

Building a Distributed Full-Text Index for the Web

SERGEY MELNIK, SRIRAM RAGHAVAN, BEVERLY YANG
and HECTOR GARCIA-MOLINA
Stanford University

We identify crucial design issues in building a distributed inverted index for a large collection of Web pages. We introduce a novel pipelining technique for structuring the core index-building system that substantially reduces the index construction time. We also propose a storage scheme for creating and managing inverted files using an embedded database system. We suggest and compare different strategies for collecting global statistics from distributed inverted indexes. Finally, we present performance results from experiments on a testbed distributed Web indexing system that we have implemented.

Categories and Subject Descriptors: H.3.1 [**Information Storage and Retrieval**]: Content Analysis and Indexing—*Indexing methods*; H.3.4 [**Information Storage and Retrieval**]: Systems and Software—*Distributed systems*

General Terms: Design, Performance

Additional Key Words and Phrases: Distributed indexing, Text retrieval, Inverted files, Pipelining, Embedded databases

1. INTRODUCTION

Various access methods have been developed to support efficient search and retrieval over text document collections. Examples include *suffix arrays* [Manber and Myers 1990], *inverted files* or *inverted indexes* [Salton 1989; Witten et al. 1999], and *signature files* [Faloutsos and Christodoulakis 1984]. Inverted files have traditionally been the index structure of choice on the Web. Commercial search engines use custom network architectures and high performance

A shorter version of this paper will appear in the Proceedings of the 10th International WWW Conference, May 2001.

Authors' address: Stanford University, Gates Computer Science Department, Room 430, 353 Serra Hall, Stanford, Calif., 94305-9040; email: {melnik, rsram, byang, hector}@db.stanford.edu.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 2001 ACM 1046-8188/07/0100-0217 \$5.00

hardware to achieve sub-second query response times using such inverted indexes.¹

An inverted index over a collection of Web pages consists of a set of *inverted lists*, one for each occurring word (or *index term*). The inverted list for a term is a sorted list of *locations* where the term appears in the collection. A location consists of a page identifier and the position of the term within the page. When it is not necessary to track each term occurrence within a page, a location will include just a page identifier (and as an option, the number of occurrences within the page). Given an index term w , and a corresponding location l , we refer to the pair (w, l) as a *posting* for w .

Conceptually, building an inverted index involves processing each page to extract postings, sorting the postings first on index terms and then on locations, and finally writing out the sorted postings as a collection of inverted lists on disk. When the collection is small and indexing is a rare activity, optimizing index-building is not as critical as optimizing run-time query processing and retrieval. With a Web-scale index however, index build time also becomes a critical factor for two reasons.

Scale and growth rate. The Web is so large and growing so rapidly [Lawrence and Giles 1999; Inktomi 2000] that traditional build schemes become unmanageable, requiring huge resources and taking days to complete (and becoming more vulnerable to system failures). As a measure of comparison, the 40 million page (220 GB) WebBase repository [Hirai et al. 2000] represents only about 4% of the estimated size of the *publicly indexable Web* as of January 2000 [Inktomi 2000], but is already larger than the 100 GB *very large TREC-7 collection* [Hawking and Craswell 1998], the benchmark for large IR systems.

Rate of change. Since the content on the Web changes extremely rapidly [Cho and Garcia-Molina 2000], there is a need to periodically crawl the Web and update the inverted index. Indexes can either be updated incrementally or rebuilt periodically, after every crawl. With both approaches, the key challenge is to handle the large wholesale changes commonly observed between successive crawls of the Web. For efficiency and simplicity, most commercial Web search systems employ the rebuilding approach [M. Burrows, personal communication]. In this case, it is critical to build the index rapidly to provide access to the new data.

To study and evaluate index building in the context of the special challenges imposed by the Web, we have implemented a testbed system that operates on a cluster of *nodes* (workstations). As we built the testbed, we encountered several challenging problems that are typically not encountered when working with smaller collections. In this paper we report on some of these issues and the experiments we conducted to optimize build times for massive collections.

— We propose the technique of constructing a *software pipeline* on each indexing node to enhance performance through intra-node parallelism (Section 3).

