

Palo Alto Research Center

**Using Collaborative Filtering to Weave an
Information Tapestry**

**David Goldberg, David Nichols, Brian Oki, and Douglas
Terry**

XEROX

Using Collaborative Filtering to Weave an Information Tapestry

David Goldberg, David Nichols, Brian Oki, and Douglas Terry

CSL-92-10 September 1992 [P92-00156]

© Copyright 1992 Xerox Corporation. All rights reserved.

XEROX

Xerox Corporation
Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, California 94304

Using Collaborative Filtering to Weave an Information Tapestry

David Goldberg, David Nichols, Brian Oki, and Douglas Terry

CSL-92-10 September 1992

© Copyright 1992 Xerox Corporation. All rights reserved

Abstract: Tapestry is an experimental system that manages an incoming stream of electronic documents, including electronic mail, news wire stories and NetNews articles. In common with some recent mail systems, Tapestry uses filtering to cope with large volumes of incoming documents. Where Tapestry differs from these systems is in its philosophy that humans provide the most reliable evaluation of documents, and so it uses *collaborative filtering*, which it implements by having users annotate documents, and then allowing filtering using those annotations. Because annotations are not known at the time that documents arrive, Tapestry filters must not only test incoming documents, but must also run repeatedly over the entire database of documents. This paper is a report on the Tapestry design and the status of its implementation.

A version of this paper will appear in *Communications of the ACM*, December 1992.

CR Categories and Subject Descriptors: H.2.4 [Database Management]: Systems - *Query processing*; H.2.8 [Database Management]: Database Applications; H.3.4 [Information Storage and Retrieval]: Systems and Software - *Current awareness systems (selective dissemination of information-SDI)*.

Additional Keywords and Phrases: continuous query, information filtering, active database, query rewrite, bounding monotonic query, incremental query, query language

XEROX

Xerox Corporation
Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, California 94304

1.0 Introduction

Tapestry is an experimental mail system developed at the Xerox Palo Alto Research center. The motivation for *Tapestry* comes from the increasing use of electronic mail, which is resulting in users being inundated by a huge stream of incoming documents [2][7][12]. One way to handle large volumes of mail is to provide mailing lists, enabling users to subscribe only to those lists of interest to them. However, as illustrated in Figure 1, the set of documents of interest to a particular user rarely map neatly to existing lists. A better solution is for a user to specify a *filter* that scans all lists, selecting interesting documents no matter what list they are in. Several mail systems support filtering based on a document's contents [3][5][6][8]. A basic tenet of the *Tapestry* work is that more effective filtering can be done by involving humans in the filtering process.

In addition to content-based filtering, the *Tapestry* system was designed and built to support *collaborative filtering*. Collaborative filtering simply means that people collaborate to help each other perform filtering by recording their reactions to documents that they read, such as the fact that a document was particularly interesting (or particularly uninteresting). These reactions, more generally called *annotations*, can be accessed by others' filters. One application of annotations is in support of moderated newsgroups. Currently, moderated groups have a single moderator, who selects a subset of messages to be posted to the moderated group. With annotations, a group can have many moderators. To see the newsgroup as it would be moderated by (say) Smith, simply filter for those articles that Smith endorsed with an annotation.

