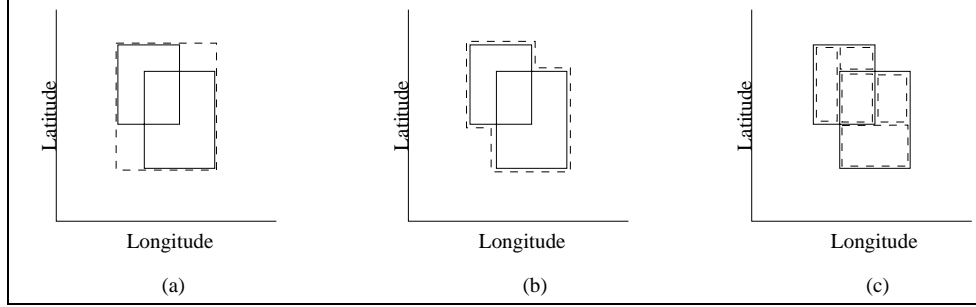


Figure 4: Three Different Merge Procedures

## 5.2 Cost Model for Geographic Queries

In the geographic scenario, the queries are only selection and projection queries. In this case no intermediate results are generated to compute the answers. Therefore, with the use of clustering indices, the cost of the server is directly proportional to the number of messages and the size of the answers:

$$Cost_{server} = k_1 \cdot |\mathcal{M}| + k_2 \cdot size(\mathcal{M})$$

By incorporating  $k_1$  and  $k_2$  into the values of  $K_M$  and  $K_T$ , we can simplify the cost model to:  $Cost_{total} = K_M \cdot |\mathcal{M}| + K_T \cdot size(\mathcal{M}) + K_U \cdot U(Q, \mathcal{M})$ . In Section 7 we illustrate how values for the model parameters can be obtained in a particular scenario.

## 5.3 The 2-Query Merging Problem for Geographic Queries

The 2-Query Merging Problem, is the special case of the query merging problem when  $|Q| = 2$ . Our geographic query example is convenient for illustrating why the 2 query merge problem is simple, but why it is hard for more than 3 queries.

In the 2-Query Merging Problem we want to decide if it is worthwhile to merge two queries  $q_1$  and  $q_2$  into a merged query  $q_3$ . For compactness, in the following discussion, let us denote  $size(q_i)$  as  $S_i$ . Therefore, the cost of processing and transmitting queries  $q_1$  and  $q_2$  separately will be  $K_M + K_T \cdot S_1$  and  $K_M + K_T \cdot S_2$  respectively, for a total cost of  $2K_M + K_T(S_1 + S_2)$ . If we merge the queries into a single query  $q_3$ , the total cost will be  $K_M + K_T \cdot S_3 + K_U \cdot U(Q, \mathcal{M})$ , where  $U(Q, \mathcal{M}) = 2 \cdot S_3 - S_1 - S_2$ . We derive the  $U(Q, \mathcal{M})$  term in the following way: if we send a message with only  $ans(q_1)$ , the client receives an answer with size  $S_1$ . If we send a message with  $ans(q_3)$  instead, the client will receive an answer of size  $S_3$ . The difference  $(S_3 - S_1)$  is the size of the irrelevant results received by the client. We can use a similar derivation for the other client and conclude that the size of the irrelevant information for the other client is  $(S_3 - S_2)$ . Therefore the total size of irrelevant information is  $U(Q, \mathcal{M}) = 2 \cdot S_3 - S_1 - S_2$ .

From these expressions, it is easy to derive a decision rule that tells us exactly when it is beneficial to merge two queries (this is, if the second cost we computed is less than the first cost). Therefore, it is beneficial to merge  $q_1$  and  $q_2$  if  $K_M + K_T \cdot [S_1 + S_2 - S_3] + K_U \cdot [S_1 + S_2 - 2 \cdot S_3] > 0$ .

Unfortunately, the general problem ( $|Q| > 2$ ) is significantly harder, since there are many ways



to combine a set of queries into merged queries. For example, if we have three queries as input, it could be the case that it is not worthwhile merging any pair of them, but it is worthwhile merging the three queries into a single query. On the other hand, it could be the case that it is worthwhile merging one specific pair, but not the other pair and not the three queries. In conclusion, we would have to consider all possible ways to partition the input queries into subsets. For each possible partition we compute a cost, and then we pick the partition with minimum cost. This approach leads to an exponential algorithm. In fact, in the extended version of this paper [10] we show that the query merging problem is NP-hard.

Let us use our geographical database scenario to show a case when merging three queries is optimal, although merging any pair is not. In Figure 5, we show three queries over our geographical database.




Figure 5: 3-Query Merging Example

In the following discussion, we will assume that the answer of a query over each square unit in the diagram has size  $S$  and that we are using the bounding rectangle merging procedure. Therefore,  $size(q_1) = size(q_2) = 2S$ ,  $size(q_3) = S$ , and  $size(mrg(q_1, q_3)) = size(mrg(q_2, q_3)) = size(mrg(q_1, q_2)) = size(mrg(q_1, q_2, q_3)) = 4S$ . Since there are three queries, there are five ways to merge them: we can merge 2 of them (3 combinations), we can merge all of them, or we can keep them separately. By deriving the costs of all the five possible ways of merging the queries, we conclude that merging all of the queries is advantageous, although merging any pair is not, when all of the following equations are satisfied:

$$S > \frac{K_M}{4K_U} \quad \text{and} \quad S > \frac{K_M}{5K_U + K_T} \quad \text{and} \quad S < \frac{2K_M}{7K_U - K_T}. \quad (1)$$

These equations are satisfiable; for instance, if we pick  $S = 1$ ,  $K_M = 10$ ,  $K_T = 9$ , and  $K_U = 4$  all the equations will be true.

## 6 Algorithms for the Query Merging Problem

In this section we introduce heuristic algorithms for the query merging problem. However, we first briefly summarize two exhaustive algorithms that serve as reference points. We will compare the performance of all algorithms in Section 7.3.