¹Even though the Web link structure is being utilized to produce high quality results, text-based retrieval continues to be the primary method for identifying the relevant pages. In most commercial search engines, a combination of text and link-based methods are employed.




Fig. 1. Testbed architecture.

- We argue that the use of an embedded database system such as *The Berkeley Database* [Olson et al. 1999] for storing inverted files, has a number of important advantages. We propose an appropriate format for inverted files that makes optimal use of the features of such a database system (Section 4).
- Any distributed system for building inverted indexes needs to address the issue of collecting global statistics (e.g., *inverse document frequency—IDF*). We examine different strategies for collecting such statistics from a distributed collection (Section 5).
- For each of the above issues, wherever appropriate, we present experiments and performance studies to compare the alternatives.

We emphasize that the focus of this paper is on the process of building an inverted index, not on using this index to process search queries. As a result, we do not address issues such as ranking functions, relevance feedback, query expansion, [Salton 1989; Witten et al. 1999], and distributed query processing [Jeong and Omiecinski 1995; Tomasic and Garcia-Molina 1993b].

The focus of this paper is not on presenting a comprehensive performance or feature-list comparison of our testbed indexing system with existing systems for indexing Web and nonWeb collections. Rather, we use our experience with the testbed to identify some key performance issues in building a Web-scale index and propose generic techniques that are applicable to any distributed inverted index system.

2. TESTBED ARCHITECTURE

Our testbed system for building inverted indexes operates on a distributed shared-nothing architecture consisting of a collection of nodes connected by a local area network. We identify three types of nodes in the system (Figure 1):

Distributors. These nodes store the collection of Web pages to be indexed. Pages are gathered by a *Web crawler* and stored in a repository distributed across the disks of these nodes [Hirai et al. 2000].

Indexers. These nodes execute the core of the index building engine.

Query servers. Each of these nodes stores a portion of the final inverted index and an associated *lexicon*. The lexicon lists all the terms in the corresponding portion of the index and their associated statistics. Depending on the organization of the index files, some or all of the query servers may be involved in answering a search query.

Note that many traditional information retrieval (IR) systems do not employ such a 3-tier architecture for building inverted indexes. In those systems, the pages or documents to be indexed are placed on disks directly attached to the machines that build the index. A 3-tier architecture, however, provides significant benefits in the context of a Web search service. A Web search service must perform three resource intensive tasks—crawling, indexing, and querying—simultaneously. Even as existing indexes are used to answer search queries, newer indexes (based on a more recent crawl) must be constructed, and in parallel, the crawler must begin a fresh crawl. A 3-tier architecture clearly separates these three activities by executing them on separate banks of machines, thus improving performance. This ensures that pages are indexed and made available for querying as quickly as possible, thereby maximizing *index freshness*.

Overview of indexing process. The inverted index is built in two stages. In the first stage, each distributor node runs a *distributor process* that disseminates the collection of Web pages to the indexers. Each indexer receives a mutually disjoint subset of pages and their associated identifiers. The indexers parse and extract postings from the pages, sort the postings in memory, and flush them to intermediate structures on disk.

In the second stage, these intermediate structures are merged together to create one or more inverted files and their associated lexicons. An (inverted file, lexicon) pair is generated by merging a subset of the sorted runs. Each (inverted file, lexicon) pair is transferred to one or more query servers. In this paper, for simplicity, we assume that each indexer builds only one such pair.

Distributed inverted index organization. In a distributed environment, there are two basic strategies for distributing the inverted index over a collection of query servers [Martin et al. 1986; Ribeiro-Neto and Barbosa 1998; Tomasic and Garcia-Molina 1993b]. One strategy is to partition the document collection so that each query server is responsible for a disjoint subset of documents in the collection (called *local inverted files* by Ribeiro-Neto and Barbosa [1998]). The other option is to partition based on the index terms so that each query server stores inverted lists only for a subset of the index terms in the collection (called *global inverted files* by Ribeiro-Neto and Barbosa [1998]). Performance studies described in Tomasic and Garcia-Molina [1993b] indicate that for large collections, the local inverted file organization uses system resources effectively and provides good query throughput in most cases. We therefore employ the local inverted file organization in our testbed.

