

Partial Collection Replication For Information Retrieval

Zhihong Lu Kathryn S. McKinley *

Department of Computer Science, University of Massachusetts

Abstract

The explosion of content in distributed information retrieval (IR) systems requires new mechanisms in order to attain timely and accurate retrieval of unstructured text. In this paper, we investigate using partial replication to search a terabyte of text in our distributed IR system. We use a replica selection database to direct queries to relevant replicas that maintain query effectiveness, but at the same time restricts some searches to a small percentage of data to improve performance and scalability, and to reduce network latency. We first investigate query locality with respect to time and replica size using real logs from THOMAS and Excite. Our evidence indicates that there is sufficient query locality to justify partial replication for information retrieval and partial replication can achieve better performance than caching queries, because the replica selection algorithm finds similarity between non-identical queries, and thus increases observed locality. We then extend the inference network model to rank and select partial replicas. We compare our new selection algorithm to previous work on collection selection over a range of tuning parameters. For a given query, our replica selection algorithm correctly determines the most relevant of the replicas or original collection, and thus maintains the highest retrieval effectiveness while searching the least amount of data as compared with the other ranking functions. We use a validated simulator to report on performance of partial collection replication as a function of locality. We compare collection partitioning to partial replication with load balancing, and find partial replication is much more effective at decreasing query response time than partitioning, even with significantly *fewer* resources, and it requires only modest query locality. We also demonstrate the average query response time under 10 seconds for a variety of work loads with partial replication on a terabyte of text.

Categories and Subject Descriptors: H.3.4 [Information Storage and Retrieval]: Systems and Software—Distributed systems, Performance evaluation (efficiency and effectiveness).

General Terms: Experimentation, Performance.

Additional Key Words and Phrases: Partial replication, Replica Selection, Distributed information retrieval architectures.

*Authors's address: Department of Computer Science, University of Massachusetts, Amherst, MA 01003; Email: {zlu, mckinley}@cs.umass.edu

1 Introduction

As information and users proliferate through the Internet and intranets, distributed information retrieval (IR) systems must cope with the challenges of scale. Distributing excessive workloads and searching as little data as possible while maintaining acceptable retrieval accuracy are two ways to improve performance and scalability of a distributed IR system. A *partial replica* of a collection, which includes query logic as well as a subset of the documents, serves these two purposes. Replicating collections on multiple servers can distribute excessive workloads posed on a single server. Partial replication on servers that are closer to their users can improve performance by reducing network traffic and minimizing network latency. Replicating a small percentage of a collection and directing queries to a relevant partial replica further improves performance by searching less data.

Distributed databases and the Web use *caching and replication* to distribute workloads and improve system performance (see Section 2). Both caching and replication, of course, require query and document *locality*, repeated access to the same part of the database. Our investigation of partial replication is unique from previous work on caching and replication for several reasons. First, structured databases and web caching use a simple, exact test of set membership to determine if the requested record or document matches the query. We could use the same approach, but because users do not all think the same, many distinct queries return the same documents and thus exact match unnecessarily degrades locality and ultimately performance. We are the first to demonstrate this degradation: up to 15% on our traces, even with a very restrictive definition of query similarity.

We examine server logs from THOMAS [THOMAS, 1998] and Excite [Excite, 1997] for our experiments. Although others report on query locality [Croft et al., 1995, Holmedahl et al., 1998], there exists no widely available, shared, or standard query sets with locality properties. We report the locality properties of our traces, and find locality that remains high (above 20%) over time (weeks). These results suggest we may update the replicas infrequently and/or use events to trigger updates, and they will continue to satisfy many queries and improve performance.

To use replicas rather than caches and maintain retrieval effectiveness, a selection function must determine whether replicas contain all, some, or none of the relevant documents for a query. We describe such a function in the context of inference networks and demonstrate its effectiveness using 20 GB TREC VLC collections and TREC queries. For a given query, our replica selection algorithm correctly determines the most relevant of the replicas or original collection, and thus maintains the highest retrieval effectiveness while searching the least amount of data as compared with the ranking functions for collection ranking.

We also report on performance as a function of locality using a validated simulator [Lu, 1999]. The

simulator closely matches our prototype system that uses InQuery for the basic IR functionality on all collections, original and replicated [Callan et al., 1995a]. We compare the performance of searching a terabyte of text using partial replication to partitioning, and find partial replication is more effective at reducing execution time, even with many fewer resources, and it requires only modest query locality to achieve better, sometimes much better, performance. We are the first to report performance for a terabyte of text; whereas previous work is on the scale of gigabytes.

The remainder of this paper is organized as follows. The next section further compares our work to related work. In Section 3, we characterize the locality and access patterns of our test logs from THOMAS and Excite. We also define our notion of query similarity, and show it increases locality up to 15%. Section 4 describes the replication architecture. Section 5 describes how to select a partial replica based on relevance, and compares its effectiveness with the ranking functions for collection ranking. Section 6 reports on the performance of searching a terabyte of text using partial replication and compares performance with collection partitioning. Section 7 summarizes our results and concludes.

2 Related Work

Both research and commercial distributed database systems have used replication for a long time to improve system performance and availability (e.g., [Ahamad and Ammar, 1989, Huang and Wolfson, 1993, Oracle Company, 1995, Wolfson and Jajodia, 1992]). They use structured data, such as objects, and present algorithms for updating read-write data to ensure consistency of different copies of data. In a structured database, query logic is set membership, a straightforward test. Text IR systems instead return a user parameterized range of responses with varying belief values, and thus the system cannot summarize queries into set membership tests.

Partial replicas in IR systems may contain all, some, or none of the relevant documents for a given query. How to select a partial replica based on relevance for a *query* where the answer is not well defined is a problem that does not exist in distributed non-text database systems. As far as we know, the work presented in this paper is the first that attacks this problem. The closest work to selecting a relevant partial replica which is a subset of the original collection is *collection selection*, i.e., locating the most relevant collections [Callan et al., 1995b, Chakravarthy and Haase, 1995, Danzig et al., 1991, Fuhr, 1996, Gravano et al., 1994, Voorhees et al., 1995], where collections are disjointed. Replica selection differs also because it directs as many queries as possible to relevant partial replicas in order to obtain performance improvements.

Danzig et al. [Danzig et al., 1991] use a hierarchy of brokers to maintain indices for document abstracts as a representation of the contents of primary collections, and support Boolean keyword matching to locate

the primary collections. If users' queries do not use keywords in the brokers, they have difficulty finding the right primary collections.

Voorhees et al. exploit similarity between a new query and relevance judgments for previous queries to compute the number of documents to retrieve from each collection [Voorhees et al., 1995]. Netserf extracts structured, disambiguated representations from the queries and matches these query representations to hand-coded representations [Chakravarthy and Haase, 1995]. Voorhees et al. and Netserf require manual intervention which limits them to relatively static and small collections.

Fuhr proposes a decision-theoretic approach to solve collection selection problem [Fuhr, 1996]. He makes decisions by using the expected recall-precision curve, expected number of relevant documents, and cost factors for query processing and document delivery. He does not report on effectiveness.

GLOSS uses document frequency information for each collection to estimate whether, and how many, potentially relevant documents are in a collection [Gravano and Garcia-Molina, 1995, Gravano et al., 1994]. The approach is easily applied to large numbers of collections, since it stores only document frequency and total weight information for each term in each collection. However its effectiveness remains unknown due to limited evaluation.

Callan *et al.* adapt the document inference network to ranking collections by replacing the document node with the collection node [Callan et al., 1995b]. Similar to GLOSS, the information stored in the collection ranking inference network is document frequencies and term frequencies for each term in each collection. Experiments using the InQuery retrieval system and the 3 GB TREC 1+2+3 collection show that using this method to select the top 50% of subcollections attains similar effectiveness to searching all subcollections.

None of the above collection selection algorithms consider replicas and partial replicas. Among the approaches for collection selection, the collection inference network model [Callan et al., 1995b] is the most thoroughly tested and effective. In this paper, we modify this technique to rank partial replicas and the original collection, propose a new algorithm for replica selection, and show that it is effective and improves performance.

Recently, researchers have used replication in the Web for document access [Baentsch et al., 1996, Bestavros, 1995, Yu and MacNair, 1998]. Generally, they cache popular documents in a hierarchy of proxy servers which are located between clients and Web servers to reduce Internet traffic. Although we also adopt a replication hierarchy to store the documents of the most frequently used queries, our replicas are searchable and thus can speed up both query and document processing. In addition, our replica selection function finds locality among non-identical queries, which increases observed locality. Since current Web caches and

database replicas use exact match, they cannot find this locality.

A number of studies have investigated the performance of distributed information retrieval (IR) systems. Most of the previous work experiments with a collection less than 1 GB and focuses on speedup when a collection is distributed over several servers [Burkowski, 1990, Harman et al., 1991, Lin and Zhou, 1993, Macleod et al., 1987, Martin et al., 1990, Martin and Russell, 1991]. Only two systems, Couvreur *et al.* [Couvreur et al., 1994] and Cahoon and McKinley [Cahoon and McKinley, 1996], use simulation to experiment with more than 100 GB of data. To our knowledge, no one has reported the performance of partial replication for searching a terabyte of text in distributed information retrieval systems. None of these previous studies include partial replication, and of course they do not compare collection partitioning with replication.

InQuery, our base system, is not the fastest text retrieval system available today [Hawking et al., 1998]. For the experiments in this paper, we model and validate against a 3 processor 250MHz Alpha which can maintain response times of under 10 seconds with 4 to 5 disks on a collection size of up to 16 GB for a heavily loaded system. Other multiprocessor systems [Hawking et al., 1998] have recently reported results for a single query (rather than a loaded system) that are less than a second on 100 GB collection. We have simulated such response times, and found the same trends we report here. We thus believe our results on replication versus partitioning are directly applicable to these systems, although they store more data on a single machine. Localizing significantly more data in a single machine coupled with faster CPU and disk processing correspondingly decreases the number of CPUs and disks per CPU we would need to achieve similar performance, and serves to decrease communication over the local network, but it will not change the general performance trends we report here.

3 Access Characteristics in Real Systems

In this section, we examine query locality in two logs from real systems. We examine query similarity versus exact match, how locality changes over time, and its effect on the size of replicas. We also suggest mechanisms for keeping replicas up to date.