## 6.1 Exhaustive Algorithm

An exhaustive algorithm for solving the query merging problem is presented in [10]. Unfortunately, exhaustive approaches have a doubly exponential complexity on the number of queries. This high complexity order makes exhaustive algorithms impractical for all but the smallest  $|Q|$ .

There exists a better algorithm for exhaustively solving the query merging problem when the cost model ensures the *single-allocation property*: each  $q_i$  in the solution is in one and only one element of  $\mathcal{M}$ . The single-allocation property means that if we want to process a set of queries  $\{q_1, q_2, q_3, q_4\}$ , we do not need to consider merged queries such as  $\mathcal{M} = \{\{q_1, q_2, q_3\}, \{q_1, q_4\}\}$  where a  $q_i$  (in this case  $q_1$ ) is in more than one element of  $\mathcal{M}$ . In the extended version of this paper, we present the Partition Algorithm that exploits this property and has a complexity of  $O(n^n)$ .

Although this may seem as a small improvement over the general exhaustive algorithm of the previous section; it significantly extends the values of  $|Q|$  for which we can use an exhaustive algorithm. For example, if we have a limit of 5 minutes, we could find the optimal solution for up to  $|Q| = 12$  using the Partition Algorithm, but only up to  $|Q| = 4$  using the Exhaustive Algorithm.

## 6.2 Pair Merging Algorithm

The Pair Merging Algorithm takes a greedy approach to solve the query merging problem. The foundations of this algorithm are two simplifying assumptions. First, we assume that the cost model has the single-allocation property (as defined in Section 6.1). Second, and more important, we assume that pair-wise decisions (i.e., deciding which pairs of queries to merge) will lead to the correct global solution. The second assumption is in general incorrect (as shown in Section 5.3). However, the assumption allows us to efficiently obtain solutions that, in practice, are very close to the real “optimal” solution (see Section 7.3). At the end of this section, we will show that the complexity of the algorithm is  $O(|Q|^2)$ .

The Pair Merging Algorithm maintains a set of sets of queries. Initially, each set contains each single query. Then for all pairs of sets, the algorithm computes the change in the total cost if each pair is merged. The pair that produces the largest positive decrease in cost is chosen and the sets are replaced by their union. The algorithm continues picking and merging sets until no merging of any pair decreases the total cost.

We evaluate the cost achieved by merging sets using our cost model. For instance, for the geographic cost model we are using as our running example (Section 5.2), we can use a generalization of the formula in Section 5.3 for solving the 2-query merging problem. In particular, it can be shown [10] that when merging two sets,  $M_a$  and  $M_b$ , containing the queries  $\{q_{a_1}, q_{a_2}, \dots, q_{a_p}\}$  and  $\{q_{b_1}, q_{b_2}, \dots, q_{b_r}\}$  respectively, the expression for solving the 2-query merging problem is:

$$Cost_{sep} - Cost_{merge} = K_M + K_T \cdot (size(mrg(M_a)) + size(mrg(M_b)) - size(mrg(M_a \cup M_b))) + K_U \cdot \{p \cdot size(mrg(M_a)) + r \cdot size(mrg(M_b)) - (p + r)size(mrg(M_a \cup M_b))\}.$$

Note, that we can obtain the expression for solving the 2-query merging problem given in Section 5.3 by making  $size(mrg(M_a)) = S_1$ ,  $size(mrg(M_b)) = S_2$ ,  $size(mrg(M_a \cup M_b)) = S_3$ , and the number of queries in each set equal to one ( $p = r = 1$ ).

The Pair Merging Algorithm, as presented, needs to compute the cost of doing all possible merges in every step. However, in each step, only two of the sets (the ones that we decide to merge) have changed. The other sets remain the same so we can use all the computation involving them in the next step. Specifically, in step  $k$  of the algorithm, there are  $(|Q| - k + 1) \cdot (|Q| - k)/2$  possible pairs; of those, only  $|Q| - k - 1$  are new pairs. The rest were all candidate pairs that were considered in the previous iteration. Note, that the fact that those candidates were not chosen in a previous iteration, does not preclude them to be chosen later (as long as they have a positive benefit). To avoid computing the costs again for those sets, in each step, we save all computed costs in a *profit table*. Before computing the cost of merging a set, we check in the profit table to see if the cost was already computed; if it was, we take it from the table; otherwise, we compute it and add it to the table. After selecting the pair of sets to be merged, we remove all entries of the profit table that are related with those sets. Using the profit table, the number of cost model evaluations is  $|Q|^2 + \sum_{i=1}^{|Q|-1} i - 1$ . Therefore, the complexity of the algorithm is  $O(|Q|^2)$ .

### 6.3 Directed Search Algorithm

The Pair Merging Algorithm works in only one direction, that is, it starts with a set  $\mathcal{M}$  where all the queries are single elements, and tries to merge those sets as much as possible. A potential weakness of this approach is that the Pair Merging Algorithm can be trapped into a local minimum of the cost function and miss the global minimum. This weakness is not unique of the Pair Merging Algorithm. Under a general cost function there are no polynomial algorithms that can avoid this weakness. However, in this section, we introduce the Directed Search Algorithm, a variation of the Pair Merging Algorithm that attempts to ameliorate this weakness.

The Directed Search Algorithm is based on two changes to the Pair Merging Algorithm. First, in addition to merging sets in  $\mathcal{M}$ , the Directed Search Algorithm may split one set in  $\mathcal{M}$  into two sets, one containing only one element and another with the remaining elements. The rationale behind this change is that splitting sets, allows the algorithm to “undo” a bad decision made earlier. We limit one of the sets to have only one element to reduce complexity. If we allow a more general splitting function, the splitting step of the algorithm becomes exponential. The second change is to use multiple initial states and choose the one that leads to the minimum cost. In this way, if one initial state leads to a local minimum from which we cannot escape, there is a good chance that a different initial state will avoid that minimum.