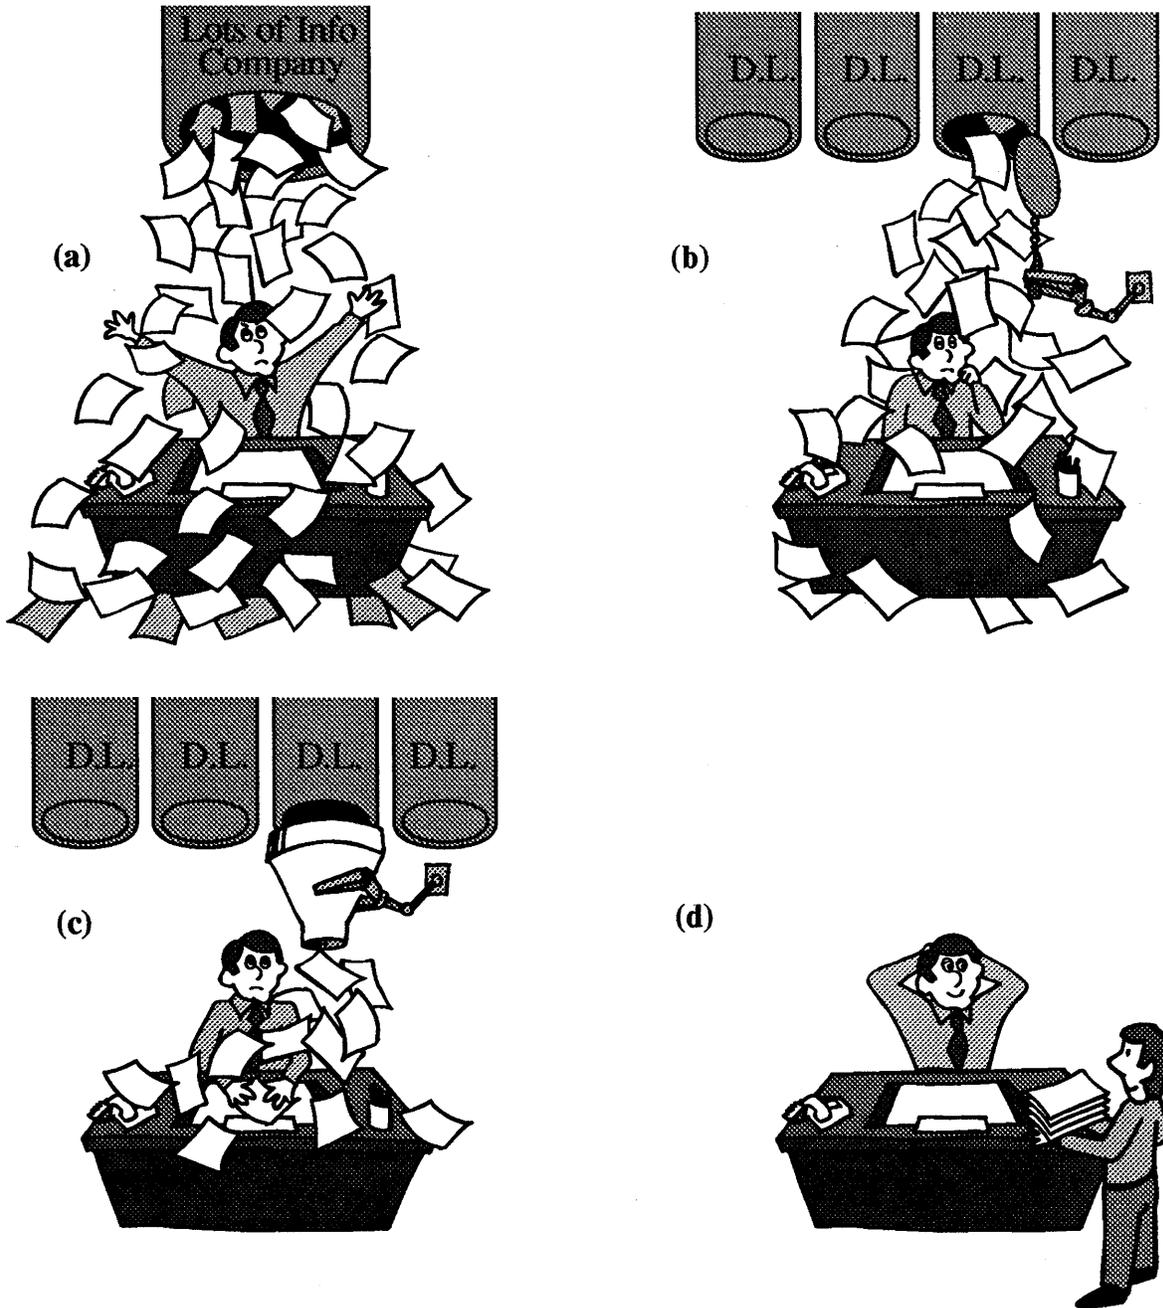
Implicit feedback from users, such as the fact that some user sent a reply to a document, can also be utilized in the filtering process. For example, suppose that you would like to receive "interesting" documents from the NetNews newsgroup comp.unix-wizards in the mail, but you don't know how to write a search expression that characterizes them, and you don't have time to read them all yourself. However, you do know that Smith, Jones and O'Brien read all of comp.unix-wizards, and reply to the more interesting documents. *Tapestry* allows you to filter on "documents replied to by Smith, Jones or O'Brien."

Collaborative filtering is novel because it involves the relationship between two or more documents, namely a message and its reply, or a document and its annotations. Unlike current filtering systems, *Tapestry* filters cannot be computed by simply examining a document when it arrives, but rather require (potentially) repeatedly issuing queries over the entire database of previously received documents. This is because sometime after a document arrives, a human (say Smith) may read that document and decide it is interesting. At the time he replies to it (or annotates it), you want your filter to trigger and send you the original document.

Tapestry is more than a mail system, because it is designed to handle any incoming stream of electronic documents. Electronic mail is only one example of such a stream: others are news wire stories and NetNews [10] articles. Moreover, *Tapestry* is not only a mechanism for filtering mail, it is also a repository of mail sent in the past. *Tapestry* unifies *ad hoc* queries over this repository with the filtering of incoming data.

A typical scenario of *Tapestry* system usage is as follows. A user decides she is interested in "mail filtering." To find documents on this topic, she issues an *ad hoc* query, perhaps by searching for the keyword "filtering". This returns too many documents. She eventually discovers that search-

Figure 1.
(a) electronic mail overload
(b) using distribution lists
(c) conventional filtering
(d) collaborative filtering



ing, either for documents containing both “information” and “filtering”, or for documents containing “filtering” that received at least three endorsements, works much better. Having tested this out, she now installs this search as a query filter, and from now on, all new documents satisfying this filter will be delivered to her mailbox.