Since currently there exists no widely available, shared, or standard set of queries with locality properties, we obtained our own sets of server logs from THOMAS [THOMAS, 1998] and Excite [Excite, 1997]. The THOMAS system is a legislative information service of the U.S. Congress through the Library of Congress. THOMAS contains the full text Congressional Records and bills introduced from the 101st Congress to 105th Congress. We analyze the logs of THOMAS between July 14 and September 13, 1998, during which the Starr Report became available. We obtained full day logs for 40 days, and partial logs for remaining 22 days due to lack of disk space in the mailing system of the library of Congress. The Excite

Num. queries	Num. unique queries	Topics			
		total	occurring once	more than once	more than one unique query
8143 (7703)	4876 (4651)	4069	2888 (71%)	1181 (29%)	412
percentages of queries that top topics account for					
100	200	500	1000	2000	
21.2%	28.7%	41.5%	54.1%	73.0%	
percentages of queries that top unique queries account for					
100	200	500	1000	2000	
18.1%	24.5%	36.4%	49.4%	64.5%	

(a) Query locality in the THOMAS log

Num. queries	Num. unique queries	Topics			
		total	occurring once	more than once	more than one unique query
499836 (444899)	365276 (320987)	249405	196672 (79%)	52733 (21%)	32750
percentages of queries that top topics account for					
500	1000	5000	10000	20000	
12.3%	16.0%	27.9%	34.4%	42.0%	
percentages of queries that top unique queries account for					
500	1000	5000	10000	20000	
7.9%	10.4%	18.4%	23.0%	28.2%	

(b) Query locality in the Excite log

Table 1: Query locality in the logs

system provides online search for more than 50 million Web pages. The Excite log we obtained contains one day of log information for September 16, 1997.

Since the logs do not contain document identifiers returned from query evaluation, we built our own test databases to cluster similar queries. We define a *topic* as all queries whose top 20 documents completely overlap. This definition of query similarity is arbitrary and restrictive; a looser definition would further improve the locality we observe. For queries from the THOMAS log, we reran all queries against a test database that uses the Congress Record for 103rd Congress (235 MB, 27992 documents). For queries from the Excite log, we reran all queries against a test database using downloads of the websites operated by ten Australian Universities (725 MB, 81334 documents).

3.1 Query Locality

Table 1 shows query locality statistics for our THOMAS and Excite logs. We collect the average number of queries, unique (singleton) queries, topics, topics occurring once, and topics that contain more than one unique query. We also present the percentages of queries that correspond to the top topics and top unique queries, respectively. Table 1(a) shows the average numbers in the THOMAS logs over 40 days with full day logs. The numbers of queries that actually find matching documents from our test database are in the parentheses in columns 1 and 2. Some queries do not find any matching documents, due to misspelling, or because query terms do not exist in the test database. The statistics show that on the average, 29% of

topics occur more than once, and they account for 63% ((7703-2888)/7703) of queries. Among the topics occurring more than once, 35% (412) contain more than one unique query. The top 100 topics (2.5% of topics) and the top 500 topics (12% of topics) account for 21.2% and 41.5% of queries, while the top 100 unique queries and top 500 unique queries account for 18.1% and 36.4%.

The Excite log on September 16, 1997, shown in Table 1(b), demonstrates that the Excite queries also have high query locality: 21% of topics occur more than once, and they account for 56% ((444899-196672)/444899) of queries. Among the topics occurring more than once, 62% (32750) contain more than one unique query. The top 1000 topics and the top 10000 topics account for 16.0% and 34.4% of queries, while the top 1000 unique queries and top 10000 unique queries account for 10.4% and 23.0%. Both sets of logs see a drop in locality between 3 and 14% if we require an exact match.

3.2 Locality as a Function of Time

We also examine the THOMAS logs to see how many queries on a given day match a topic or a query that appears on a previous day or week, in order to examine the overlap as a function of time. Table 2 shows for a number of days between 7/15 and 9/11 the percentage of queries that match a top query through topic match and exact query match on a previous day or week. Column 1 lists date. Columns 2 through 4 list the query overlap when we build a replica using top topics or top unique queries on the previous day and update it daily. Columns 5 through 7 list the query overlap when we build a replica using top topics or top unique queries on July 14, 1998, without update afterwards. Columns 8 through 10 list the query overlap when we build a replica using top topics or top unique queries in the week from July 14 to July 20, 1998, without update afterwards. Replicas may actually satisfy more queries than we report, because we do not include queries whose top documents appear in the replica because the response is a combination of two or more other topic queries. Since the logs do not contain document identifiers and our test database is pretty small, we can not obtain accurate figures about this situation.

The statistics also show that topic matching finds up to 15% more overlap than exact query match for the same size replicas over time. For example on 9/11, we saw many distinct queries such as “Starr,” “Starr Report,” “Bill Clinton,” and “Monica Lewinsky,” all trying to access the Starr Report.

For topic match, we build a replica with the top documents for the top topics. For exact query match, we build a replica with the top documents for the top unique queries. For example, when we build replicas using top 1000 topics or unique queries of the week of 7/14-7/20, topic match on 7/23 increases the overlap from 28.6% (query exact match) to 35.9%, which means a replica can distract 7.3% more queries from the original server. Replicating more topics further widens this difference.

date	Overlap with								
	the previous day			7/14			the week on 7/14-7/20		
	Topic match			Topic match			Topic match		
	all	top 500	top 1000	all	top 500	top 1000	all	top 500	top 1000
7/15	43.3%	24.8%	30.1%	43.3%	24.8%	30.1%	n/a	n/a	n/a
7/16	44.4%	24.4%	30.4%	42.6%	24.0%	29.2%	n/a	n/a	n/a
7/23	45.0%	27.3%	31.5%	41.4%	23.4%	28.7%	60.8%	29.0%	35.9%
7/31	n/a	n/a	n/a	38.5%	21.9%	26.4%	58.0%	26.0%	32.3%
8/14	36.6%	21.9%	26.1%	38.1%	21.3%	26.0%	54.9%	25.6%	31.0%
8/28	32.9%	19.1%	23.0%	34.3%	18.3%	23.4%	51.9%	22.8%	28.4%
9/11	78.1%	69.2%	71.7%	44.0%	8.7%	22.2%	58.6%	11.2%	27.0%

date	Exact query match			Exact query match			Exact query match		
	all	top 500	top 1000	all	top 500	top 1000	all	top 500	top 1000
7/15	33.1%	18.9%	23.3%	33.1%	18.9%	23.3%	n/a	n/a	n/a
7/16	34.5%	19.3%	23.0%	32.9%	18.1%	22.4%	n/a	n/a	n/a
7/23	36.4%	21.1%	24.6%	32.3%	17.9%	22.3%	49.4%	23.7%	28.6%
7/31	n/a	n/a	n/a	29.4%	16.9%	20.3%	46.5%	20.8%	25.1%
8/14	28.2%	16.3%	20.0%	29.0%	16.4%	20.0%	43.4%	20.2%	24.1%
8/28	25.4%	14.5%	17.6%	25.9%	14.0%	17.5%	41.2%	18.2%	22.2%
9/11	71.8%	63.6%	65.2%	24.9%	6.6%	18.7%	43.2%	8.2%	19.3%

Table 2: Overlap over time in the THOMAS log: Topics vs. exact query match

Top topics	% of queries	Replica Size (top 200 documents per query)		
		(2 KB per doc)	(3 KB per doc)	(9 KB per doc)
1000	16.0%	400 MB	600 MB	1.8 GB
5000	27.9%	2 GB	3 GB	9 GB
10000	34.4%	4 GB	6 GB	18 GB
20000	42.0%	8 GB	12 GB	36 GB

Table 3: The Replica Size Based on the Excite log

3.3 Estimating the Size of Replicas

Based on query locality, we may estimate the replica size, which is a function of average document size, query locality, and number of top documents per query we chose to store, as shown in Table 3. The average document size varies from source to source. For example, the average document sizes of the USENET News, Wall Street Journal, and the websites operated by 10 Australia Universities are 2 KB, 3 KB, and 9 KB, respectively [Harman, 1997]. The average document size of the 20 GB TREC VLC collection is 2.8 KB [Harman, 1997]. The TREC VLC text collection consists of data from 18 sources, such as news, patents, and Web sites. Our estimation uses three different numbers: 2 KB, 3 KB, and 9 KB. For query locality, we use the statistics obtained from the Excite log, since its workloads are at the level of the system we investigate. We obtain the top 200 documents for each query. We overestimate these sizes because we assume there is no overlap among the documents, although as shown above, distinct queries often result in overlapping documents. In Table 3, columns 1 and 2 show the query locality from the Excite log; columns 3 through 5 show the estimated replica size when we vary the average document size. For example, a 4 GB, 6 GB, and 18 GB replica satisfies at least 34.4% of queries with an average document size of 2 KB, 3 KB,

and 9 KB, respectively.

3.4 When to Build or Update a Replica

Query overlap tends to decrease very gradually as time elapses. For example, 35.9% of queries on 7/23 and 32.3% of queries on 7/31 matched a topic in the replica covering top documents for top 1000 topics of the week of 7/14-7/20, respectively. These statistics suggest that we do not need to update the replica daily on a typical day, since significant numbers of queries match a top query that appeared several days ago. However we do need some mechanism to deal with a bursty event like the Starr Report, as shown by the sharp decrease in locality on 9/11. Regular, daily updating would catch this event, but it may be too costly, react too slowly, or unnecessarily degrade performance when the system experiences the expected gradual degradation of locality. We propose two on-demand updating strategies as follows:

- Event triggered updating: watch for bursty events, and trigger the updating procedure when some special events happen.
- Performance triggered updating: watch the percentage of workloads the replica selector sends to the replicas, and trigger the updating procedure when the percentage falls below some threshold.

For event triggered updating strategy, we can simply use human intervention. When the system manager anticipates or observes a special event, and increasingly many users issue queries on it, she initiates the updating procedure. Automatic event detection is an on-going research topic. When it becomes effective, we suggest using it to trigger the updating procedure automatically. Instead of rebuilding a replica, we could add documents into replicas without deleting others for quicker updates, which means we need to save some extra space for bursty events.

Performance triggered updating is very easy to implement in the current system. We let the replica selector record the percentage of queries that it sends to each replica, when the percentage falls below a threshold, the system informs the system manager. Performance triggered updating also works for bursty events, if a lot of users search for an event that does not exist in the replicas.