Choosing the set of initial states,  $\mathcal{J}$ , is critical for the performance of the algorithm. The best set of initial states would be one that allows the algorithm to explore as much as possible of the search space. In our experiments, we included the initial states where all the queries are separate ( $|\mathcal{M}| = |Q|$ ), the initial state where all queries are merged ( $|\mathcal{M}| = 1$ ), as well as random partitions of  $Q$ . The Directed Search Algorithm is presented in the extended version of this paper [10].

The worst case complexity of the algorithm is  $O(|\mathcal{J}||Q|^{|Q|})$ . This is because it is possible for the algorithm to explore the entire search space. Nevertheless, the algorithm is guaranteed to finish, as it only advances when there is a lower cost option. Although, the worst case performance is exponential, in our experiments this bound was never reached. Furthermore, in the experiments the algorithm showed a polynomial average running time.

## 6.4 Clustering Algorithm

The Clustering Algorithm takes a “divide and conquer” approach to the query merging problem. The foundation of the algorithm is the definition of a “distance” metric between queries. Basically, if the distance between two queries is “far enough,” we can ignore all combinations of merged queries that contain those queries. In this section, we will describe the algorithm, independently of the distance metric. In the following section, we will define the distance metric and introduce concrete examples.

A graphical intuition of the Clustering Algorithm is presented in Figure 6. In the figure there are five queries,  $q_1$  to  $q_5$ . Queries  $q_1$ ,  $q_3$ , and  $q_5$  are very close together, and therefore are good candidates for merging among themselves. However, queries  $q_2$  and  $q_4$  are far and it may not make sense to even consider merging with them. Specifically, the Clustering Algorithm works by computing the “distance” between each pair of queries and if this distance is below a certain threshold, it puts the two queries in the same cluster. In the figure, the algorithm may start by finding that the distance between the  $q_2$  and  $q_4$  is below the threshold and therefore should be in the same cluster (drawn as a dotted line). Then, the algorithm may find that  $q_1$  and  $q_3$  are also below the threshold and will put them in the same cluster. Then, the algorithm may find that the distance between  $q_5$  and  $q_3$  is also below the threshold, but because  $q_3$  already belongs to a cluster, the algorithm adds  $q_5$  to that cluster, instead of creating a new one. If  $q_5$  were close to several clusters (that is,  $q_5$  is close to at least a member of each of those clusters), a single cluster would be formed containing  $q_5$  and all of those clusters.

After this, the algorithm cannot find more pair of queries with distances belong the threshold and the clustering phase ends. In the following phase, we need to apply within each cluster any of the algorithms previously studied to merge its queries.




Figure 6: Clustering Algorithm Scenario

As the Clustering Algorithm considers pairs of queries, its performance is the greater of  $O(|Q|^2)$  and  $O(mrg_{alg}(|Q'|))$ , where  $mrg_{alg}$  is the performance of the merging algorithm used in the second phase and  $Q'$  is the cluster with the maximum number of queries. Obviously, the clustering algorithm works well only if there are clusters in the data. There are two cases for non-clustered data: queries will either be very close together or they will be very far apart. In the first case, the algorithm will identify just a single cluster that contains all the queries (and therefore will not improve the running time of the second phase). In the second case, the algorithm will indicate that there are not opportunities for merging queries, and the second phase need not be run.

In the next section, we will define the distance metric. Depending on this metric, the Clustering Algorithm behaves as an exact or as a heuristic algorithm. Due to limited space, we will only present a heuristic distance metric. In the extended version of this paper [10], an exact metric is presented.

### 6.4.1 A Heuristic Distance Metric

Given queries  $q_1$  and  $q_2$ , we want to determine if it will ever be advantageous to place them in the same partition, even if  $q_1$  or  $q_2$  have already been merged with other queries. To check, we can compute the maximum benefit that may be obtained by combining  $q_1$  (or a merged query containing it) and  $q_2$  (or a merged query containing it). In doing so, we will assume that costs not directly related to  $q_1$  and  $q_2$  are not affected by the decision (this may be false). If the maximum benefit is non-negative, we place  $q_1$  and  $q_2$  in the same cluster; otherwise, we leave them separate (though they may be eventually be combined due to another query that is close to both). This gives us a heuristic rule, but as we will see, our experimental results show it performs very well.

To illustrate, consider the cost model for geographic queries. To obtain the maximum possible benefit, we consider three components:

- $K_M \cdot |\mathcal{M}|$ : By merging two queries together, we reduce the size of  $|\mathcal{M}|$  by one. Note that we are ignoring the possible effects on  $\mathcal{M}$  after merging the two queries (which may reduce  $|\mathcal{M}|$  even more).
- $K_T \cdot size(\mathcal{M})$ : In the best case (this is when the result of  $q_1$  is contained in the result  $q_2$ ), the benefit  $size(\mathcal{M})$  will be at most  $min(size\{q_1\}, size\{q_2\})$ .
- $K_U \cdot U(Q, \mathcal{M})$ : In the best case,  $U(Q, \mathcal{M})$ , the benefit will be reduced by  $2 \cdot size(mrg(\{q_1, q_2\})) - size(\{q_1\}) - size(\{q_2\})$ .

Therefore, we should leave queries in separate clusters when:

$$K_M + K_T \cdot min(size\{q_1\}, size\{q_2\}) - K_U \cdot 2 \cdot size(mrg(\{q_1, q_2\})) - size(\{q_1\}) - size(\{q_2\}) < 0.$$

## 7 Performance Evaluation

In order to test the efficiency of the algorithms developed, a simulator has been implemented for geographic queries. It simulates an environment in which the queries are given on a two-dimensional database (see Figure 7). The database elements consist of two search attributes and the queries received by the server are range queries. The simulator consists of three main modules. The first module provides input to the simulator. The user specifies certain parameters like the dimensions of region covered in database, the size and number of queries, the cost parameters ( $K_M$ ,  $K_T$ ,  $K_U$ ) and the merge algorithm used. Given these parameters, this module generates a set of queries which is used as an input for the algorithms. The second module runs one of the algorithms described previously (i.e., Pair Merging, Directed Search, or Clustering algorithms). Finally, the last module evaluates the savings of the heuristic algorithms and quantifies their deviation from the optimal solution.