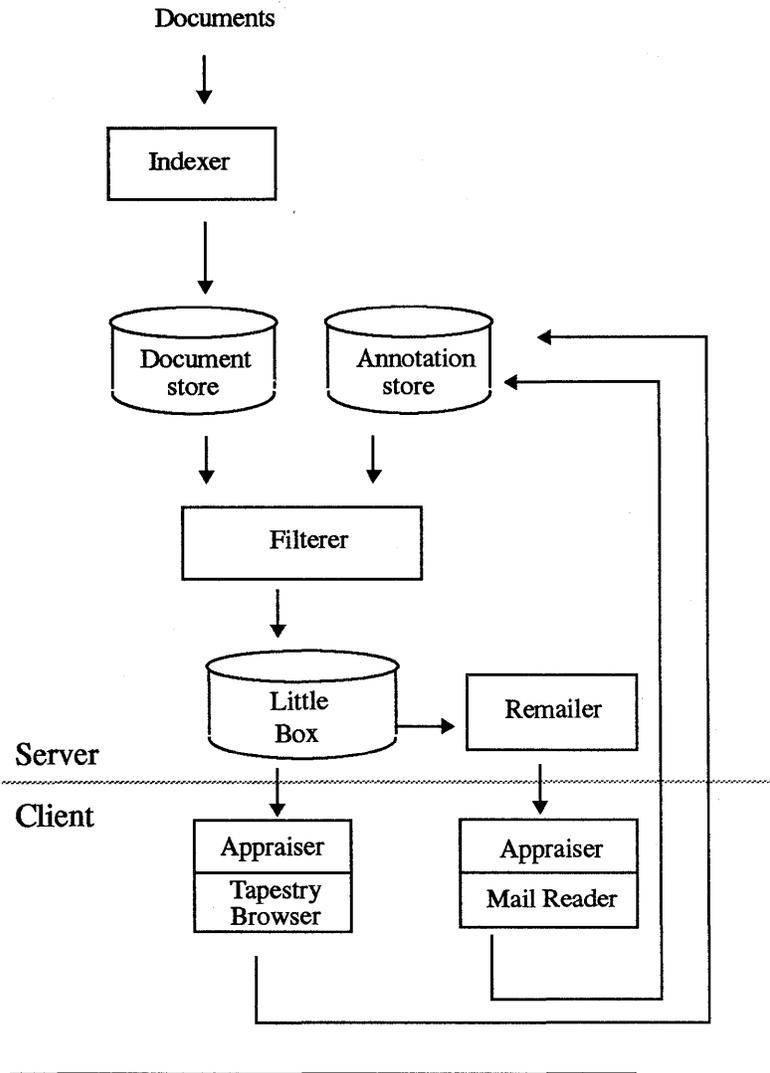
The rest of this paper discusses the design and implementation of Tapestry. Section 2.0 gives an overview of the Tapestry architecture; section 3.0 discusses TQL, the Tapestry query language used for specifying filters; and section 4.0 gives the details of the repeated execution of filter queries. Then section 5.0 details the existing Tapestry implementation.

2.0 Architecture

Figure 2 shows the flow of documents through the major architectural components of Tapestry. These components are:

- *Indexer*. Reads documents from external sources such as electronic mail, NetNews, or news-wires and adds them to the document store. The indexer is responsible for parsing documents into a set of indexed fields that can be referenced in queries.
- *Document store*. Provides long term storage for all Tapestry documents. Also maintains indices on the stored documents so that queries over the document database can be efficiently executed. The document store is append-only.
- *Annotation store*. Provides storage of annotations associated with documents. The annotation store is also append-only.
- *Filterer*. Repeatedly runs a batch of user-provided queries over the set of documents. Those documents matching a query are placed in the little box of the query’s owner.
- *Little box*. Queues up documents of interest to a particular user. Each user has a little box, where documents are deposited by the filterer and removed by a user’s document reader.
- *Remailer*. Periodically sends the contents of a user’s little box to the user via electronic mail. This is intended for users who wish to access Tapestry with their current mail reader.
- *Appraiser*. Applies personalized classification to a user’s documents, i.e. to those documents in the user’s little box. This function can automatically prioritize and categorize documents.
- *Reader/Browser*. Provides the user interface for accessing Tapestry services. This includes facilities for adding/deleting/editing filters, retrieving new documents, displaying documents, organizing documents into folders, supplying annotations, running *ad hoc* queries, etc.

Tapestry uses a client/server model. Two styles of interaction with the server are envisioned. The preferred mode of interaction is via a reader/browser which provides users with easy access to the full range of filtering and annotation functions. Users that do not want to, or are not able to, use the Tapestry browser can access Tapestry services from a conventional mail reader by having a remailer daemon send documents that match a user’s filters to the user via electronic mail. Users can also send mail documents to the Tapestry server to invoke any of its operations, such as adding filters, adding annotations, and even running *ad hoc* queries. The Tapestry architecture is flexible about the location of the client/server split. Figure 2 illustrates one possible division.

Figure 2. The flow of documents through Tapestry

Most of the Tapestry architecture follows naturally from the goal of providing collaborative filtering. For example, to support filters involving relationships between documents, there must be a document store. In order for users to communicate their ratings of documents, annotations are provided. The following subsections provide a more detailed rationale for some of the architectural components.

2.1 Document And Annotation Stores

Ideally, the Tapestry store will save documents forever. With the decreasing price of disk storage, this is becoming increasingly practical. As will be explained in section 4.0, documents as seen by

the filter query language must be *immutable*. This means that once a document arrives in the document store, it is never modified. Thus Tapestry documents can be conveniently stored using newer technologies such as WORM (write-once, read-many) disks.

Annotations are stored separately from documents, with links connecting each annotation to its associated document. It might seem more natural to combine documents and annotations into a single store, with the annotations to a document appended as additional fields. There are several reasons why this was not done. First, since annotations for a document arrive after the document itself, appending annotations as additional fields would violate the immutability requirement. Second, some annotations are themselves complex objects, and those annotations are more simply stored as separate records with pointers back to the document they annotate. The issue of complex annotations also arises in the Tapestry query language (see section 3.2).

2.2 Appraisers

Tapestry users want more than a binary sieve that can only accept or reject a document. For example, a user might want to assign priorities to messages, giving messages that announce meetings a higher priority than messages that announce promotions. And it would be handy to be able to change priorities. For example, the arrival of a message containing an update about a meeting (perhaps announcing a new meeting room) might cause the previous announcement to be given a lower priority, but probably not deleted, since it may contain details not repeated in the updating message.