4 Replication Architecture

In this section, we describe how to build a hierarchy of replicas in our system. We replicate the top documents and their index in a partial replica which helps queries about the same and similar topics but that use different terms to find their relevant documents in these partial replicas. We determine which documents to replicate as follows: for a given query, we tag all top n documents that query processing returns as “accessed” and increment their access frequencies, regardless of whether the user requests the text of these

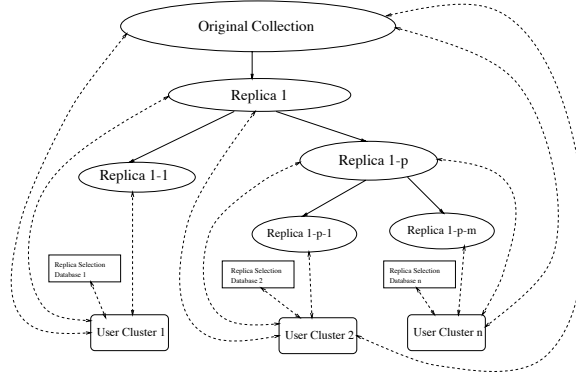


Figure 1: The replication hierarchy

documents. We keep the access frequency for each document within a time period, e.g., a week, and then replicate the most frequently accessed documents the most.

We organize replicas as a hierarchy, illustrated in Figure 1. The top node represents an original collection that could be a single collection residing on a network node or a virtual collection consisting of several collections distributed over a network. The bottom nodes represent users. We may divide users into different clusters, each of which reside within the same domain, such as an institution, or geographical area. The inner nodes represent partial replicas. The replica in a lower layer is a subset of the replicas in upper layers, i.e., $\text{Replica 1-1} \subset \text{Replica 1} \subset \text{Original Collection}$. The replica that is closest to a user cluster contains the set of documents that are most frequently used by the cluster. An upper layer replica may contain frequently used documents for more than one cluster of users. The solid lines illustrate data is disseminated from the original collection to replicas. Along the arcs from the original collection, the most frequently used documents are replicated many times.

The replica selection database directs queries to a relevant partial replica. It may send user requests to any replica or the original collection along the arcs from the top node depending on relevance and other criteria, such as server load. The dotted lines illustrate the interaction between users and data. If we do not divide the users into different groups, the hierarchy is simply a linear hierarchy. Replica selection is a two-step process in this architecture: it ranks replicas based on relevance, and then selects one of the most relevant replicas based on load.

In the next section, we will show that the inference network model is very effective at selecting a relevant replica. We implement the replica selection inference network as a pseudo InQuery database, where each pseudo document corresponds to a replica or collection, and its index stores the document frequency and term frequency for each term in any of the replicas. Since the replica selection database stores document frequency and collection term frequency for each term that occurs in any of replicas, its size is determined




Figure 2: The collection retrieval inference network.

by the number of unique terms in the largest replica. Based on our observation in our system, the size of the replica selection database is approximately 6 MB for every 100,000 unique terms. We know the 20 GB TREC VLC collection has 13,880,064 unique terms. If our largest replica is 20 GB, the estimated size of the replica selection database is around 1.2 GB. Based on these statistics, we estimate the replica selection database for 1 terabyte of text is between 1 and 2 GB.

5 Partial Replica Selection Based on Relevance

The first step of replica selection is how to find a partial replica that contains enough relevant documents for a given query. In this section, we investigate how to do this task with inference networks, and evaluate the effectiveness of our replica selection approach using the InQuery retrieval system [Callan et al., 1992], and the 2 GB TREC 2+3 collection and the 20 GB TREC VLC collection. We use queries developed for TREC topics 51-350 in our experiments. We compare our proposed replica selection function with the collection ranking function. We measure the system’s ability to pick the expected partial replica, and the precision of the resulting response as compared with searching the original collection.

The rest of the section is organized as follows: Section 5.1 investigates how to rank partial replicas and the original collection using the inference network model, Section 5.2 describes the experimental settings, Section 5.3 compares our proposed replica selection function with the collection selection function, Section 5.4 and Section 5.5 further demonstrate the effectiveness of our approach for both replicated queries and unreplicated queries, and Section 5.6 summarizes the results of this section.

5.1 Ranking Partial Replicas with the Inference Network Model

We adapt the collection retrieval inference networks [Callan et al., 1995b] to rank partial replicas and the original collection. The collection retrieval inference network model consists of two component networks: a

$$T = \frac{df_{ij}}{df_{ij} + k \cdot ((1 - b) + b \cdot \frac{cw_i}{ave_cw}}$$

$$I = \frac{\log(\frac{|N|}{cf} + 0.5)}{\log(|N| + 1.0)}$$

$$P(r_j|D_i) = \alpha + (1 - \alpha) \cdot T \cdot I$$

where

df_{ij} is the number of documents that contain term r_j in collection D_i ,
 cw_i is the number of words in collection D_i ,
 ave_cw is the average number of words,
 N is the number of collections,
 cf is the number of collections that contain r_j .
 k is a constant that controls the magnitude of df (the default is 200),
 b is a constant varying from 0 to 1 used to control the sensitivity of the function to cw (the default is 0.75), and
 α is a default belief (set to 0.4).

Figure 3: The collection ranking function in InQuery.

collection network and a query network, illustrated in Figure 2. The D_i nodes correspond to collections, and the r_j nodes correspond to concepts in the collections. The Q node represents a query, and the ϵ nodes correspond to query concepts in the query. By using the collection retrieval inference network, collection ranking becomes an estimate of $P(I|D_i)$ from combining the conditional probabilities through the network. When we adapt the collection retrieval inference network model to rank replicas, we use D_i nodes to represent the original collection and partial replicas, where D_n represents the original collection, $D_i, i = 1, 2, \dots, n - 1$ represent partial replicas, and $D_1 \subset D_2 \subset \dots \subset D_n$. (In the collection retrieval inference network, D_i nodes do not have a subset relationship.) The purpose of ranking partial replicas is to find a *single* replica that satisfies a given query instead of a subset of collections in the collection retrieval inference network. We refer to this inference network as to the replica selection inference network. As in the collection retrieval inference network model, $P(\alpha_k|r_j)$ is set to 1.0. The central work of applying this inference network to replica selection is to develop an effective replica ranking function to estimate $P(r_j|D_i)$.

Since we adapt the collection retrieval inference network, we first examine whether the InQuery collection ranking function works well with ranking partial replicas. The InQuery collection ranking function uses df (the document frequency of each term) as the basic metric, as shown in Figure 3 [Callan et al., 1995b]. In our experiment settings in Section 5.2, the default InQuery collection ranking function directs more than 70% of the replicated queries to the original collection, however, since we use these replicated queries to build replicas, the replica selector should direct them to the replicas instead of the original collection. Although we can tune the parameters of the InQuery collection ranking function to direct more queries to the

$$ave_tf = \frac{ctf_{ij}}{df_{ij}}$$

$$cutoff_i = cutoff_1 \frac{\log(DN_i)}{\log(DN_1)}$$

$$AT = \begin{cases} ave_tf & \text{if } df_{ij} > cutoff_i \\ ave_tf \cdot \frac{df_{ij}}{cutoff_i} & \text{otherwise} \end{cases}$$

$$T = \frac{AT}{AT + k \cdot ((1 - b) + b \cdot \frac{ave_doclen_i}{ave_ave_doclen})}$$

$$I = \frac{\log(\frac{|N|}{rf} + 0.5)}{\log(|N| + 1.0)}$$

$$P(r_j|D_i) = \alpha + (1 - \alpha) \cdot T \cdot I$$

where

ctf_{ij}	number of occurrences of term r_j in replica/collection D_i ,
df_{ij}	number of documents that contain term r_j in D_i ,
DN_i	number of documents in D_i ,
$cutoff_1$	cutoff value for the smallest replica D_1 , which we set as the number of top documents for each query,
$cutoff_i$	cutoff number of documents in D_i ,
N	number of replicas plus the original collection,
rf	number of replicas and the collection that contain r_j ,
ave_doclen_i	average document length in D_i ,
ave_ave_doclen	average ave_doclen_i ,
k	constant that controls the magnitude of AT ,
b	constant varying from 0 to 1 used to control the sensitivity of the function to ave_doclen , and
α	default belief (set to 0.4).

Figure 4: The replica selection function.

replicas, the precision drops too much, for example, the precision drops approximately 25% when the function directs 80% of replicated queries to the replicas (see Section 5.3 for the details). The InQuery collection ranking function does not work well with replica selection, because it favors collections with larger df , but partial replicas typically have smaller df than the original collection.

Since a partial replica contains the top documents of the most frequently used queries, by examining the document ranking function, we know that the top documents are ranked as the top, just because query terms occur more often in these documents than the others. Therefore if a replica contains the top documents for a query, the average term frequency of each query term in the replica should be higher than in the original collection. Based on this heuristic, we construct a replica selection function based on the average term frequency. In addition, we find a term is important in selecting replicas if it occurs often (with middle or high term frequency) in that replica/collection and it also occurs in a certain number of documents (above a cutoff for document frequency). A term occurring in too few documents does not help even though it has high term frequency. We need to ignore these terms. Figure 4 illustrates our replica selection function which

uses the average term frequency and penalizes the terms that appear less than a given cutoff number in the corresponding replica/collection. We compare this function with the InQuery collection ranking function in Section 5.3, and demonstrate its effectiveness using 350 TREC queries on a 2 GB collection and a 20 GB collection in Section 5.4 and Section 5.5.

We implement the replica selection inference network as a pseudo InQuery database, where each pseudo document corresponds to a replica or collection, its index stores the df (document frequency) and ctf (replica/collection term frequency) for each term. We do not store any proximity information in order to minimize the space requirements of the replica selection database. As in the collection retrieval inference network, all proximity operators are replaced with Boolean AND operators.

5.2 Experimental Settings

We evaluate the effectiveness of our replica selection approach using InQuery [Callan et al., 1992], and a 2 GB TREC collection that contains TREC 2+3 collections, and a 20 GB collection that contains all TREC 6 VLC collections. We use queries developed for TREC topics 51-350 in our experiments. We measure the system’s ability to pick the relevant partial replica, and the precision of the resulting response as compared with searching the original collection. We use TREC queries instead of the queries from the logs, because some of TREC queries have relevance judgments that enable us to produce precision and recall figures for evaluating the effectiveness.

By using the 2 GB collection, we compare the effectiveness of our replica selection function with the InQuery collection ranking function using short queries, and demonstrate the effectiveness of our replica selection function using both short queries and long queries. A short query is simply a sum of the terms in the corresponding description field of the topic. Long queries are automatically created from TREC topics using InQuery query generation techniques [Callan et al., 1992], which consist of terms, phrases and proximity operators. Generally, a long query for a topic is more effective than the short query [Callan et al., 1992]. The average number of terms per query for the set of short queries is 8 after removing the stopwords, and the average number of terms per query for the set of long queries is 120. For each set of queries, we divide queries into two categories: replicated queries and unreplicated queries, where *the replicated queries are those whose top documents are used to build the replicas*. Since only topics 51-150 and topics 202-250 have relevance judgment files for the 2 GB TREC collection, a single trial contains 50 random unreplicated queries from these 149 topics, and use them report the effectiveness for these topics.