### 7.1 Generating Input

In most environments, it is quite likely that the input given by clients generates a pattern which creates groups of queries that are located near each other. Some portions of the database are likely

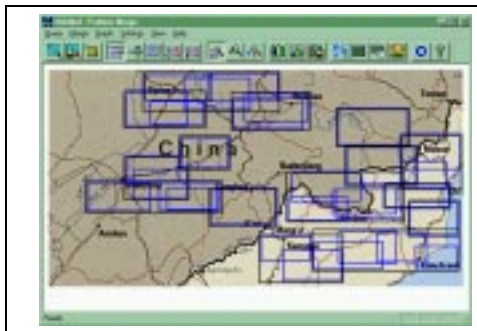


Figure 7: Screen Shot of the Simulator

to be accessed more frequently than others. For instance if the database specifies atmospheric conditions, the portions of the database denoting the regions with higher population are more likely to get queries. Similarly, in a battlefield situation, the number of queries in regions which have more combat troops will be much higher. Therefore, rather than generating random input queries, a clustering effect has been added to the input generating module, in which some queries tend to create small clusters.

Queries are generated in two ways; randomly and using clustering. The parameter  $cf$  is the ratio of the number of queries generated using clustering to the number of total queries. So  $cf|Q|$  gives the number of queries generated using clustering and  $(1 - cf)|Q|$  gives the number of queries generated randomly. The parameter  $sf$  is the ratio of number of queries in a cluster to the number of queries generated using clustering. We can compute the number of queries in a cluster and the number of clusters by  $sf \cdot cf \cdot |Q|$  and  $1/sf$  respectively. For each cluster, a cluster origin is generated randomly in the database. The distances of the queries in a cluster to the origin is computed using a normal distribution  $N(0, df^2)$  where  $df$  is the density of the clusters. The direction of each query from the cluster origin is generated randomly (i.e., a random value between  $0^\circ$  and  $360^\circ$ ). Another input parameter gives the size of the queries. The minimum and maximum ranges for both attributes are given. The size of each query is selected randomly in these ranges.

## 7.2 Cost Parameters

To select values for our cost parameters  $K_M$ ,  $K_T$  and  $K_U$ , we need to consider the specifics of the system we are studying, e.g., how costly it is to transmit data, relative to the other costs. The parameters impact not just the solution obtained, but how hard it is to find the optimal solution. For instance, there are some values for which it is trivial to find the optimal solution (consider  $K_M = 1$ ,  $K_T = 0$ ,  $K_U = 0$ ). For other values, algorithms such as Pair Merging may not find the optimal solution.

In this subsection we briefly illustrate how these parameters can be estimated in a given scenario. Please keep in mind that this is simply an illustration. We measure costs in dollars, since this makes it easier to compare processing and network costs.

At the server, we can first estimate the dollar cost of one second of processing as follows. We estimate that the server and its database system cost \$100,000, and this system will be amortized

for 2 years, giving a cost of \$50,000 per year. In addition, operating costs are \$50,000 per year, say. Thus, the cost per server second is

$$C_S = (50,000[\text{dollar/year}] + 50,000[\text{dollar/year}])/31,536,000[\text{sec/year}]$$

$$C_S = 0.003171[\text{dollar/sec}]$$

(In this and the expressions that follow we show the units in square brackets.) We test the server with queries that yield no answer (null queries), and discover that the server can process 1 query per second. We then test queries with different result sizes and discover that each additional answer object adds 1/100 second to the query time. Thus, we estimate the dollar cost at the server as

$$Cost_{server} = C_S[\text{dollars/sec}] * (1[\text{sec/query}] * |\mathcal{M}|[\text{query}] + 0.01[\text{sec/object}] * size(\mathcal{M})[\text{object}])$$

For the network, we compute the dollar cost per megabyte transmitted using numbers for a DSL service provided by Stanford. A DSL modem including installation cost is \$1165, or \$48.50 per month if we amortize over 2 years. The monthly fee is \$235, and the maximum bandwidth is 1.1Mbs. This gives us a dollar cost per MB of

$$C_N = (48.50[\text{dollar/month}] + 235[\text{dollar/month}])/((1.1\text{Mbs}/8[\text{b/B}]) * 2592000[\text{sec/month}])$$

$$C_N = 0.0007956[\text{dollar/MB}]$$

Using the results of [8], we estimate that setting up the connection for each query answer consumes 100KB. We also estimate that each answer object is 1KB in size. Hence, the network cost is

$$Cost_{network} = C_N[\text{dollar/MB}] * (0.1[\text{MB/query}] * |\mathcal{M}|[\text{query}] + 0.001[\text{MB/object}] * size(\mathcal{M})[\text{object}])$$

For computing client costs, we assume the client is a handheld device, and that the major cost is due to battery use. Using the PalmPilot as an example, a daily usage of 30 minutes results in a battery lifetime of 2 months. In other words, we get 1800 minutes of processing per battery change. We estimate the cost of batteries (including labor) at \$5. Thus, the cost per second is

$$C_C = 5[\text{dollars/replacement}]/(1800[\text{min/replacement}] * 60[\text{sec/min}])$$

$$C_C = 0.0000463[\text{dollar/sec}]$$

By performing experiments, we estimate that processing each answer object, whether useful or not, takes 0.1 second of processing on the client device. Thus, the client cost is

$$Cost_{client} = C_C[\text{dollar/sec}] * 0.1[\text{sec/object}] * (\sum_{q_i \in Q} size(q_i)[\text{object}] + U(Q, \mathcal{M})[\text{object}])$$

Combining the constants we have estimated, we obtain the following values for the overall cost model:  $K_M = 0.003251$ ,  $K_T = 0.000032$ ,  $K_U = 0.00000463$ .

As we have stated, our main objective in this section has been to illustrate the process by which one can estimate these constants. The performance experiments that must be performed to estimate costs, and the actual values obtained, can of course vary, but in the end, one can obtain cost proportionality constants that make it possible to compare the costs incurred at each stage of the multicast process.