To support classification of documents, Tapestry provides *appraiser* functions. Fitting appraisers into the overall architecture is not completely straightforward. At first it would seem simplest to run each user's appraiser on the server as documents arrive. However, this has a potentially serious drawback. Filtering on incoming documents is a very computationally intensive task. Imagine a Tapestry system with hundreds of users, each with dozens of filter queries, running on a document stream of tens of documents per minute. Running appraisers directly on the incoming document stream would put them on the critical performance path. To avoid this, the Tapestry architecture performs filtering in two steps. The first level of filtering is performed by filter queries, which are binary: they either accept or reject a document. The accepted documents for a user are then placed into that user's *little box*. The second level of filtering is done by appraiser functions that run only over the contents of the little box. Unlike the "big box" (the global Tapestry database), the little box will have few enough messages to allow them to be copied to the workstation. This allows the users's mail reading program or browser to provide more complex appraiser functions than could be supported in the server.

2.3 Browsers

The Tapestry architecture supports browsers that combine the functions of a mail reader and a traditional document browser. Corresponding to the role of mail reader, such a browser should supply "new mail" functionality. The server supports this by delivering the results of filter queries (new mail) to the little box, leaving it up to the client to remove the results. Browsers periodically run the appraiser over the documents in the little box, record their document identifiers, and then delete them from the little box. *Ad hoc* queries are another way to get documents into the browser.

Ad hoc queries are made to the server in the same query language as filter queries and may return documents that were not previously in the browser.

In traditional mail systems, each mail reader obtains and stores its own copy of each message. Thus messages sent to a large mailing list are stored many times. Since Tapestry provides an immutable document store, Tapestry browsers need only keep a document identifier (i.e. pointer). When a user deletes a message from the browser, the document still exists and can be recovered using an *ad hoc* query.

Users of a browser would like to be able to issue queries that involve both document fields and *private fields*. Private fields store information such as whether a document has been read yet, which folders it is in, etc. A browser can store private fields along with other document fields, making them easily available for *ad hoc* queries. However, since documents must appear immutable to filter queries, and private fields are mutable, private fields can only be referenced by *ad hoc* queries, not filter queries.

3.0 Tapestry Query Language (TQL)

A key part of Tapestry is filtering documents, with the filters specified as queries. Hence, choosing the language in which filter queries are written was one of the important design decisions. One obvious choice was to use SQL[1], the widely-used standard query language for relational databases. Adopting it as the Tapestry query language would have had the additional advantage of simplifying the implementation, because Tapestry is implemented on top of a commercial database which supports SQL.

We rejected using SQL as our query language for two reasons. First, there is a serious mismatch between the relational model and the Tapestry model of documents. The set of fields in a document is extensible, whereas SQL schemas have a fixed set of fields. Also, SQL does not directly support sets, whereas many document fields are set-valued. Examples are the “To:” field of mail messages, and the “Newsgroups:” field of netnews articles. Second, we wanted to make it easy for users to type in *ad hoc* filter queries, and we felt that the amount of boiler plate in SQL made that difficult.¹

Thus Tapestry has its own query language TQL (for *Tapestry Query Language*). The next two subsections describe TQL informally by the use of examples. Even though TQL is easier to use than SQL, we expect most users will not use TQL directly, but instead will issue queries from a browser using predefined (but possibly parametrized) queries.

3.1 Basic Examples

A TQL query is a boolean expression. It select those documents that satisfy the expression. The set of allowable TQL expressions are similar to statements in first order predicate calculus. They

1. This is not meant as a criticism of SQL. Tapestry filter queries are much more specialized than general SQL queries, which is why they can be written with less boiler plate.

Table 1. Common Fields and Their Types

to	set of strings
date	date
sender	string
cc	set of strings
subject	string
newsgroups	set of strings
in-reply-to	set of documents
words	set of strings

combine “atomic formulas” with boolean operators, and they can have free variables quantified by EXISTS or FORALL. However, unlike predicate calculus, TQL supports sets.

The simplest Tapestry queries are atomic formulas, which involve relational operators like = and < as well as the wildcard matching operator LIKE. An example is

```
m.subject = "Next Tapestry Meeting"
```

which selects exactly those documents (or *messages*) *m* whose subject field *m.subject* is “Next Tapestry Meeting.”

TQL queries reference the fields of documents using *m.field*, where *field* is the name of a document field. Each field has a type. Some common fields and their types are listed in Table 1. Most correspond to fields of mail messages and newsgroup articles. One exception is “words”, which is the set of all words occurring in the body of the document.

More complex TQL queries are built up by combining atomic formulas with boolean operators as in the following query:

```
(m.sender = "Smith" OR m.date < "April 15, 1991")
AND m.subject LIKE "%Tapestry%"
```

This query selects messages that were either from “Smith” or else sent before April 15, and whose subject field included the word “Tapestry”. As in SQL, % is a wild card symbol that matches any number of characters.

The major difference between TQL and predicate calculus is TQL’s support for sets. A simple example of a Tapestry query using set-valued fields is the atomic formula

```
m.to = {"Joe", "Tom"}
```

which matches documents whose *m.to* fields include “Joe” and “Tom” (and possibly others). Sets can involve operations other than =, such as the query

```
m.to = { "Joe", LIKE "%Bill%" }
```

which asks for an *m.to* field containing at least ‘Joe’ and a name containing ‘Bill’.

Quantified variables are needed for collaborative queries. An example is

```
EXISTS (m1: m1.sender = "Joe" AND m1.in-reply-to = {m})
```

which selects all documents *m* that Joe has replied to.