We conduct our experiments by repeating the following procedure 5 times, each trial uses a different number as the seed to produce random numbers, and thus picks different queries for a query set. In each

trial, we randomly choose 50 queries from queries $\{51-150, 202-250\}$ as our *unreplicated* query set T , and randomly divide the remaining 250 queries in queries 51-350 into 5 sets: $\{Q, i = 1, 2, 3, 4, 5\}$, each set containing 50 queries. We then build a 6-layer replication hierarchy by using the 2 GB TREC collection or the 20 GB collection as the original collection C , and collecting the top n documents resulting from searching the original collection for each query in $\{Q, i = 1, 2, 3, 4, 5\}$ to build 5 partial replicas $\{D_i, i = 1, 2, 3, 4, 5\}$, where D_i contains at most $n * i$ documents, consisting of the top n documents for each query in query sets $\{Q_j, j = 1, \dots, i\}$. Clearly, $D_1 \subset D_2 \subset D_3 \subset D_5 \subset C$. This structure mimics 5 replicas that increase in size and thus includes more of the top queries. We build a replica selection inference network to rank these five replicas and the original collection. The queries in query sets $\{Q, i = 1, 2, 3, 4, 5\}$ are called *replicated* queries.

By using the 20 GB collection, we examine how the size of collection affects the effectiveness of our replica ranking function. Since we do not have relevance judgment files for topics 51-150 and topics 202-250 against the 20 GB collection, and the 2 GB collection is a subset of the 20 GB collection, we use the relevance judgment files for the 2 GB collection to produce the precision figures. We also conduct another set of experiments in order to make up for insufficient relevance judgments for topics $\{51-150, 202-250\}$. We use queries 301-350 as our unreplicated query set T , since these 50 topics are more thoroughly judged against the 20 GB VLC collection than topics $\{51-150, 202-250\}$. We use queries 51-100 as Q_1 , 101-150 as Q_2 , 151-200 as Q_3 , 202-250 as Q_4 , and 251-300 as Q_5 .

When using the 2 GB collection as the original collection, the size of replicas ranges from 0.3% to 1.5%, 1% to 5%, 2% to 10%, and 5% to 20% of the original collection when replicating the top 30, 100, 200, and 500 documents, respectively. When using the 20 GB collection as the original collection, the size of replicas ranges from 0.1% to 0.5%, 0.2% to 1%, and 0.5% to 2% of the original collection when replicating the top 100, 200, and 500 documents, respectively.

When we evaluate a document or collection ranking function, we say a function is better than others if and only if it can produce higher precision at selected numbers of documents or at all standard levels of recall. In the case of replica selection, we need to add another criterion for the ranking function: directing as many queries as possible to the relevant replicas in order to improve system response time. We can tune the parameters of our functions to control the percentage of replicated queries to the replicas (as shown in Section 5.3). The range varies from 0% to 90%. None of the function we tested can direct 100% of replicated queries to the replicas. However when we direct more queries to the replicas, we have to tolerate a larger precision loss. In our experiments, we compare the precision of each function when it directs more than 80% of replicated queries to the replicas.

For a replicated query, since we know which replica contains its top documents, we define its *expected replica* as the smallest replica that is built with the top documents for the query. For an unreplicated query, since replicas may contain some relevant documents, we expect our replica selector will direct some of these queries to a relevant replica. We define the *expected replica* for an unreplicated query as the smallest replica that causes a precision drop less than 5%. For both kinds of queries, especially unreplicated queries, we expect we will have to tolerate some loss in precision in order to avoid searching the entire collection. We choose a drop in precision between 0 and 10% for a query as our acceptable range, i.e., searching the selected replica retrieves at most one less relevant document for every 10 documents as compared with searching the entire original collection.

We define *collection precise queries* as those queries that can achieve the precision above 10% when searching the original collection for the top n documents, i.e., the query finds at least one relevant document for every ten documents. We exclude collection imprecise queries when we present the ability of a replica selector to pick the relevant replicas for unreplicated queries, because a replica with zero relevant documents is probably an acceptable choice for a query whose precision is below 10% in the original collection. We define *replica precise queries* as those for which searching the selected replica causes a precision loss less than 5% of the precision attained by searching the original collection.

5.3 Comparing Ranking Functions

In this section, we compare the effectiveness of the InQuery collection ranking function illustrated in Figure 3 and our replica selection function illustrated in Figure 4 by varying k and b for short queries in test trial 1 when we replicate the top 200 documents for each query. (We also performed experiments replicating the top 100 and 500 documents with similar results.) We will show that our replica selection function is comparable to the collection ranking function in ability to pick the expected replica for replicated queries, but that it significantly improves precision and finds the expected replica much more consistently for unreplicated queries.

Table 4 lists the results of replica selection by counting the number of queries and to which replica or collection each function directs the queries, when the parameters, k and b , vary. Table 4(a) lists the results for 99 replicated queries for which we have relevance judgments. Table 4(b) lists the results for 37 unreplicated collection precise queries, 18 of which are replica precise queries. In both tables, columns 1 through 3 list the name of functions, the values of parameters k and b , and the function abbreviations. For replicated queries, columns 4 through 9 contain the number of queries that the replica selector sends to each of the replicas (D_i) as well as the original collection (C); columns 10 through 13 contain the percentages

Ranking Function	Parameters k, b	Func. code	Replica					C	% to replicas			C
			D ₁	D ₂	D ₃	D ₄	D ₅		right	smaller	larger	
Expected		E	18	16	25	21	19	0	100%	0%	0%	0%
Random		Ran	17	17	17	16	16	16	16%	35%	33%	16%
InQuery Collection Ranking Function	200, 0.25	I1	0	0	0	0	0	99	0%	0%	0%	100%
	200, 0.75	I2	0	1	4	5	20	69	14%	0%	16%	70%
	200, 1	I3	28	14	22	14	10	11	65%	16%	8%	11%
	100, 1	I4	28	15	22	14	10	10	65%	17%	8%	10%
	400, 1	I5	29	14	23	14	8	11	64%	17%	8%	11%
Replica Selection Function	2, 0	R1	22	12	10	12	32	11	59%	7%	23%	1%
	2, 0.2	R2	20	15	11	17	24	12	57%	9%	22%	12%
	2, 0.8	R3	6	3	3	5	1	81	15%	1%	2%	82%
	1, 0.2	R4	17	6	7	14	28	27	47%	5%	20%	27%
	4, 0.2	R5	21	10	10	11	26	21	54%	7%	18%	21%

(a) Replicated queries (99 queries)

Ranking Function	Parameters k, b	Func. code	Replica					C	Precision loss			% of replica precise queries to C
			D ₁	D ₂	D ₃	D ₄	D ₅		C + < 5%	5%–10%	> 10%	
Expected		E	1	6	3	6	2	19				
Random		Ran	7	8	5	7	4	6	35%	19%	46%	22%
InQuery Collection Ranking Function	200, 0.25	I1	0	0	0	0	0	37	100%	0%	0%	100%
	200, 0.75	I2	0	0	0	0	6	31	89%	3%	8%	89%
	200, 1	I3	6	5	11	7	2	6	40%	30%	30%	11%
	100, 1	I4	5	6	12	7	2	5	38%	30%	32%	11%
	400, 1	I5	6	6	11	11	6	2	40%	30%	30%	11%
Replica Selection Function	2, 0	R1	8	7	1	2	9	10	51%	19%	30%	6%
	2, 0.2	R2	7	6	2	2	8	12	68%	16%	16%	11%
	2, 0.8	R3	0	0	0	0	1	36	100%	0%	0%	94%
	1, 0.2	R4	4	2	1	1	13	16	73%	14%	14%	22%
	4, 0.2	R5	7	2	2	1	10	15	64%	14%	22%	22%

(b) Unreplicated queries (37 collection precise queries)

Table 4: Comparing ranking functions using short queries on the 2GB TREC 2+3 collection (replicas built with top 200 documents)

of queries that are directed to the expected replica (right), smaller replica, larger replica, and the original collection. The “expected” (E) row lists the number of judged queries that a perfect replica selector would direct to each replica and to the original collection. For unreplicated queries, columns 4 through 9 contain the number of collection precise queries the replica selector sends to each of the replicas as well as the original collection; column 10 contains the percentages of collection precise queries that are directed to the original collection and the replicas that cause a precision loss less than 5%; columns 11 through 12 contain the percentages of collection precise queries that are directed to replicas that cause a precision loss from 5% to 10%, and more than 10%. Column 13 contains the percentage of 18 replica precise queries that are directed to the original collection. The “expected” (E) row for the unreplicated queries contains the number of queries that we expect to go to each of replicas and the original collection because the result will cause a 5% or less drop in precision.

For the InQuery collection ranking function, varying k : from 100 to 400 does not significantly change

at m docs	Precision of Replicated Queries (%)						
	C	random	I3	I4	I5	R1	R2
10	48.7	40.4 (-17.2)	47.7 (-2.3)	48.2 (-1.2)	47.2 (-3.3)	48.4 (-0.8)	48.3 (-1.6)
20	44.8	36.1 (-19.6)	43.2 (-3.7)	43.3 (-3.4)	42.9 (-4.4)	44.4 (-1.0)	44.2 (-1.4)
30	40.7	32.4 (-20.4)	39.3 (-3.4)	39.4 (-3.0)	39.1 (-4.0)	40.2 (-1.1)	40.2 (-1.2)
100	31.1	23.9 (-23.1)	29.4 (-5.6)	29.5 (-5.5)	29.3 (-6.1)	30.5 (-2.2)	30.5 (-2.1)
200	25.1	18.7 (-25.3)	23.0 (-8.4)	23.0 (-8.4)	22.9 (-8.7)	24.8 (-2.8)	24.3 (-3.2)

(a) 99 Replicated queries

at m docs	Precision of Unreplicated Queries (%)						
	C	random	I3	I4	I5	R1	R2
10	39.8	24.6 (-38.2)	30.0 (-24.6)	29.4 (-26.1)	30.0 (-24.6)	33.6 (-15.6)	35.9 (-9.6)
20	36.8	23.7 (-35.6)	27.2 (-26.1)	26.6 (-27.7)	27.3 (-25.8)	32.3 (-12.2)	34.4 (-7.9)
30	33.4	22.3 (-33.1)	24.9 (-25.4)	24.3 (-27.2)	24.9 (-25.4)	30.8 (-7.8)	31.9 (-4.6)
100	26.4	15.0 (-43.1)	16.9 (-35.9)	16.2 (-38.7)	16.9 (-35.9)	22.8 (-13.6)	23.5 (-10.8)
200	21.1	10.4 (-50.5)	11.7 (-44.7)	11.1 (-47.3)	11.7 (-44.7)	17.1 (-19.2)	18.1 (-14.5)

(b) 50 Unreplicated queries

Table 5: Effectiveness of different ranking functions using short queries on the 2 GB TREC2+3 collection (replicas built with top 200 documents)

effectiveness (compare I3-I5). When we set k to 200 (the default of the InQuery collection ranking function) and increase the value of b , the replica selector directs more queries to the replicas. For the replicated queries, the default InQuery collection ranking function ($k=200, b=0.75$) only directs 30% of queries to the replicas, which is not our choice. When we tune the parameters to $k=200$ and $b=1$, the function directs 89% of queries to the replicas.