Testbed environment. Our indexing testbed uses the large repository of Web pages provided by the WebBase project [Hirai et al. 2000] as the test corpus for the performance experiments. The storage manager of the WebBase system receives pages from the Web crawler [Cho and Garcia-Molina 2000] and populates the distributor nodes. The indexers and the query servers are single




Fig. 2. Logical phases.

processor PCs with 350–500 MHz processors, 300–500 MB of main memory, and are equipped with multiple IDE disks. The distributor nodes are dual-processor machines with SCSI disks housing the repository. All the machines are interconnected by a 100 Mbps Ethernet LAN.

3. PIPELINED INDEXER DESIGN

The core of the indexing system is the *index-builder* process that executes on each indexer. The input to the index-builder process is a sequence of Web pages and their associated identifiers.² The output of the index-builder is a set of *sorted runs*. Each sorted run contains postings extracted from a subset of the pages received by the index-builder.

The process of generating these sorted runs can logically be split into three phases, as illustrated in Figure 2. We refer to these phases as *loading*, *processing*, and *flushing*. During the loading phase, some number of pages are read from the input stream. The processing phase involves two steps. First, the pages are parsed to remove HTML tagging, tokenized into individual terms, and stored as a set of postings in a memory buffer. In the second step, the postings are sorted in place, first by term, then by location. During the flushing phase, the sorted postings in the memory buffer are saved on disk as a sorted run. These three phases are executed repeatedly until the entire input stream of pages has been consumed.

Loading, processing and flushing tend to use disjoint sets of system resources. Processing is obviously CPU-intensive, whereas flushing primarily uses secondary storage, and loading can be done directly from the network, tape, or a separate disk. Indexing performance can, therefore, be improved by executing these three phases concurrently. Since the execution order of loading, processing and flushing is fixed, these three phases together form a *software pipeline*.

Figure 3 illustrates the benefits of pipelined parallelism during index construction. The figure shows a portion of an indexing process that uses three concurrent threads, operates on three reusable memory buffers, and generates six sorted runs on disk.

The key issue in pipelining is to design an execution schedule for the different indexing phases that will result in minimal overall running time (called *makespan* in the scheduling literature). Our problem differs from a typical *job scheduling* problem [Chakrabarti and Muthukrishnan 1996] in that we can vary the sizes of the incoming *jobs*, that is, in every loading phase we can

²The URLs are normally replaced by numeric identifiers for compactness.




Fig. 3. Multi-threaded execution.

choose the number of pages to load. In the rest of this section, we describe how we make effective use of this flexibility. First we derive, under certain simplifying assumptions, the characteristics of an *optimal indexing pipeline schedule* and determine the theoretical speedup achievable through pipelining. Next, we describe experiments that illustrate how observed performance gains differ from the theoretical predictions.

3.1 Theoretical Analysis

Let us consider an indexer node that has one resource of each type—a single CPU, a single disk, and a single network connection over which to receive the pages. How should we design the pipeline shown in Figure 2 to minimize index construction time?

First, notice that executing concurrent phases of the same kind, such as two disk flushes, is futile, since we have only one resource of each type. Consider an index-builder that uses N executions of the pipeline to process the entire collection of pages and generate N sorted runs. By an *execution of the pipeline*, we refer to the sequence of three phases—loading, processing, and flushing—that transform some set of pages into a sorted run. Let B_i , $i = 1 \dots N$, be the buffer sizes used during these N executions. The sum $\sum_{i=1}^N B_i = B_{total}$ is fixed for a given amount of text input and represents the total size of all the postings extracted from the pages. Our aim is to come up with a way of choosing the B_i values so as to minimize the overall running time.