Finally, a user's filter queries can reference the queries of another user. For example, the TQL query

```
m IN Terry.Baseball AND m.words = { "Dodgers" }
```

returns all the messages selected by Terry's 'Baseball' query that contain the word 'Dodgers'.

3.2 Annotations

The design of TQL presented so far follows rather naturally once the decision is made to have the query language match the form of electronic documents such as mail messages and NetNews articles. It is not so straightforward to decide how to handle annotations. As explained in section 2.1, annotations are not stored as fields of the document they annotate. However, this does not preclude TQL treating them as additional document fields, and indeed this is the most natural representation for annotations such as priority. A notation like "m.a.priority" could be used to access the priority of a document, the 'a' serving to map out a separate name space for annotations. Similarly, the folders to which a document belongs could be a set valued field, "m.a.folders".

Things do not work smoothly for the more complex annotations used to support collaborative filtering. Consider trying to implement voting using additional document fields. If vote is to be an annotation field, then "m.a.vote" would have to be a set of votes, each of which has a structure of its own, such as who the voter was, and the value of his vote. So a query such as "messages voted for by weiser" would be expressed as something like "the set m.a.vote must have a member v with v.owner = weiser", and this would require extending the set notation of the previous section.

The way this query is written in TQL is

```
a.type = "vote" AND a.owner = "weiser" AND a.msg = m
```

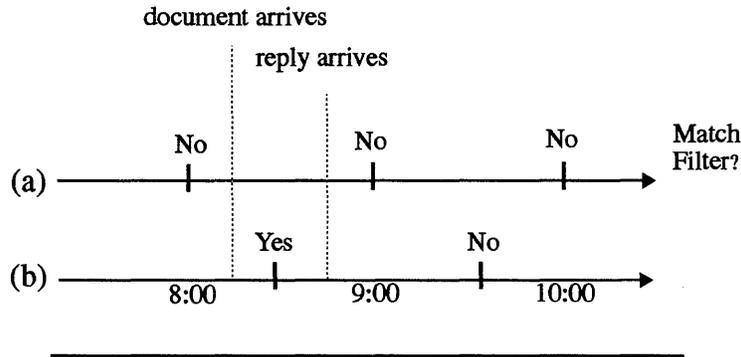
By introducing an annotation object, which always has a field msg that links it to a document, the kind of queries that support collaborative filters become simpler. We mentioned earlier that collaborative queries use EXISTS. The query above has an implicit EXISTS, and can also be written as

```
EXISTS (a: a.type = "vote" AND a.owner = "weiser" AND a.msg = m)
```

The cost introducing separate annotation objects is that simple queries like "documents of priority 10" become slightly more complex

```
a.type = "priority" AND a.value = 10 AND a.msg = m
```

Since one of the major design goals of Tapestry was to support collaborative filtering, we felt that the design with separate annotation objects was preferable.

Figure 3. Nondeterministic Behavior of Filters

4.0 Filter Queries

The heart of the Tapestry server is the Filterer, which executes users' filter queries. A straightforward method of implementing a filter query is to periodically execute it, say once every hour. This approach has the problem of returning all the old messages that matched the query the last time it ran, so something must be done to suppress these messages. Moreover, there is another more serious problem, namely that periodic execution can exhibit unpredictable behavior.

Consider the query: "select documents to which nobody has sent a reply." When a document is added to the database, it matches the query. However, once a reply document arrives, the document being replied to no longer matches the query. If a particular document were to arrive in the database at 8:15 and a reply to it arrived at 8:45, then the document would not be returned by a system that simply ran the filter query every hour on the hour [Figure 3(a)], but would be returned by a system that ran it every hour on the half hour (b), since the document would match at 8:30. This raises the general question: What are reasonable semantics for a filter query that executes repeatedly?" In other words: What guarantees can be provided to users about the set of documents returned by a filter query?

Users should not need to understand the implementation of the system in order to know what results to expect as the result of a filter query. The semantics should be independent of how the system operates internally and when it chooses to perform various operations such as executing queries. Two users with the same filter query should see the same result data. This implies that the semantics of filter queries should be time-independent.

4.1 Continuous Semantics

Tapestry gives filter queries *continuous semantics*, which is defined as follows:

The results of a filter query are the set of data that would be returned if the query were executed at every instant in time.

That is, the system guarantees to show the user any document that would be selected by the query at any time. The system may implement this behavior in any number of ways, such as collecting results and presenting them to the user periodically, but the actual set of results eventually seen by the user is well-defined and time-independent.

Rewriting the above in symbols, let $Q(t)$ denote the set of documents that would be returned by the execution of query Q over the database that existed at time t . That is, $Q(t)$ is the result of running Q at time t . When a query Q is executed with *continuous semantics*, it returns not $Q(t)$, but rather:

$$\bigcup_{s \leq t} Q(s)$$

Filter queries are qualitatively different from one-time queries. Consider the user who wants to see all the documents that do not receive replies. The obvious formulation: “select documents to which nobody has sent a reply,” when executed as a filter query, would return every document to the user, since every document has no replies when it first arrives. This is undoubtedly not what the user intended. The problem does not lie with continuous semantics, but rather with the user’s imprecise specification of his filter query. Finding the documents that *never* receive a reply would require waiting forever, but in practice a short wait will return a good approximation, since most messages are replied to quickly. Thus a more precise query would be something like: “select documents that are more than two weeks old and to which nobody has sent a reply.” This illustrates the point that some queries only make sense when executed on a one-time basis, and are not suitable as filter queries that are repeatedly executed.