For the replica selection function, $k = 2$ gets better results than $k = 1$ and $k = 4$ (compare the functions R3-R5). When we decrease the value of b , the replica selector directs more queries to the replicas. For $k=2$ and $b=0.2$, the function directs 88% of queries to the replicas.

Among the functions listed in Table 4, six functions *random*, I3, I4, I5, R1, and R2 direct more than 80% of replicated queries to the replicas. We compare the precision of these six functions in Table 5. The first column lists the number of documents at which we present the precision. Column 2 lists the precision when all queries go to the original collection, i.e., what percent of the top m documents is relevant when searching the original collection. Columns 3 through 8 list the results using random selection and each ranking function. The numbers in parentheses show the precision percentage difference as compared with searching the original collection. Table 5(a) lists the results for replicated queries, and Table 5(b) lists the results for unreplicated queries. Replicated queries produce much better results than unreplicated queries, because their top documents are stored in at least one of the replicas.

It is not surprising that random selection performs poorly, because it has high probability of picking a replica with few relevant documents. For unreplicated queries, it causes a precision percentage loss ranging from 38% to 50% as compared with searching the original collection, C. For replicated queries, it causes a

precision loss ranging from 17% to 25%.

For the other five functions in Table 5, when we examine the precision for replicated queries (Table 5(a)), all these functions are acceptable, since the precision drops less than 8.7%. However, when we examine the precision for unreplicated queries (Table 5(b)), the precision difference is significant. Using InQuery collection ranking function **I3** where we set $k = 200$ and $b = 1$, the precision loss of unreplicated queries range from 24.6% to 44.7%. We get our best result using our replica selection function **R2** with $k = 2$ and $b = 0.2$. The precision of the replicated queries drops less than 3.2% of the original collection, and is better when fewer documents are returned. The precision loss of the unreplicated queries range from 4.8% to 14.5%. For the top 30 documents, the precision loss of unreplicated queries range from 4.8% to 9.6%.

In the remaining experiments, the replica selector uses the replica ranking function with $k = 2$ and $b = 0.2$, because it sends appropriate queries to replicas with an acceptable precision loss of at most 9.6% for the top 30 documents in this test suit.

5.4 Effectiveness with Replicated Queries

This section evaluates our proposed replica selection function for replicated queries on a wider range of queries and collections. For replicated queries, we want to test whether the replica selector directs most of them to an expected replica. Note it is possible for a replica smaller than the expected one to contain all top documents for a given query, since the top documents of other queries could include the top documents for this query. Although we use 250 queries to build replicas, we only present the results for 99 replicated queries which have relevance judgment files in this section.

Finding the Expected Relevant Replica

This section measures the ability of the replica selector to pick the expected replica by counting the number of queries that are directed to different replicas and the original collection, as shown in Table 6. In Table 6, columns 1 and 2 indicate the size of collection and the type of queries we use in our experiments. Column 3 indicates the number of top documents for each query. The remaining columns are the same as Table 4(a).

For short queries on the 2 GB collection, on average, our replica selector directs 85.6% ($59.0\% + 7.7\% + 18.9\%$) of replicated queries to the replicas, and 66.7% of queries to the expected replica or a replica smaller than we expect. Increasing the number of replicated documents increases the accuracy of replica selection, because the replicas contain more relevant documents for replicated queries. For example, using the top 500 documents for each query to build replicas, the replica selector directs 90.3% of queries to the replicas on the average, while using top 100 documents directs 86.1% of queries to the replicas on average.

For long queries on the 2 GB collection, on average, our replica selector directs 89.0% ($57.4\% + 8.3\% +$

Size	Query Type	Top n	Average Num. of Queries to Replica					C	% to Replica			C
			D ₁	D ₂	D ₃	D ₄	D ₅		right	smaller	larger	
	Expected		21.6	16.6	21.4	20.6	18.8	0				
2 GB	short	30	16.0	11.6	13.2	13.2	24.6	20.4	52.7%	4.2%	22.6%	20.6%
		100	17.2	13.0	15.8	15.2	24.0	13.8	58.4%	5.9%	21.8%	13.9%
		200	20.0	13.2	14.2	17.0	21.0	13.6	59.8%	8.1%	18.4%	13.7%
		500	25.0	14.4	15.8	18.6	15.6	9.6	65.1%	12.7%	12.5%	9.7%
		Ave.	19.5	13.0	14.8	16.0	21.3	14.3	59.0%	7.7%	18.9%	14.4%
2 GB	long	100	15.8	13.4	13.2	13.2	25.8	17.6	52.7%	5.8%	23.6%	17.8%
		200	18.0	17.4	12.6	14.4	28.2	8.4	57.0%	8.5%	26.0%	8.5%
		500	21.8	17.6	14.0	15.0	24.0	6.6	62.4%	10.7%	20.2%	6.7%
		Ave.	18.5	16.1	13.3	14.2	26.0	10.9	57.4%	8.3%	23.3%	11.0%
20 GB	short	100	15.6	14.2	12.8	15.8	20.2	20.4	54.3%	4.2%	21.0%	20.6%
		200	15.4	13.2	12.0	15.4	24.6	18.4	56.5%	4.2%	20.6%	18.6%
		500	18.2	14.4	12.2	16.0	25.8	12.4	64.2%	4.2%	19.0%	12.5%
		Ave.	16.4	13.9	12.3	15.7	23.5	17.1	58.3%	4.2%	20.2%	17.3%

Table 6: Replica selection for replicated queries

23.3%) of replicated queries to the replicas, and 65.7% of queries to the expected replica or a replica smaller than we expect. Increasing the number of replicated documents also increases the accuracy of replica selection, as for the short queries.

For short queries on the 20 GB collection, on average, our replica selector directs 82.7% (58.3% + 4.2% + 20.2%) of replicated queries to the replicas, and 62.5% of queries to the expected replica or a replica smaller than we expect. Increasing the number of replicated documents increases the accuracy of replica selection, as against the 2 GB collection.

Precision of Replica Selection versus the Original Collection

Since the replica selector directs a few queries to a replica that is smaller than expected, we compare the effectiveness of executing queries against replicas or the original collection selected by the replica selector with against the original collection. Table 7 compares the average precision of replica selection over 5 test trials with searching the original collection for short queries on the 2 GB collection, long queries on the 2 GB collection, and short queries on the 20 GB collection. In these tables, column 1 lists the number of documents at which we present the precision figures. Column 2 lists the precision figures when all queries go to the original collection. Columns 3 through 6 list the precision figures when building replicas using different numbers of top documents. The numbers in the parentheses show the precision percentage difference.

For short queries on the 2 GB collection, replica selection results in a precision percentage loss less than 3.1% of searching the original collection for the same number of responses or fewer. For long queries on the 2 GB collection, replica selection results in a precision percentage loss less than 2.4%.

at <i>m docs</i>	Precision				
	orig.	Top 30	Top 100	Top 200	Top 500
10	47.3	46.9 (-0.8)	47.0 (-0.6)	47.4 (+0.3)	46.9 (-0.8)
20	43.5	43.0 (-1.2)	43.0 (-1.1)	43.3 (-0.4)	42.9 (-1.3)
30	39.6	39.0 (-1.5)	39.1 (-1.3)	39.4 (-0.7)	39.2 (-0.9)
100	30.8		29.9 (-2.7)	30.1 (-2.0)	30.1 (-2.1)
200	24.7			24.0 (-3.0)	24.0 (-3.0)
500	16.5				16.0 (-3.1)

(a) short queries on the 2 GB collection

at <i>m docs</i>	Precision			
	orig.	Top 100	Top 200	Top 500
10	56.3	56.1 (-0.4)	56.4 (+0.1)	56.2 (-0.2)
20	54.6	54.2 (-0.7)	54.3 (-0.6)	54.1 (-0.9)
30	51.7	51.3 (-0.8)	51.1 (-1.1)	51.2 (-1.0)
100	41.5	41.1 (-1.1)	41.0 (-1.2)	41.0 (-1.1)
200	34.1		33.4 (-2.1)	33.6 (-1.6)
500	22.9			22.4 (-2.4)

(b) long queries on the 2 GB collection

at <i>m docs</i>	Precision			
	orig.	Top 100	Top 200	Top 500
10	15.5	15.5 (-1.2)	15.4 (-1.2)	15.5 (-0.3)
20	15.0	15.0 (-0.3)	15.0 (-0.5)	15.0 (+0.0)
30	14.0	14.0 (+0.0)	14.0 (-0.1)	14.0 (+0.2)
100	11.5	11.4 (-0.9)	11.5 (-0.6)	11.5 (-0.3)
200	9.8		9.7 (-0.9)	9.7 (-0.1)
500	7.5			7.4 (-0.8)

(c) short queries on the 20 GB collection

Table 7: Effectiveness of replica selection for replicated queries (each trial has 99 judged queries)

For short queries on the 20 GB collection, replica selection results in a precision percentage loss less than 1.2% as compared to searching the original collection, and sometimes the precision improves a little, because the replica does not contain some top-ranked irrelevant documents. In other words, selecting a smaller replica occasionally does no harm.

5.5 Effectiveness with Unreplicated Queries

This section evaluates our proposed replica selection function on a wider range of queries and collections for unreplicated queries. See Section 5.2 for detailed experimental setting.