Loading and flushing take time linear in the size of the buffer. Processing time has a linear component (representing time for removing HTML and tokenizing) and a linear-logarithmic component (representing sorting time). Let $l_i = \lambda B_i$, $f_i = \varphi B_i$, and $p_i = \delta B_i + \sigma B_i \log B_i$ represent the durations of the loading, flushing, and processing phases for the i th execution

Table I. Measured Constants

Constant	Value
λ	1.26×10^{-3}
φ	4.62×10^{-4}
δ	6.74×10^{-4}
σ	2.44×10^{-5}

of the pipeline.³ For large N , the overall indexing time is determined by the scarcest resource (the CPU, in Figure 3) and can be approximated by $T_p = \max\{\sum_{i=1}^N l_i, \sum_{i=1}^N p_i, \sum_{i=1}^N f_i\}$.

It can be shown (see A) that T_p is minimized when all N pipeline executions use the same buffer size B , where $B = B_1 \cdots = B_N = \frac{B_{total}}{N}$. Let $l = \lambda B$, $f = \varphi B$, and $p = \delta B + \sigma B \log B$ be the durations of the loading, processing, and flushing phases, respectively. We must choose a value of B that maximizes the speedup gained through pipelining.

We calculate speedup as follows. Pipelined execution takes time $T_p = N \max(l, p, f)$ ($6p$ in Figure 3) and uses 3 buffers, each of size B . In comparison, sequential execution using a single buffer of size $3B$ will take time $T_s = \frac{N}{3}(l' + p' + f')$, where $l' = \lambda(3B)$, $f' = \varphi(3B)$, and $p' = \delta(3B) + \sigma(3B) \log(3B)$. Thus, in a node with a single resource of each type, the maximal theoretical speedup that we can achieve through pipelining is (after simplification):

$$\theta = \frac{T_s}{T_p} = \frac{(l + p + f)}{\max(l, p, f)} + \frac{\sigma \log 3}{\max(\lambda, \varphi, \delta + \sigma \log B)} = \theta_1 + \theta_2$$

Now, $\theta_1 \geq 1$ whereas $\theta_2 \leq \frac{\sigma \log 3}{\max(\lambda, \varphi)} \ll 1$ for typical values of λ , φ , and σ (refer to Table I). Therefore, we ignore θ_2 and concentrate on choosing the value of B that maximizes θ_1 . The maximum value of θ_1 is 3, which is reached when $l = p = f$, that is, when all three phases are of equal duration. We cannot guarantee $l = f$ since that requires $\lambda = \varphi$. However, we can maximize θ_1 by choosing $p = \max(l, f)$ so that $\theta_1 = 2 + \frac{\min(l, f)}{\max(l, f)}$.

For example, in Figure 3, the ratio between the phases is $l:p:f = 3:4:2$. Thus, θ_1 for this setting is $\frac{3+4+2}{4} = 2.25$. We could improve θ_1 by changing the ratio to 3:3:2, so that $\theta_1 = 2 + \frac{2}{3} \approx 2.67$. In general, setting $\delta B + \sigma B \log B = \max\{\lambda B, \varphi B\}$, we obtain

$$\lg B = \frac{\max\{\lambda, \varphi\} - \delta}{\sigma} \quad (1)$$

This expression represents the size of the postings buffer that must be used to maximize the pipeline speedup on an indexer with a single resource of each type. If we use a buffer of size less than that specified by equation 1, loading or flushing (depending on the relative magnitudes of λ and φ), will be the bottleneck, and the processing phase will be forced to periodically wait for the other phases to complete. An analogous effect will take place for buffer sizes greater than

³ $\lambda = \lambda_1 \lambda_2$, where λ_1 is the rate at which pages can be loaded into memory from the network and λ_2 is the average ratio between the size of a page and the total size of the postings generated from that page.




Fig. 4. Optimal buffer size.

the one prescribed by equation 1. We can generalize equation 1 for an indexer with c identical processors, d identical disks, and i input streams, to obtain

$$\lg B = \frac{\max\{\lambda[\mu/i], \varphi[\mu/d]\} - \delta[\mu/c]}{\sigma[\mu/c]}, \quad (2)$$

where $\mu = \max\{c, d, i\}$.

3.2 Experimental Results

To study the impact of the pipelining technique on indexing performance, we conducted a number of experiments on our testbed, using a single indexer supplied with a stream of Web pages from a distributor.