4.2 Implementation

How can continuous semantics be realized in a practical system? Certainly, running a query at every instant in time is not possible, and if it were possible, would not be practical. The rest of this section gives an overview of techniques for providing continuous semantics in an effective and efficient manner. An earlier paper gives full details of how this is done [13].

The key to providing efficient continuous semantics is the following observation: Given a query whose result set is non-decreasing over time, the simple technique of periodically executing the query and returning the new results yields continuous semantics. Such a query is said to be *monotone*. The frequency with which a monotone query is executed simply affects the size of each batch of results, not the collective set of results.

Tapestry implements filter queries with continuous semantics in two stages. First a query is rewritten as a monotone query that returns at least all documents that currently match the original query or else matched it at some time in the past. If the rewritten query is Q , and Tapestry has pre-

Figure 4. Filter Query Execution

```

Set  $\tau = -\infty$ 
FOREVER DO
  set  $t :=$  current time
  Execute query  $Q^I(\tau, t)$ 
  Return result to user
  set  $\tau := t$ 
  Sleep for some period of time
ENDLOOP

```

viously evaluated Q at time τ , then at time t Tapestry can implement continuous semantics by returning $Q(t) - Q(\tau)$ to the user, where “-” stands for set difference.

In general, the sets $Q(t)$ and $Q(\tau)$ are almost the same, and contain mostly documents that have already been returned to the user. Computing $Q(t) - Q(\tau)$ is very inefficient, since $Q(t)$ and $Q(\tau)$ both return large sets, but then most of these documents ‘cancel’ when $Q(t) - Q(\tau)$ is computed. So Tapestry has a second stage, in which the monotone query Q is rewritten as an *incremental query*, $Q^I(\tau, t)$, that can quickly compute an approximation to $Q(t) - Q(\tau)$.

To summarize the discussion so far, when a filter query is submitted to Tapestry, it is first rewritten to a monotone query Q , and then Q is further rewritten to an incremental Q^I . This incremental query is what is used by the Tapestry filterer. The filterer repeatedly runs the incremental query, queues up the selected documents for delivery to users, records the time at which each query was run, waits some period of time, and then repeats this process using the recorded times as parameters to the incremental queries. This algorithm is shown in Figure 4.

We can now explain why Tapestry does not allow documents to be deleted (that is, uses an append-only document store). Because the filterer runs at discrete times, if documents could be removed, then different users could receive different results from the same filter, depending on when the filter ran relative to document deletion. This would be a violation of continuous semantics.

4.3 Examples

A couple of examples should give the flavor of the query transformations. Consider the query “show messages sent by Joe,” which can be expressed in TQL as

```
m.sender = “Joe”
```

This query is already monotone since the set of messages sent by Joe is strictly non-decreasing over time. Therefore, the query simply needs to be converted into an incremental form. Recall that the incremental query $Q^I(\tau, t)$ should return messages that began matching the original between times τ and t . For the above example, the incremental query considers all messages that arrived in this time range:

```
m.sender = “Joe” AND ( $\tau < m.ts$  AND  $m.ts \leq t$ )
```

The “ts” field is a timestamp added by Tapestry when the message arrives in the document store.

As a more complicated example, consider the query “show bug reports that are more than 2 weeks old and have not been answered”. In TQL, this can be written as

```
m.to = “BugReports” AND m.ts + [2 weeks] < now()
AND NOT EXISTS (mreply: mreply.in_reply_to = {m})
```

This query is not monotone since it may select a message after the message becomes two weeks old and stop selecting the message when a reply arrives. Tapestry converts it into the following monotone query:

```
m.to = “BugReports” AND m.ts + [2 weeks] < now()
AND NOT EXISTS (mreply: mreply.in_reply_to = {m} AND mreply.ts < m.ts + [2 weeks])
```

This monotone query has a slightly different meaning than the original query, but one that is consistent with continuous semantics. Specifically, it says “show bug reports that are not answered within 2 weeks”.

The incremental version of this query considers all messages that became two weeks old in the time between τ and t :

```
m.to = “BugReports” AND m.ts + [2 weeks] < now()
AND ( $\tau$  < m.ts + [2 weeks] AND m.ts + [2 weeks]  $\leq t$ ) AND
NOT EXISTS (mreply: mreply.in_reply_to = {m} AND mreply.ts < m.ts + [2 weeks])
```

Because t and $\text{now}()$ are the same time in practice, this query can be simplified a bit by removing the “AND $m.ts + [2 weeks] < \text{now}()$ ” clause.

5.0 The Current System

A system that embodies the architecture presented in section 2.0 is currently in use by a small number of researchers. This section describes the implementation of various components of the current Tapestry system.

5.1 Database manager

Tapestry stores documents, annotations, and filter queries in a commercial relational database management system provided by Sybase[11]. Information about messages is stored in a set of relational tables. A single table does not suffice since this information does not fit cleanly into the relational model. In particular, there is no single collection of attributes that apply to all messages, and some of the attributes, such as the set of recipients or newsgroups for a message, are set-valued. Information that is common to all messages, and is not set-valued, is stored in a table that has one entry per message. Other information that varies from message to message is stored in an auxiliary table. Each message may occupy one or more rows in this table. Similarly, set-valued attributes are stored in a special table in which each value of a set occupies a single row. Annotations, which, like messages, have an extensible set of attributes, are stored in several tables as

well. As stated earlier, one of the principal motivations behind the design of TQL was to hide this complex database schema from Tapestry users.