### 7.3 Experiments

In this section we study the performance of the Pair Merging, Directed Search, and Clustering algorithms. Sample geographical queries were generated and both the exhaustive and heuristic algorithms were used to evaluate the results. Obviously the exhaustive algorithms give the optimal solution for a given set of queries. In order to evaluate the efficiency of our algorithms, we wish to address the following questions:

- What is the probability that the heuristic algorithms find the optimal solution?
- If the algorithms do not find the optimal solution, how far are the solutions to the optimal ones?
- In what scenarios does query merging pay off, and what are the potential gains?

Since we want to focus on how well our algorithms perform, rather than on predicting performance of a particular system, we select a scenario where it is particularly hard to find the optimal solution, and where we will stress the algorithms. This scenario was obtained by running the simulator over many different sets of parameters, and selecting the values (for  $cf$ ,  $sf$ ,  $df$ ,  $K_M$ ,  $K_T$ ,  $K_U$ ) where the heuristic algorithms found solutions further from the optimal. Thus, the results we will present here are *pessimistic* for the heuristic algorithms. As we will see, the results are rather good for the high-stress scenario, so this means the algorithms will perform even better in almost any other scenario.

Cost Parameters		Query Generating Parameters		Other Parameters	
Parameter	value	Parameter	value	Parameter	value
	$\times 10^{-5}$	$cf$	0.80	Database Size	
$K_M$	365	$sf$	0.30	$ Q  \leq 12$	100x100
$K_T$	3.25	$df$	100	$ Q  > 12$	400x400
$K_U$	0.16	$ Q $	100	Sample Size	10000
		Maximum Query size	40x40	$ J $	50
		Minimum Query size	20x20		

Figure 8: Base values

The input parameters and their base values are given in Figure 8. The values used for  $K_M$ ,  $K_T$ ,  $K_U$  are close to those illustrated in Section 7.2, but adjusted slightly to stress the algorithms more. Incidentally, note that only the ratio between  $K_M$ ,  $K_T$ ,  $K_U$  matters when comparing algorithms. If we multiply each of these values with a constant factor, the overall cost would change but the goodness of the solutions would not change. In our experiments, we use two different database sizes depending on the number of queries. When having a small number of queries, we used a smaller database size, so queries were not so dispersed that merging was never advantageous. The *Sample Size* specifies the number of times the simulator was run to generate the results shown in the graphs.

In Figure 9, we can see the sensitivity of the Pair Merging Algorithm to the cost parameters  $K_M$ ,  $K_T$  and  $K_U$ . The y-axis gives us the resulting number of merged queries ( $|\mathcal{M}|$ ). The parameter values on the x-axis are normalized. For instance, in the central graph in Figure 9, the number 1.5



on the x-axis indicates that  $K_T$  was chosen to be 1.5 times the base value, i.e,  $1.5 \cdot 0.0000325 = 0.00004875$ . As we can see from these graphs, the algorithm seems to be most sensitive to  $K_M$  since the overall cost is directly related to  $|\mathcal{M}|$ . We can also observe that for high values of  $K_T$ , a change in this parameter does not have a great influence on the algorithm since the network cost dominates.




Figure 9: Sensitivity Analysis

In Figure 10, the performance of the Pair Merging and the Directed Search algorithms are compared. This graph gives the fraction of the runs where the Directed Search Algorithm performs better than the Pair Merging Algorithm. The x-axis indicates the number of initial states for the Directed Search Algorithm. For instance, if we use 50 initial states, Directed Search finds a better solution than Pair Merging in about 20% of the cases, and in the remaining 80% of the cases both find the same solution. Note that Pair Merging never finds a better solution since it considers a subset of the merge configurations considered by Directed Search. Thus, the figure quantifies how much better Directed Search becomes as we increase the number of initial states ( $|\mathcal{J}|$ ).

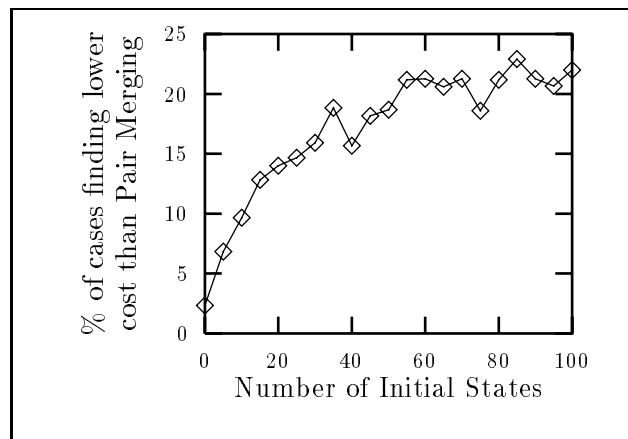


Figure 10: Performance of Directed Search Algorithm Compared to Pair Merging Algorithm




Figure 11: Percentage of Cases Finding the Optimal Solution

Figure 11 shows the fraction of runs where the Pair Merging and Directed Search algorithms find the optimal solution (as found by the Partition Algorithm). We only consider 12 queries or less as the Partition Algorithm must evaluate as many combinations as the  $|Q|$ th Bell number. For  $|Q| = 12$ , this is 4,213,597 combinations, and beyond that it grows to unmanageable sizes. We

omitted the trivial case  $|Q| = 2$ , as both algorithms are guaranteed to find the best solution.

As we can see, the chances of reaching the optimal solution decreases as the number of queries increase in both algorithms. The Directed Search Algorithm is more likely to reach the optimal solution. Extrapolating the curves, the results seem to imply that for large numbers of queries, it will be very unlikely that the optimal solution is found with heuristic algorithms. This is bad news, except that our next graphs will show that the deviation from optimal is very small.