Finding the Relevant Replica

Table 8 lists the average expected number of collection precise queries in each replica over five test trials and shows the results of replica selection by collecting the average number of collection precise queries that are directed to different replicas as well as the original collection. We list results for short queries on the 2 GB TREC 2+3 collection, long queries on the 2 GB TREC 2+3 collection, and short queries on the 20 GB

Size	Top Type	Query n	Coll. Precise queries	Precision loss less than 5%					
				Ave. Queries Expected					C
				D ₁	D ₂	D ₃	D ₄	D ₅	
2 GB	short	30	38.2	2.6	1.0	2.2	1.2	1.4	29.8
		100	37.8	3.8	2.8	3.2	2.2	1.0	24.8
		200	37.8	4.6	3.8	2.6	3.2	1.2	22.4
		500	37.8	6.2	3.8	4.4	3.4	1.4	18.6
		Ave.	37.9	4.3	2.9	3.1	2.5	1.3	23.9
2 GB	long	100	42.4	3.6	3.4	3.4	2.2	1.2	28.6
		200	42.4	5.2	3.8	3.4	2.2	2.0	25.8
		500	42.4	8.2	5.4	4.0	2.8	2.0	20.0
		Ave.	42.4	5.7	4.2	3.6	2.4	1.7	24.8
		20 GB	short	100	18.4	0.8	0.6	0.8	1.0
200	19.0			0.4	1.0	1.2	1.2	0.6	14.6
500	19.2			2.2	2.0	2.0	1.0	1.0	11.0
Ave.	18.9			1.1	1.2	1.3	1.1	0.5	13.7
20 GB	301-350 short			100	36	0	1	1	0
		200	35	0	0	1	0	4	30
		500	35	0	1	0	1	4	29
		Ave.	35.3	0	0.6	0.6	0.3	3.7	30.0

(a) Expected number of collection precise queries in each replica

Size	Query Type	Top n	Ave. Queries to Replica					C	Precision Loss			% of Repl.Prec. queries to C
			D ₁	D ₂	D ₃	D ₄	D ₅		C+< 5%	5% - 10%	> 10%	
2 GB	short	30	2.6	2.0	2.0	2.8	5.6	23.2	72.4%	7.4%	16.2%	27.9%
		100	2.6	3.0	3.2	3.4	8.2	17.4	70.5%	12.2%	17.4%	13.8%
		200	4.0	3.4	2.4	4.8	6.8	16.4	71.5%	11.2%	17.3%	15.3%
		500	4.8	3.0	4.6	6.4	5.8	13.2	66.2%	22.7%	11.1%	15.4%
		Ave.	3.5	2.8	3.1	4.4	6.6	17.5	70.2%	13.4%	15.5%	18.1%
2 GB	long	100	3.6	2.4	1.8	2.8	6.6	25.2	77.5%	13.1%	9.4%	25.9%
		200	4.8	4.0	1.8	3.0	11.2	17.6	70.6%	17.6%	11.8%	13.1%
		500	5.4	4.8	3.2	5.0	10.2	13.8	78.5%	17.4%	4.2%	5.7%
		Ave.	4.6	3.7	2.3	3.6	9.3	18.9	75.5%	16.0%	8.5%	14.9%
		20 GB	short	100	0.4	0.4	0.6	0.8	2.2	14.0	84.7%	4.5%
200	0.6			0.6	1.2	1.4	1.8	13.4	84.2%	6.3%	9.5%	31.6%
500	0.4			0.8	1.8	1.2	2.6	12.4	86.4%	6.2%	7.3%	29.0%
Ave.	0.5			0.6	1.2	1.1	2.2	13.3	85.1%	5.7%	9.2%	32.5%
20 GB	301-350 short			100	2	1	0	1	1	31	86.1%	0.0%
		200	2	1	0	1	1	30	85.7%	0.0%	14.3%	40.0%
		500	1	1	0	1	2	30	85.7%	5.8%	8.5%	40.0%
		Ave.	1.7	1.0	0.0	1.0	1.3	30.3	85.8%	1.9%	12.2%	43.3%

(b) Results of replica selection for collection precise queries

Table 8: Replica selection for unreplicated queries

TREC VLC collection. In Table 8, columns 1 and 2 indicate the size of collection and the type of the query sets. Column 3 indicates the number of documents stored for each query. The remaining columns are the same as Table 4(b).

For short queries on the 2 GB collection, on the average, our replica selector directs 83.6% (70.2% + 13.4%) of collection precise queries to the replicas that cause a precision loss less than 10% (our acceptable level) as well as the original collection, and only directs 18.1% of queries which are replica precise to the original collection.

For long queries on the 2 GB collection, on the average, our replica selector directs 91.5% (75.5% + 16.0%) of collection precise queries to the replicas that cause a precision loss less than 10% as well as the original collection, and only directs 14.9% of replica precise queries to the original collection.

For short queries on the 20 GB collection, when we experiment with the same setting as the 2 GB collection, on the average, our replica selector directs 90.8% (85.1%+5.7%) of collection precise queries to the replicas that cause a precision loss less than 10% as well as the original collection. When we experiment with queries 301-350 as our unreplicated queries, our replica selector directs 87.7% (85.8%+1.9%) of collection precise queries to the replicas that cause a precision loss less than 10% as well as the original collection.

Precision of Replica Selection versus the Original Collection

We compare the retrieval precision of executing unreplicated queries against replicas or the original collection selected by our replica selector with only searching the original collection.

Table 9 lists average precision over 5 test trials for short queries on the 2 GB TREC 2+3 collection, long queries on the 2 GB TREC 2+3 collection, and short queries on the 20 GB TREC VLC collection.

For the 2 GB collection using short queries, the precision losses range from 6.8% to 17.1%. Increasing the number of replicated documents for each query improves the precision, because the replicas contain more relevant documents for each replicated query, which helps determining the similarity between unreplicated and replicated queries. When the number of top retrieved documents is less than 30 documents, which are the retrieval levels that concern online users most, our replica selector causes an average precision percentage loss within 14.6% and 10% of searching the original collection, when we only replicate the top 30 documents and the top 100 documents for each replicated query, respectively.

For the 2 GB collection using long queries, the precision losses range from 0.9% to 17.4%. For the top 30 retrieved documents, on the average, the precision drops less than 8.7% when we replicate more than 100 documents for each replicated queries, which is slightly better than short queries.

at m docs	Precision				
	orig.	Top 30	Top 100	Top 200	Top 500
10	42.8	37.8 (-11.9)	38.9 (-9.2)	39.4 (-8.0)	39.9 (-6.8)
20	39.4	34.0 (-13.8)	35.8 (-9.1)	36.1 (-8.4)	35.6 (-9.7)
30	35.5	30.3 (-14.6)	32.6 (-8.2)	32.7 (-7.8)	32.7 (-7.9)
100	27.2		23.3 (-14.0)	24.0 (-11.6)	24.2 (-10.8)
200	21.8			18.3 (-16.5)	18.8 (-13.9)
500	14.3				11.8 (-17.1)

(a) short queries on the 2 GB collection

at m docs	Precision			
	orig.	Top 100	Top 200	Top 500
10	55.3	52.9 (-4.3)	52.0 (-6.0)	54.8 (-0.9)
20	52.7	49.4 (-6.2)	48.6 (-7.8)	50.9 (-3.4)
30	50.0	46.2 (-7.6)	45.7 (-8.7)	47.9 (-4.3)
100	40.4	35.3 (-12.6)	35.1 (-13.2)	36.8 (-8.8)
200	33.1		27.3 (-17.4)	29.1 (-12.0)
500	21.8			18.2 (-16.6)

(b) long queries on the 2 GB collection

at m docs	Precision			
	orig.	Top 100	Top 200	Top 500
10	12.8	12.4 (-3.1)	12.6 (-1.6)	12.4 (-3.4)
20	12.3	11.6 (-5.5)	11.9 (-3.1)	11.8 (-3.4)
30	11.8	11.4 (-3.1)	11.9 (+0.8)	11.8 (+0.3)
100	10.0	9.1 (-9.6)	9.6 (-4.2)	10.1 (+0.2)
200	8.4		7.7 (-7.9)	8.3 (-1.0)
500	6.4			5.9 (-7.7)

(c) short queries on the 20 GB collection

at m docs	Precision			
	orig.	Top 100	Top 200	Top 500
10	40.4	36.2 (-10.4)	36.2 (-10.4)	37.8 (-6.4)
20	35.4	30.7 (-13.3)	30.7 (-13.3)	31.3 (-11.6)
30	31.3	26.9 (-14.2)	26.9 (-14.2)	27.0 (-14.0)
100	20.2	17.8 (-11.9)	17.8 (-11.9)	17.2 (-15.0)
200	14.4		12.8 (-10.9)	12.2 (-14.0)
500	7.8			6.7 (-14.0)

(d) short queries (topics 301-350) on the 20 GB collection

Table 9: Effectiveness of unreplicated queries (each trial has 50 queries)

For the 20 GB collection using short queries, when we experiment with the same setting as the 2 GB collection and use the relevance files for the 2 GB collection, the precision ranges from losing 9.6% to improving 0.8%. For the top 30 retrieved documents, the precision loss is less than 5.5%. When we use short queries 301-350 as our unreplicated queries, the precision loss for the top 30 documents is less than 14.2%. Since topics 301-350 were much more thoroughly judged than topics {51-150, 202-250} for the 20 GB VLC collection, although still only the top 30 documents of each query were judged, we think the results using topics 301-350 are more accurate, which means our replica selection performs slightly worse on the 20 GB collection than on the 2 GB collection. However, the precision percentage loss of 14.2% in our context only means we retrieve one less relevant document for the top 30 documents.

5.6 Summary

We investigated how to select a relevant partial replica using the inference network model. Our approach enables a system to efficiently rank partial replicas and select a replica based on relevance for a given query. We developed a replica selection function, and demonstrated its effectiveness and superiority over a collection ranking function using the InQuery retrieval system and TREC collections. Our results show that the inference network model is a very promising tool for selecting a relevant replica. By using our proposed replica selection function, the replica selector can direct at least 82% of replicated queries to a relevant partial replica rather than the original collection, and it achieves a precision percentage loss less than 10% for the 2 GB collection and 14.2% for the 20 GB collection for the top 30 retrieved documents of each query, when we build replicas using more than 100 top documents for each replicated query.

6 Performance of Partial Replication for Searching a Terabyte of Text

In this section, we further describe the architecture of our distributed IR system, and compare the performance of partial replication with collection partitioning when searching a terabyte of text.

6.1 Our Distributed Information System

In our distributed IR system illustrated in Figure 5, clients, InQuery servers, and the connection broker reside on different machines. Clients are user interfaces to the retrieval system. InQuery servers store the original collection and partial replicas, and perform IR service such as query evaluation, obtaining summaries, and document retrieval. A collection or a replica may be distributed over several InQuery servers. The connection broker keeps track of all the InQuery servers for replicas or otherwise, outstanding client requests, and organizes response from InQuery servers. For partial replication, the connection broker also performs replica selection based on both relevance and load.