5.2 Indexer

The indexing program is responsible for understanding a given document format, extracting attributes from the document, and storing these in the database. Logically, a separate indexing program exists for each type of document that is added to the Tapestry system. For example, the format of NetNews articles and mail messages is very different than that of New York Times articles. Fortunately, the indexer is the only part of the Tapestry system that is sensitive to the format of a document. New sources of documents can be added simply by writing new indexing programs.

For NetNews, the indexer takes all the header fields in the message and translates them into tapestry message fields. In addition, the words in the body of the message are added to a set-valued Tapestry field named "words." Words on a stop list of common English words are not added, and each word is stemmed to eliminate inflected forms (e.g. "ran" is indexed as "run"). No proximity of frequency information is kept for words in the body.

As of this writing, we are indexing a subset of NetNews (the "comp" subtree), keeping the last 100 MB of data around at any given time. This is about 12 days worth of data, or 43000 messages. Our Sybase tables and indices occupy an additional 300 MB of storage.

5.3 TQL-to-SQL translator

Before a TQL query can be executed over the Tapestry database, it must be converted to SQL, the query language used by the Sybase database manager. The Tapestry system compiles (or translates) each TQL query into an equivalent SQL query. For *ad hoc* queries, this translation is done directly on the query provided by the Tapestry user. For filter queries, the TQL statement is first converted into its bounding monotone query and incrementalized, as described in section 4.0, and then translated into SQL. The SQL query for a filter is then maintained in the Sybase database as a stored procedure. A stored procedure is more efficient than an *ad hoc* query since the query optimization overhead is amortized over the many executions of the query.

Because information about messages and annotations is distributed throughout several tables in the Tapestry database, the SQL equivalents of most TQL queries involve one or more database join operations. Hence, the SQL queries can be quite complicated. Figure 5 shows a sample TQL query along with the resulting SQL query. Studies have shown that a good query optimizer, provided with suitable database indices, can produce query plans that allow these complex queries to run efficiently. In particular, the execution cost of an incremental query produced by our translator is proportional to the number of messages added to the database since the query last ran and not dependent on the overall size of the database. See our paper on continuous queries for more details [13].

Figure 5. Example of translation from TQL to SQL

```
EXISTS(m1:(( $\tau$  < m.ts AND m.ts  $\leq$  Now()) OR ( $\tau$  < m1.ts AND m1.ts  $\leq$  Now()))
AND m1.sender = "Joe" AND m1.in_reply_to = {m})
```

TQL

SQL

```
SELECT m.id FROM msgs m WHERE
EXISTS(SELECT * FROM msgs m1 WHERE
  ((@tau < m.ts AND m.ts <= getdate()) OR
  (@tau < m1.ts OR m1.ts <= getdate())) AND (m1.sender = "joe") AND
EXISTS(
  SELECT * FROM reply_to t1, msgs tm1 WHERE
  t1.id = m1.id AND t1.reply_ref = 1 AND
  t1.msg_id = tm1.msg_id AND tm1.id = m.id))
```

5.4 Remailer

Messages that are selected by a user's filter queries are queued up for delivery to that user. These queues, which constitute the users' "little boxes", are also stored in the Tapestry database. Eventually, we plan to build Tapestry clients that access these queues directly, including a Tapestry browser. In the meantime, we have built a remailing agent that periodically retrieves all of the messages that have been selected for a user and sends each message to that user via electronic mail. Each message is modified to include an extra header field that indicates which filter(s) selected the messages. This is used as input to the appraiser, permits a user to understand why the message was selected and provides valuable feedback for debugging or refining a filter query.

5.5 Mail readers

Having the Tapestry server send selected messages to users electronically eliminates the need to build special clients. Importantly, users can continue to use their favorite mail readers to manage both their private mail and the Tapestry documents selected by their filter queries. While we do not believe this to be the ideal means of interacting with the Tapestry service, it has allowed us to quickly make use of the filtering capabilities.

Some Tapestry clients use the Andrew Messages reader developed at Carnegie Mellon University [9]. Like most modern mail readers, it provides a nice user interface for reading messages and moving them into mail folders. Moreover, it supports the "FLAMES" language, which allows users to write a simple form of "appraisers" that automatically move messages matching a given

predicate or rule into a given folder. In particular, users can write FLAMES rules to identify and process messages that were sent by the Tapestry service and selected by a certain filter query.

To experiment with a different type of appraiser function, we added prioritizing queries to the Cedar-based mail reader developed at Xerox PARC called Walnut [4]. Users can supply a set of queries that can be applied to all incoming messages. As with the FLAMES rules, these queries can look for the special header field indicating that a message is from the Tapestry service. Each query assigns a numerical priority to messages that match the query. If a message matches several queries, then it is assigned the maximum of the priorities. Walnut will display messages within a folder in various orders including priority order. This allows users to quickly see the high-priority messages (and ignore the low priority ones). To date, our experience with prioritizing queries has been quite positive. They have convinced us of the value of having appraisers that further classify and organize messages selected by filter queries.

5.6 Name canonicalizer

It is very common for queries to involve the names of mail senders and receivers. There are two problems with these names. First, a given person usually has multiple electronic names. Second, if a name has any chance of being unique, it must be highly qualified, and that works against our goal having making it easy to type an *ad hoc* query. This subsection presents our design (not yet implemented) for dealing with naming.

The second problem is the easiest to solve. In the “official” TQL query language, names are fully qualified. However, users will normally enter queries via a browser. Thus, the browser can offer an expand command, which takes a shorthand and expands it to be fully qualified. This not only saves typing, but also serves to verify that the name was expanded as expected.