To compute the deviation from optimum, let us say that the cost of disseminating a given set of queries without any merging is  $Cost_{initial}$ . Let  $Cost_{optimal}$  be the optimal cost obtained by an exhaustive algorithm, and let  $Cost_{heuristic}$  the cost reached by a heuristic algorithm. We measure the distance of the heuristic solution to the optimal solution as follows:

$$Distance = 100 \cdot \frac{Cost_{heuristic} - Cost_{optimal}}{Cost_{initial} - Cost_{optimal}}$$

This formula gives the deviation from optimum, relative to the maximum costs that may be saved through merging. For instance, a value of 0.0% indicates the solution is optimal, and a value of 100.0% indicates the cost is the same as with no merging.

Figure 12 shows the distance of the solutions. As expected, the distance for Pair Merging and Directed Search increase as we increase the number of queries. The Directed Search Algorithm has distances equal to zero for  $|Q| \leq 7$  because it is acting almost as an exhaustive algorithm. Recall that, in our experiments, we used 50 initial states for the Directed Search Algorithm. The total number of cases for  $|Q| = 3$  to 7 ranges between 5 and 877; thus, there is a high probability that Directed Search will be able to search the whole space. As the number of queries increase, the Directed Search Algorithm is no longer able to do an exhaustive search and its distance increases.

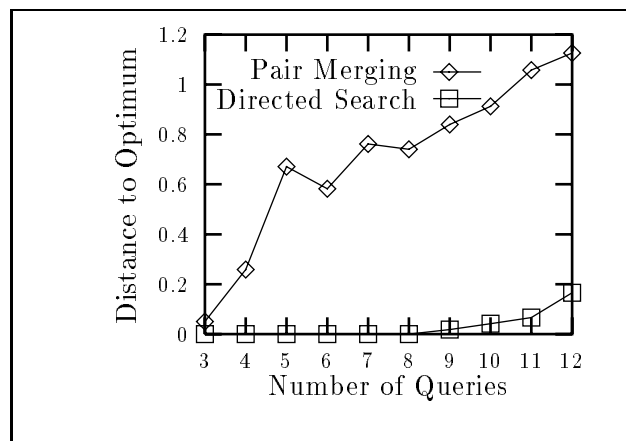


Figure 12: Distance to Optimum of the Pair Merging and Directed Search Algorithms

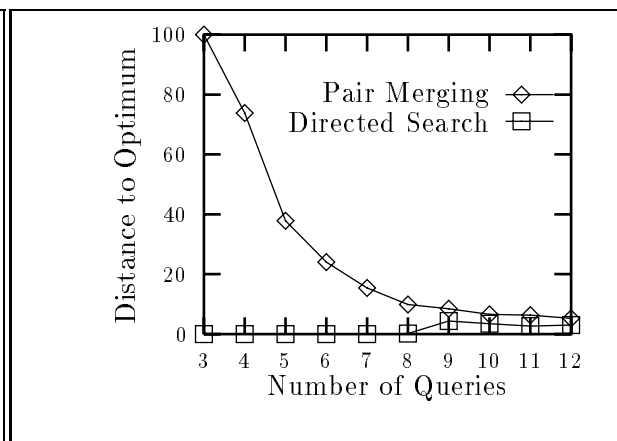


Figure 13: Distance of Suboptimal solutions to Optimum

Although we would need more data points to confirm this, we hypothesized that the the curves in Figure 12 have a logarithmic growth. So we expect that the heuristic algorithms will continue to have small distances even for relatively large number of queries. Furthermore, recall that we

selected our base parameters to stress our algorithm, so they are likely to perform much better in most other scenarios. Each data point in Figure 12 gives the distance averaged over all 10,000 runs.

In Figure 13, on the other hand, we show the average distance for those runs where the optimal solution was *not* found. (Note the change of scale.) The distance for Directed Search *increases* a bit around 9 queries, but this is simply because at this point the algorithm stops being close to exhaustive. Beyond 9 queries, the distance for Directed Search should start decreasing. To see this, recall that the distance of Directed Search will always be lower or equal to that of Pair Merging. Since Pair Merging shows a clear decreasing trend, Directed Search must also be decreasing. This is good news, for it predicts that distances (errors) will be small for large numbers of queries. As we increase the number of queries, the search space becomes huge, but the number of solutions that are very close to the optimal one also grows rapidly, so the heuristic algorithms have an excellent chance of finding a very good solution. Again, recall that our base scenario is one where it is hard to find the optimal solution, so in many other cases, distances will be even smaller.

In Figure 14 we can see the total cost after merging obtained by the Pair Merging, Directed Search and the Clustering with Heuristic Distance Metric algorithms. Since the costs given by the Pair Merging and Directed Search algorithms are very close, we showed them as a single line. To compute the cost of the Clustering Algorithm, the Pair Merging Algorithm was run on each of the clusters generated. The *no merge cost* is the cost of processing the queries without any merging. If our cost parameters are based on dollars (see Section 7.2), then the costs in Figure 14 are in dollars. The actual costs shown are small (a few dollars), but keep in mind that this is for a small number of queries. As the number of queries grows, so will the savings introduced by query merging. Furthermore, if the multicast is repeated say every hour, then the savings will be multiplied by  $24 \times 365$  in a year, and we can see that the savings can be significant.