Figure 5: Our Distributed Information Retrieval System

In this paper, we focus on three basic IR commands: *query*, *summary*, and *document* commands. A **query command** consists of a set of words or phrases (terms), such as “information retrieval”, or “distributed system.” Query responses consist of a list of document identifiers ranked by belief values which estimate the probability that the document satisfies the information need. For each query, a client may obtain one or several summaries on relevant documents by sending **summary commands** which consist of a set of document identifiers and their collection identifiers. The summary information of a document typically consists of the title and the most relevant passages in the document. It may also include information such as source and organization. A client may also retrieve complete documents by sending a **document command** which consists of a document identifier and a collection identifier. In response, the system returns the complete text of the document from the collection.

When a client sends a query to the connection broker, the connection broker first uses a replica selector to determine whether there is a partial replica that is not only relevant to the query, but also not overloaded. If there is one, the connection broker sends the query to the InQuery server(s) that maintain the relevant replica, otherwise it sends the query to the InQuery servers that maintain the original collection. After each involved InQuery servers returns the results, the connection broker merges results and returns them to the client. For a summary command, the connection broker sends the command to the InQuery servers whose identifiers are described in the command. The connection broker merges the summary information responses and sends a single message back to the client. For a document command, the connection broker sends the command to the InQuery server that contains the document, and then forwards the document to the client as soon as it receives the document from the InQuery server.

In our system, if query locality is high, the replica selector may send too many queries to a replica which results in load imbalance. We load balance by predicting the response time of each replica and the




Figure 6: Performance validation of simulator with partial replication.

original collection using the average response time and the number of the outstanding commands. When the replica selector chooses a replica based on relevance, we calculate the predicted response time p_resp_j of the replica, any larger replica, and the original collection using $ave_resp_j \cdot (1 + num_wait_mes_j)$, where ave_resp_j is the average response time for last 200 responses for either the replica or the original collection, and $num_wait_mes_j$ is the number of the outstanding commands to which neither the original collection or the replica have responded. We send the command to the one with the least p_resp_j . The connection broker obtains information on the response time as it receives queries responses and tracks the number of outstanding messages.

We evaluate the performance of our distributed information retrieval system using a simulator with a performance model that is driven by measurements obtained using InQuery running on DEC Alpha Server 2100 5/250 with 3 CPUs (clocked at 250 MHz) and 1024 MB main memory, running Digital Unix V3.2D-1 (Rev 41). Servers are connected by a 10 Mbps Ethernet. In previous work, we showed the simulator closely matches a multithreaded implementation of InQuery [Cahoon et al., 1999, Lu, 1999]. In addition, we report on the validation of some of our simulation results below, comparing partitioning and replication with varying degrees of locality for a 16GB collection on a single server, and again our measured times closely match our simulator. Of course, simulation enables us to explore in a controlled environment high loads and very large configurations.

6.2 Validation of Partial Replication Performance

This section compares the simulator and an implementation of partial replication for searching a 16 GB collection on a multi-tasking server using InQuery 3.1 as the query arrival rate increases on a 3-CPU Alpha Server 2100 5/250 running Digital UNIX V3.2D-1 (Rev 41). We used a multi-tasking server instead of

Parameters	Abbre.	Values
Num. of Commands	N_{com}	1000
Command Arrival Rate		0.1 2 4 6 8 10
Poisson dist. (avg. commands/sec)	λ	12 14 16 18 20
Command Mixture Ratio query:summary:document	R_{cm}	1:1.5:2
Terms per Query (average) shifted neg. binomial dist.	N_{tpq}	2
Query Term Frequency dist. from queries	D_{qtf}	Obs. Dist.
Data per Server	S_{size}	32 GB
Size of Collection	C_{size}	1 TB
Replication Percentage	P_{repl}	3% (32 GB)
Distracting Percentage	P_{disp}	10% - 100% by 10

Table 10: Configuration Parameters for Terabyte Experiments

a multithreaded server just to save us time from implementing replica selection in our legacy system, In-Query which uses too many global variables. Our earlier work showed that the multitasking server performs similarly to the multithreaded server, although the multithreaded server is always slightly faster (90% of measured response times fall within 10% of each other) [Lu et al., 1998].

In this experiment, we distribute a 16 GB collection over 4 disks and used an extra disk to store a 4 GB replica. We assume queries arrive as a Poisson process, and use 50 short queries with average of 2 terms per query. Figure 6 compares the performance of using the real system and the simulator when the replica distracts 40% of commands, and shows that two systems present the same trends and expected improvements from partial replication.

6.3 Searching a Terabyte of Text

In this section, we compare the simulated performance of partial replication with collection partitioning using one and a hierarchy of replicas. We model command arrival as a Poisson process. We use short queries with an average of 2 terms per query, and set the ratio of query commands, summary commands, and document commands to 1:1.5:2, as we observed in the THOMAS log. We assume each server stores a 32 GB collection. As our baseline, we use 32 servers to store 1 terabyte of text. These experiments use a 32 GB replica, which is sufficient to satisfy more than 40% of queries in the Excite log. We vary the **distracting percentage** which represents the percentage of queries that the replica selector sends to a partial replica. Table 10 presents the experimental parameters, their abbreviations, and values.

Partial Replication versus Collection Partitioning

In this section, we compare the performance of the following configurations with the baseline (partitioning over 32 servers):


- Partitioning over additional servers: partitioning 1 TB of text over 33, and 64 servers, each of which stores 31 GB and 16 GB of data.
- Partial Replication: building a 32 GB replica on one additional server (33 total servers).

Figure 7 illustrates the average query response time when we partition 1 TB of text over 32, 33, and 64 servers, and over 32 servers plus one server that contains a 32 GB replica. We vary the distracting percentage which indicates the percent of queries sent to the replica. Figure 7(a)-(c) illustrate when commands arrive at 4, 10, and 20 commands per second. The graphs plot query response time versus the distracting percentage. When we have one additional server, using it to store a replica performs significantly better than further partitioning over this server, especially when the commands arrive at a high rate, as shown in Figure 7(c). The improvement occurs when the replica satisfies only 3% of commands for more highly loaded systems, e.g., 10 and 20 commands per second, and the improvement increases with increases in query locality. Using one partial replica also performs similar or better than partitioning over twice as many as servers when the replica satisfies at least 20% of commands. For example, when the arrival rate is 20 commands per second, partitioning over 64 servers reduces the average query response time by a factor of 1.6, while one partial replica reduces it by a factor of 2.3 when the replica satisfies 40% of commands. When the distracting percentage becomes high, the replica selector load balances between the replica and the partitioned original collection which maintains retrieval effectiveness and quick response times.

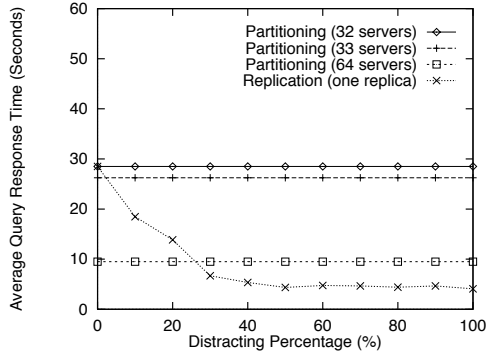
There are two major reasons that partial replication outperforms partitioning. (1) A server takes around $3/5$ the time to search half the data according to our measurements, and thus when we partition a terabyte of text over 64 servers instead of 32 servers, each server can not process twice as many commands as using 32 servers. (2) Searching a replica results in less network traffic and needs less coordination of the results from each partition in the connection broker. For example, the utilizations of the network and the connection broker for partitioning over 64 servers are 28% and 70%, while the corresponding utilizations for using 32 servers and one partial replica is 12% and 58%. For highly loaded systems, replication significantly improves performance over partitioning and uses only about half of the resources!

6.4 Partial Replication as a Hierarchy


In this section, we assume 1, 2, and 4 additional servers and organize them as a hierarchy of replicas. We examine how much improvement a hierarchy of replicas will produce. We assume the first, second, third, and fourth additional server stores 32 GB, 16 GB, 8 GB and 4 GB of data. where $R_1 \supset R_2 \supset R_3 \supset R_4$ and R_i represents the data on the i -th server. The replicas satisfy accurately a total of $p\%$ of commands. Of these $p\%$ of commands the replica selector sends to replicas, all are satisfied by the largest replica, 10% less,



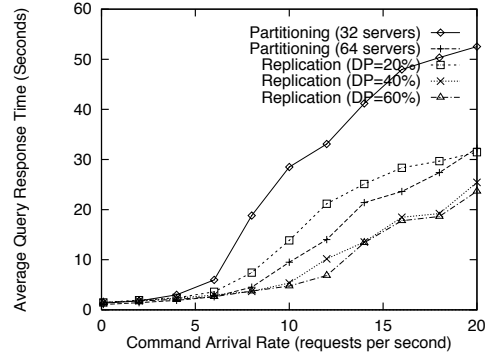
(a) $\lambda = 4$



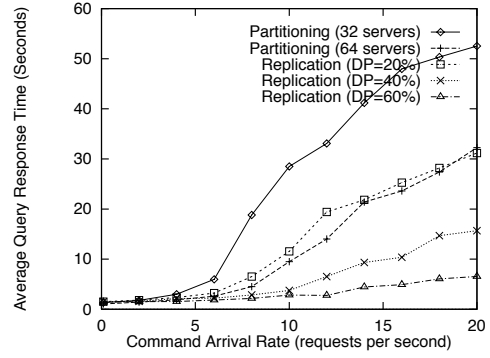
(b) $\lambda = 10$




(c) $\lambda = 20$



(a) one replica



(b) two replicas



(c) four replicas

Figure 7: Comparing partial replication with partitioning

Figure 8: Performance using a hierarchy of replicas

i.e., $(p\% - 10\%)$ are satisfied by the second largest replica, another 10% less, i.e., $(p\% - 20\%)$ by the third largest replica, and $(p\% - 30\%)$ by the fourth replica, where the commands sent to a i -th largest replica may also be satisfied at the $(i + 1)$ -th largest replica.

Figure 8 illustrates the average query response time when we build one and two replicas, where the replicas distract 20%, 40%, and 60% of commands, as well as four replicas, where the replicas distract 40% and 60% of commands. The results show that 2 replicas are sufficient to achieve large performance improvements beyond partitioning when the replicas satisfy 40% and 60% of commands. In our baseline, partitioning over 32 servers achieves an average query response time below 10 seconds at 7 commands per second. Using one replica to satisfy 20% of commands and using two replicas to satisfy 40% and 60% of commands achieve average query response time below 10 seconds at 9, 16, and more than 20 commands per second, respectively, while partitioning over 64 servers (using 32 additional servers) only achieves average query response time below 10 seconds at 10 commands per second.