The first problem is more difficult, because there is not a 1-1 mapping between names and people. Suppose we simplify the problem by assuming that each person referenced in a query can be uniquely named with an internet name of the form name@site, where name and site each are of the form part1.part2... . There is still a problem because both names and sites can have many aliases, and so the mapping is many-to-1. In other words, although a person can be specified unambiguously, it is difficult find all documents involving a given person, because of all the aliases.

Our solution involves creating a canonical form for each name, which is a fully qualified internet name, along with a program that converts names to canonical form. For the “From” field of mail originating within PARC, the canonicalizer can do a perfect job. For other names, it must use heuristics.

Once such a canonicalizer exists, it can be used when executing a query such as

```
m.sender = “weiser”
```

It would be too expensive to perform the three steps of examining the Sender field of each document, canonicalizing it, and then comparing that with canonical form of “weiser” every time an incremental query was executed. Instead the raw names in documents are processed as they arrive

in Tapestry. Although the names could simply be replaced with their canonical forms, that isn't done because the canonicalizer is imperfect, and we want to make it easy to update its translations when an error is discovered.¹

Our solution is that as documents arrive in Tapestry, for each raw name in the document, if it has not been seen before, it is run through the canonicalizer, and added to a table that contains [raw name, canonical name] pairs.

Then the query

```
m.sender = "Weiser:PARC:Xerox"
```

is converted to

```
m.sender = names.canonname AND names.rawname = "Weiser:PARC:Xerox"
```

The advantage of having the level of indirection is that we can easily compensate for incorrect heuristics in the canonicalizer by changing entries in the names table.

6.0 Summary and Future Work

Tapestry is an experimental system designed to receive, filter, file and browse electronic documents that arrive in a continuous stream. Because this class of documents includes electronic mail, Tapestry is intended to be used as a replacement for current electronic mail systems.

The novelty of Tapestry lies in its support for collaborative filtering. Users are encouraged to annotate documents, and these annotations can then be used for filtering. We envision two types of readers for various classes of documents. Eager readers will read all the documents in the class in order to get immediate access. More casual readers will wait for the eager readers to annotate, and read documents based on those reviews. Experience with NetNews suggests that there will not be a lack of readers willing to be 'eager' annotators.

When a Tapestry user installs a filter that uses annotations, documents matching that filter are returned as soon as the document receives the specified annotations. Thus Tapestry filters can be thought of as running continuously. The primary technical innovation in Tapestry is an efficient algorithm for implementing filter queries that have predictable semantics.

Future work falls into two categories. First, we need to accumulate more user experience with Tapestry so that we can better analyze how well the design actually works in practice. Second, the Tapestry design presented in this paper is missing a few important pieces. One of them is security: the integration of private mail with public information such as NetNews is unlikely to be widely accepted without a strong security scheme. Another missing piece is the browser. We have not yet done a detailed design of a browser. The integration of different information streams provided by Tapestry may enable some interesting new browser techniques.

1. This has the unfortunate side effect of destroying append-only semantics, but there doesn't seem to be any way around this problem.

7.0 Acknowledgments

We would like to thank Pavel Curtis, Doug Cutting and Maria Okasaki for carefully reading a draft of this paper.

8.0 References

- 1] *DIS 9075:199x(E)*.
Database Language SQL
ANSI, April 1991.
- 2] P. J. Denning.
Electronic Junk.
Communications of the ACM 25(3): 163-165, March 1982.
- 3] D. K. Gifford, R. W. Baldwin, S. T. Berlin, and J. M. Lucassen.
An architecture for large scale information systems.
Proceedings Tenth Symposium on Operating Systems Principles, Orcas Island, Washington, December 1985, pages 161-170.
- 4] Jack Kent, Doug Terry, and Willie-Sue Orr
Browsing electronic mail: Experiences interfacing a mail system to a DBMS. *Proceedings 14th International Conference on Very Large Databases (VLDB)*, Los Angeles, California, August 1988, pp 112-123.
- 5] E. Lutz, H. v. Kleist-Retzow, and K. Hoernig.
MAFIA—An active mail-filter-agent for an intelligent document processing support.
Multi-User Interfaces and Applications, S. Gibbs and A. A. Verrijn-Stuart (editors), North Holland, 1990, pages 16-32.
- 6] T. W. Malone, K. R. Grant, F. A. Turbak, S. A. Brobst, and M. D. Cohen.
Intelligent information sharing systems.
Communications of the ACM 30(5):390-402, May 1987.
- 7] J. Palme.
You have 134 unread mail! Do you want to read them now?
Proceedings IFIP WG 6.5 Working Conference on Computer-based document Services, Nottingham, England May 1984, pp 175-184.
- 8] S. Pollock.
A rule-based message filtering system.
ACM Transactions on Office Information Systems 6(3):232-254, July 1988.
- 9] J. Rosenberg, C. F. Everhart, and N. S. Borenstein.
An overview of the Andrew Message System.
Proceedings SIGCOMM '87 Workshop on Frontiers in Computer Communications Technology, Stowe, Vermont, August 1987, pp 99-108.

- 10] Ben Smith.
The Unix Connection.
Byte, 14(5):245-251, May 1989.
- 11] Sybase.
Transact-SQL user's guide.
Sybase, Inc., October 1989.
- 12] D. B. Terry.
7 steps to a better mail system.
Message Handling Systems and Application Layer Communication Protocols, P. Schicker and E. Stefferud (editors), North Holland, 1991, pages 23-33.
- 13] D. B. Terry, D. Goldberg, D. Nichols and B. Oki.
Continuous Queries Over Append-Only Databases.
Proceedings ACM-SIGMOD Symposium on the Management of Data, San Diego, June 1992.