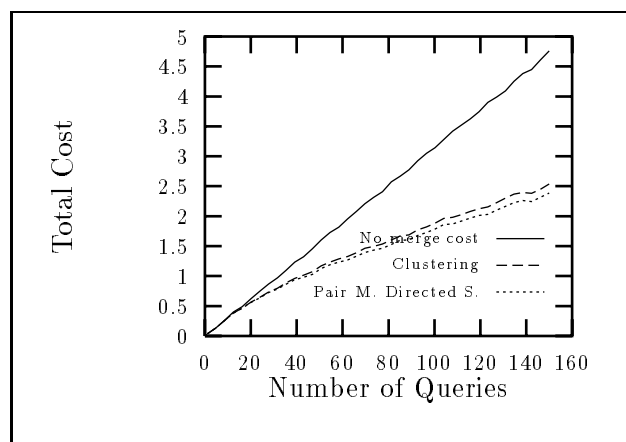


Figure 14: Final Total Cost

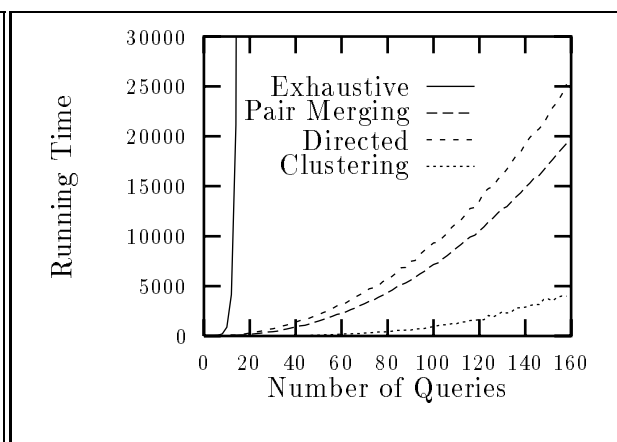


Figure 15: Running Time of Algorithms

The Clustering Algorithm does not yield better solutions, but we expect it have a smaller running time, since the merging algorithm runs on clusters with smaller number of queries. This is quantified in Figure 15, which shows the running times of the algorithms, measured by the number of merge/split operations. Recall that the complexity of the exhaustive algorithm is  $O(|Q|^{|Q|})$ ,

whereas the complexities of the Pair Merging and Directed Search algorithms is  $O(|Q|^2)$ . As we can see, the Clustering Algorithm is much faster as the number of queries increases, while still finding solutions that are close to those of the other algorithms. Thus, the Clustering Algorithm seems to be a good choice for scenarios with many queries.

In our final experiment we attempt to quantify when query merging pays off and by how much. Clearly, the most critical factor is the amount of overlap between submitted queries. If the queries are mostly disjoint, there will be little advantage to merging; if there is significant overlap, we expect significant gains. To get a sense for these gains, we can vary parameter  $cf$ , which controls the fraction of the queries that exhibit clustering. When  $cf = 1$  all queries are clustered, and when  $cf = 0$ , all queries are independent. Figure 16 shows the *total* costs incurred in processing the queries when merging is used and when it is not used, as a function of  $cf$ . Notice that even when queries are independent, merging introduces savings because there is still some random overlap among queries. As  $cf$  increases, the overlap increases, and the savings grow. Again, keep in mind that the savings of 1 to 2 dollars seen, will increase if there are more queries (100 were used for this experiment), and the savings will occur *every time* results are multicast.

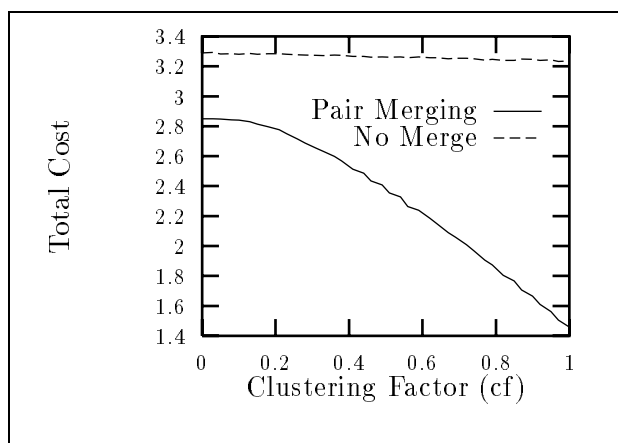


Figure 16: Total Cost Before and After Pair Merging

## 8 Related Work

The Data Dissemination Problem has been studied by a number of projects [18, 6, 21, 13, 1, 17, 12, 5]. However, none of them attempt to reduce costs by automatically merging similar queries.

The Query Merging Problem is also related to Client-side caching in client/server configurations [22]. In this approach, data is loaded into each client cache as answers to other queries are broadcast by the server. When a client is ready to make a query, it first checks in its own cache to see if the cache already contains the answer. The difference with our work, is that in the client-caching approach, queries are not known, so the server cannot optimize the global cost.

There are a number of data dissemination products and services in the market [7, 30, 3, 28]. However, as far as we know they do not attempt to do any real query merging. Most of these products are very simple, requiring clients to maintain their subscriptions and to “pull” from the server any new information. Servers normally unicast the results to each client, making this

approach non-scalable and resulting in a very high cost.

The Cellular Telephony and Telecommunication research community has also consider the problem of improving the bandwidth use on broadcast channels [20][9][24]. The difference between this effort and our work is the level of abstraction. While the telephony community focus on random memory page requests (and therefore, there is little information available to the optimizer), our work focus on queries and query answers which allows us to have more sophisticated schemes.

The BADD problem [23, 14] has generated a wealth of research in the data dissemination arena. References [4] and [26] have proposed multicast protocol, that can be used as a low level support to our algorithms. Deployment of Internet services through a satellite broadcast channel has been studied in [25] and “smart information push” by [15]. Reference [32] extends the client-side caching by considering caches not only at the client, but also at intermediate locations “close” to the clients. Finally, in [33] the data staging problem is described and heuristics to solve it are presented.

The query merging problem in a geographical database is closely related to the polygon covering problem [11], and to the set covering problem [31]. However, the special characteristics of the Query Merging Problem make it difficult to directly use the well known solutions to those problems.

## 9 Conclusions

In this paper we have studied the Query Merging Problem. We presented a very general framework and cost model for evaluating merging, and we presented a variety of merging algorithms. To illustrate and experimentally evaluate performance, we considered geographic queries as a representative example. Our results show that dissemination costs can be significantly decreased by applying a merging algorithm, and that heuristic algorithms work well.

Choosing which algorithm to use depends on the number of query subscriptions, the time available, and the precision required. If we have a small number of queries, we can use the Partition Algorithm (the practical limit is twelve queries when running in a typical workstation). If the running time is critical (e.g., in a scenario where queries and subscriptions change dynamically and hence the merge sets must be recomputed on the fly), then using the Clustering Algorithm before applying any merging algorithm improves the running time significantly. If finding the best solution is important and the number of queries is large, the Directed Search Algorithm should be used. When using the Directed Search Algorithm, the number of initial states can be chosen according to the running time limitations.