Thus, for our system (slower than the current state of the art, unfortunately!), we achieve query response times under 10 seconds for a relatively highly loaded system with 20 requests per second using 4 replicas and query locality of about 50%. With a faster base system, replication is still preferable to partitioning given even modest query locality, however fewer replicas are probably necessary to maintain fast response times.

7 Conclusions

In this paper, we investigated how to search a terabyte of text using partial replication. We build a hierarchy of replicas based on query frequency and available resources, and use the InQuery retrieval system for the replicas and the original collection. We examine queries from THOMAS [THOMAS, 1998] and Excite [Excite, 1997] to find locality patterns in real systems. We find there is sufficient query locality that remains high over long periods of time which will enable partial replication to maintain effectiveness and significantly improve performance. For THOMAS, updating replicas hourly or even daily is unnecessary. However, we need to some mechanism to deal with bursty events. We propose two simple updating strategies that trigger updates based on events and performance, instead of regular updating. In our traces, query exact match misses many overlaps between queries with different terms that in fact return the same top documents, whereas partial replication with an effective replica selection function will find the similarities. We believe this trend will hold for other query sets against text collections and for web queries.

We investigate how to select a relevant partial replica using the inference network, and demonstrate the effectiveness of our approach using the InQuery retrieval system and TREC collections. The results show that the inference network model is a very promising approach for ranking partial replicas. By using our

new replica selection function, our replica selector can direct at least 82% of replicated queries to a relevant partial replica rather than the original collection, and it achieves a precision percentage loss within 10% and 14.2% for the top 30 retrieved documents for those unreplicated queries, when sizes of replicas range from 2% to 10% for the 2 GB collection, and 0.2% to 1% for the 20 GB collection, respectively.

We demonstrate the performance of our system searching a terabyte of text using a validated simulator. We compare the performance of partial replication with partitioning over additional servers. Our results show that partial replication is more effective at reducing execution times than partitioning on significantly fewer resources. For example, using 1 or 2 additional servers for replica(s) achieves similar or better performance than partitioning over 32 additional servers, even when the largest replica satisfies only 20% of commands. Higher query locality further widens the performance differences. Our work porting and validating InQuery and the simulator from a slower processor to the Alpha, as well as experiments with faster querying times which are reported elsewhere [Lu, 1999], lead us to believe the performance trends will hold for faster systems using fewer resources.

Acknowledgments

This material is based on work supported in part by the National Science Foundation, Library of Congress and Department of Commerce under cooperative agreement number EEC-9209623, supported in part by United States Patent and Trademark Office and Defense Advanced Research Projects Agency/ITO under ARPA order number D468, issued by ESC/AXS contract number F19628-95-C-0235, and also supported in part by grants from Digital Equipment, NSF grant EIA-9726401, and an NSF Infrastructure grant CDA-9502639. Kathryn S. McKinley is supported by an NSF CAREER award CCR-9624209. We thank Ben Mealey and Library of Congress for providing the THOMAS log. We thank Doug Cutting and Excite for providing the Excite log. Any opinions, findings and conclusions or recommendations expressed in this material are the authors and do not necessarily reflect those of the sponsors.

References

- [Ahamad and Ammar, 1989] Ahamad, M. and Ammar, M. (1989). Performance characterization of quorum-consensus algorithms for replicated data. *IEEE Transaction of Software Engineering*, 15(4).
- [Baentsch et al., 1996] Baentsch, M., Molter, G., and Sturm, P. (1996). Introducing application-level replication and naming into today's Web. In *Proceedings of Fifth International World Wide Web Conference*, Paris, France.
- [Bestavros, 1995] Bestavros, A. (1995). Demand-based document dissemination to reduce traffic and balance load in distributed information systems. In *Proceedings of SPDP'95: The 7th IEEE Symposium on Parallel and Distributed Processing*, San Anotonio, Texas.
- [Burkowski, 1990] Burkowski, F. J. (1990). Retrieval performance of a distributed text database utilizing a parallel process document server. In *1990 International Symposium On Databases in Parallel and Distributed Systems*, pages 71–79, Trinity College, Dublin, Ireland.

- [Cahoon and McKinley, 1996] Cahoon, B. and McKinley, K. S. (1996). Performance evaluation of a distributed architecture for information retrieval. In *Proceedings of the Nineteenth Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 110–118, Zurich, Switzerland.
- [Cahoon et al., 1999] Cahoon, B., McKinley, K. S., and Lu, Z. (1999). Evaluating the performance of distributed architectures for information retrieval using a variety of workloads. *ACM Transactions on Information Systems* (accepted).
- [Callan et al., 1995a] Callan, J. P., Croft, W. B., and Broglio, J. (1995a). TREC and TIPSTER experiments with INQUERY. *Information Processing & Management*, 31(3):327–343.
- [Callan et al., 1992] Callan, J. P., Croft, W. B., and Harding, S. M. (1992). The INQUERY retrieval system. In *Proceedings of the 3rd International Conference on Database and Expert System Applications*, Valencia, Spain.
- [Callan et al., 1995b] Callan, J. P., Lu, Z., and Croft, W. B. (1995b). Searching distributed collections with inference networks. In *Proceedings of the Eighteenth Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Seattle, WA.
- [Chakravarthy and Haase, 1995] Chakravarthy, A. and Haase, K. (1995). Netserf: Using semantic knowledge to find internet information archives. In *Proceedings of the 18th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 4–11, Seattle, WA.
- [Couvreur et al., 1994] Couvreur, T. R., Benzel, R. N., Miller, S. F., Zeitler, D. N., Lee, D. L., Singhai, M., Shivaratri, N., and Wong, W. Y. P. (1994). An analysis of performance and cost factors in searching large text databases using parallel search systems. *Journal of the American Society for Information Science*, 7(45):443–464.
- [Croft et al., 1995] Croft, W. B., Cook, R., and Wilder, D. (1995). Providing government information on the Internet: Experiences with THOMAS. In *The Second International Conference on the Theory and Practice of Digital Libraries*, Austin, TX.
- [Danzig et al., 1991] Danzig, P. B., Ahn, J., Noll, J., and Obraczka, K. (1991). Distributed indexing: A scalable mechanism for distributed information retrieval. In *Proceedings of the Fourteenth Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 221–229, Chicago, IL.
- [Excite, 1997] Excite (1997). <http://www.excite.com>.
- [Fuhr, 1996] Fuhr, N. (1996). A decision-theoretic approach to database selection in networked ir. In *Workshop on Distributed IR*, Germany.
- [Gravano and Garcia-Molina, 1995] Gravano, L. and Garcia-Molina, H. (1995). Generalizing gloss to vector-space databases and broker hierarchies. In *Proceedings of the 21st VLDB Conference*, Zurich, Switzerland.
- [Gravano et al., 1994] Gravano, L., Garcia-Molina, H., and Tomasic, A. (1994). The effectiveness of gloss for the text database discovery problem. In *Proceedings of the SIGMOD 94*, pages 126–137.

- [Harman, 1997] Harman, D., editor (1997). *The Sixth Text REtrieval Conference (TREC-6)*. National Institute of Standards and Technology Special Publication, Gaithersburg, MD.
- [Harman et al., 1991] Harman, D., McCoy, W., Toense, R., and Candela, G. (1991). Prototyping a distributed information retrieval system that uses statistical ranking. *Information Processing & Management*, 27(5):449–460.
- [Hawking et al., 1998] Hawking, D., Craswell, N., and Thistlewaite, P. (1998). Overview of TREC-7 very large collection track. In *Proceedings of the 7th Text Retrieval Conference*.
- [Holmedahl et al., 1998] Holmedahl, V., Smaith, B., and Yu, T. (1998). Aperative caching of dynamic content on a distributed web server. In *IEEE Proceedings of 7th International Symposium on High Performance Distributed Computing (HPDC-7)*, pages 235–242, Chicago, IL.
- [Huang and Wolfson, 1993] Huang, Y. and Wolfson, O. (1993). A competitive dynamic data replication algorithm. In *IEEE Proceedings of 9th International Conference on Data Engineering*, pages 310–337, Vienna, Austria.
- [Lin and Zhou, 1993] Lin, Z. and Zhou, S. (1993). Parallelizing I/O intensive applications for a workstation cluster: A case study. *Computer Architecture News*, 21(5):15–22.
- [Lu, 1999] Lu, Z. (1999). *Scalable Distributed Architectures For Information Retrieval*. PhD thesis, University of Massachusetts at Amherst.
- [Lu et al., 1998] Lu, Z., McKinley, K. S., and Cahoon, B. (1998). The hardware/software balancing act for information retrieval on symmetric multiprocessors. In *Proceedings of Europar'98*, Southampton, U.K.
- [Macleod et al., 1987] Macleod, I. A., Martin, T. P., Nordin, B., and Phillips, J. R. (1987). Strategies for building distributed information retrieval systems. *Information Processing & Management*, 23(6):511–528.
- [Martin et al., 1990] Martin, T. P., Macleod, I. A., Russell, J. I., Lesse, K., and Foster, B. (1990). A case study of caching strategies for a distributed full text retrieval system. *Information Processing & Management*, 26(2):227–247.
- [Martin and Russell, 1991] Martin, T. P. and Russell, J. I. (1991). Data caching strategies for distributed full text retrieval systems. *Information Systems*, 16(1):1–11.
- [Oracle Company, 1995] Oracle Company (1995). Strategies and techniques for using Oracle7 replication. <http://www.oracle.com/products/servers/replication/html/collateral.html>.
- [THOMAS, 1998] THOMAS (1998). legislative information on the internet. <http://thomas.loc.gov>.
- [Voorhees et al., 1995] Voorhees, E. M., Gupta, N. K., and Johnson-Laird, B. (1995). Learning collection fusion strategies. In *Proceedings of the Eighteenth Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Seattle, WA.
- [Wolfson and Jajodia, 1992] Wolfson, O. and Jajodia, S. (1992). An algorithm for dynamic replication of data. In *Proceedings of 11th ACM Symposium on the Principles of Database Systems*, San Diego, California.

[Yu and MacNair, 1998] Yu, P. S. and MacNair, E. A. (1998). Performance study of a collaborative method for hierarchical caching in proxy servers. In *Proceedings of 7th International World Wide Web Conference*, Brisbane, Australia.