In this paper we presented a general dissemination model in which only broadcast was used. In the extended version of this paper [10] we consider multiple physical channels. We are also working on other extensions, such as subscriptions with different frequencies, and handling incremental changes to the set of queries.

## References

- [1] S. Acharya, R. Alonso, M. Franklin, and S. Zdonik. Broadcast disks: data management for asymmetric communication environments. In *ACM SIGMOD*. ACM, May 1995.
- [2] Airmedia. At <http://www.airmedia.com/>.
- [3] BackWeb. At <http://www.backweb.com/>.

- [4] F. Bauer and M. Brown. A reliable multicast transport protocol for a global broadcast service-based network. In *MILCOM 97*, volume 2, pages 988–92. IEEE, 1997.
- [5] A. Bestavros. Speculative data dissemination and service to reduce server load, network traffic and service time for distributed information systems. In *Proceedings of ICDE'96*, March 1996.
- [6] T. Bowen, G. Gopal, G. Herman, T. Hickey, K. Lee, J. Raitz, and A. Weinrib. The datacycle architecture. *CACM*, 35(12), December 1992.
- [7] A. Chan. Transactional publish/subscribe: the proactive multicast of database changes. In *ACM SIGMOD*, June 1998.
- [8] J. C.-I. Chuang and M. A. Sirbu. Pricing multicast communication: a cost based approach. Technical Report 98354, Carnegie Mellon University, 1998.
- [9] M. Corson and J. Macker. Global quality of service-based and reliable data dissemination via asymmetric direct broadcast satellite channels. In *19th Pacific Telecommunications Conference*, Jan 1997.
- [10] A. Crespo, O. Buyukkokten, and H. Garcia-Molina. Efficient query subscription processing in a multicast environment. Technical report, Stanford University, 1999.
- [11] J. C. Culberson and R. A. Reckhow. Covering polygons is hard (preliminary abstract). In *29th Annual Symposium on Foundations of Computer Science*. IEEE, oct 1988.
- [12] S. Dao and B. Perry. Efficient dissemination of information on the internet. *Bulletin of the Technical Committee on Data Engineering*, 19(3):48–54, September 1996.
- [13] S. Dao and B. Perry. Information dissemination in hybrid satellite/terrestrial networks. *Bulletin of the Technical Committee on Data Engineering*, 19(3):12–18, September 1996.
- [14] R. Douglass, J. Mork, and B. Suresh. Battlefield awareness and data dissemination (badd) for the warfighter. In *Proceedings of the SPIE*, volume 3080, pages 18–24, 1997.
- [15] J. Dukes-Schlossberg, Y. Lee, and N. Lehrer. Iids: intelligent information dissemination server. In *MILCOM 97 Proceedings*, volume 2, pages 635–9. IEEE, 1997.
- [16] M. Franklin and S. Zdonik. Dissemination-based information systems. *Bulletin of the Technical Committee on Data Engineering*, 19(3):20–30, September 1996.
- [17] M. Franklin and S. Zdonik. A framework for scalable dissemination-based systems. In *OOPSLA '97. Object Oriented Programming Systems Languages and Applications*. ACM, Oct 1997.
- [18] D. Gifford. Polychannel systems for mass digital communication. *CACM*, February 1990.
- [19] J. Graham-Cumming. Hits and miss-es: A year watching the web. In *Sixth International World Wide Web Conference*, 1997.
- [20] Q. Hu, D. L. Lee, and W.-C. Lee. Optimal channel allocation for data dissemination in mobile computing environments. In *18th International Conference on Distributed Computing Systems*, May 1998.
- [21] T. Imielinski and B. Badrinath. Mobile wireless computing. *CACM*, 37(10):18–28, Oct 1994.
- [22] A. M. Keller and J. Basu. A predicate-based caching scheme for client-server database architectures. *The VLDB Journal*, 5(1), January 1996.
- [23] M. Lazaroff and P. Sage. Any information, anywhere, anytime for the warfighter. In *Proceedings of the SPIE*, volume 3080, pages 35–42, 1997.
- [24] W.-C. Lee, Q. Hu, and D. L. Lee. Channel allocation methods for data dissemination in mobile computing environments. In *Proceedings. The Sixth IEEE International Symposium on High Performance Distributed Computing*. IEEE, Aug 1997.
- [25] R. Lindell, J. Bannister, C. DeMatteis, M. O'Brien, J. Stepanek, M. Campbell, and F. Bauer. Deploying internet services over a direct broadcast satellite network: challenges and opportunities in the global broadcast service. In *MILCOM 97*. IEEE, 1997.

- [26] N. Lu. Network interface to tactical communications. In *MILCOM*. IEEE, 1997.
- [27] M. V. Mannino, P. Chu, and T. Sager. Statistical profile estimation in database systems. *ACM Computing Surveys*, 20(3), September 1988.
- [28] MARIMBA. *Castanet*. At <http://www.marimba.com>.
- [29] B. Nadel. A kinder, gentler pointcast. *PC Magazine*, 15(18), 1996.
- [30] S. Ramakrishnan and V. Dayal. The pointcast network. In *ACM SIGMOD*, June 1998.
- [31] H. Salkin and J. Saha. Set covering: algorithms, results and codes. In *Bulletin of the Operations Research Society of America*, volume 20, suppl.2, Nov 1972.
- [32] T. Stephenson, B. DeCleene, G. Speckert, and H. Voorhees. Badd phase ii. dds information management architecture. In *Proceedings of the SPIE*, volume 3080, pages 49–58, 1997.
- [33] M. Tan, M. D. Theys, H. J. Siegel, N. B. Beck, and M. Jurczyk. A mathematical model, heuristic, and simulation study for a basic data staging problem in a heterogeneous networking environment. In *Proceedings of the 7th International Computing Workshop (HCW'98)*. IEEE, 1998.