We first ran the index-builder process in *measurement mode*, where we recorded the execution times of the various phases and determined the values of λ , φ , σ , and δ (Table I). Using the values of these constants in equation 1, we evaluated B to be 16 MB. The optimal total size of the postings buffers, as predicted by our theoretical analysis, is therefore, $3B = 48$ MB.

Impact of buffer size on performance. Figure 4 illustrates how the performance of the index-builder process varies with the size of the buffer. It highlights the importance of the analytical result as an aid in choosing the right buffer size. The optimal total buffer size based on actual experiments turned out to be 40 MB. Even though the predicted optimum size differed slightly from the observed optimum, the difference in running times between the two sizes was less than 15 minutes for a 5 million page collection. For buffer sizes less than 40, loading proved to be the bottleneck, and both the processing and flushing phases had to wait periodically for the loading phase to complete. As the buffer size increased beyond 40, however, the processing phase dominated the execution time since larger and larger buffers of postings had to be sorted.

Performance gain through pipelining. Figure 5 shows how pipelining impacts the time taken to process and generate sorted runs for a variety of input

Table II. Overall Speedup after Including Time for Merging Sorted Runs

Collection Size	Speedup in Generating Sorted Runs (Seconds)	Overall Speedup Including Merging (Seconds)
100,000	96	76
500,000	584	550
1,000,000	1264	1134
2,000,000	2505	2265




Fig. 5. Performance gain through pipelining.

sizes. Note that for small collections of pages, the performance gain through pipelining, though noticeable, is not substantial. This is because small collections require very few pipeline executions and the overall time is dominated by the time required at startup (to load up the buffers) and shutdown (to flush the buffers). This is one of the reasons that pipelined index building has not received attention; most systems have dealt with smaller collections. As collection sizes increase, however, the gain becomes more significant, and for a collection of 5 million pages, pipelining completes almost 1.5 hours earlier than a purely sequential implementation. Our experiments showed that, in general, for large collections, a sequential index-builder is about 30–40% slower than a pipelined index-builder. Note that the observed speedup is lower than the speedup predicted by the theoretical analysis described in the previous section. That analysis was based on an “ideal pipeline,” in which loading, processing and flushing do not interfere with each other in any way. In practice, however, network and disk operations do use processor cycles and access main memory. Hence, any two concurrently running phases, even of different types, do slow each other down.

Note that for a given total buffer size, pipelined execution generates sorted runs that are approximately 3 times smaller than those generated by a sequential indexer. Consequently, 3 times as many sorted runs will need to be merged in the second stage of indexing. As indicated in Table II, however, our experiments indicate that even for very large collection sizes, the potential increase in

merging time is more than offset by the time gained in the first stage through pipelining. We expect that as long as there is enough memory, at merge time, to allocate buffers for the sorted runs, merging performance will not be significantly affected.

4. MANAGING INVERTED FILES IN AN EMBEDDED DATABASE SYSTEM

When building inverted indexes over massive Web-scale collections, the choice of an efficient storage format is particularly important. There have traditionally been two approaches to storing and managing inverted files; using a custom implementation, or leveraging existing relational or object data management systems [Brown et al. 1994; Gorssman and Driscoll 1992].

The advantage of a custom implementation is that it enables very effective optimizations tuned to the specific operations on inverted files (e.g., caching frequently used inverted lists, compressing rarely used inverted lists using expensive methods that may take longer to decompress). When leveraging existing data management systems, such fine-grained control over the implementation may not be possible. There is also likely to be extra overhead in using the higher level of abstraction provided by the data management system. On the other hand, there is opportunity for reduction in development time and complexity. The challenge lies in designing a scheme for storing inverted files that makes optimal use of the storage structures provided by the data management system. The storage scheme must be space efficient and must ensure that the basic lookup operation on an inverted file (i.e., retrieving some or all of the inverted list for a given index term) can be efficiently implemented using the access methods of the data management system.

Note that even though many custom inverted file structures, such as those described by Moffat and Zobel [1996], Anh and Moffat [1998], and Brown [1995] could potentially use database systems for managing the opaque data blocks of their index, such implementations do not fully exploit the capabilities of the underlying system. For instance, such applications would typically need to employ several types of data blocks, as well as custom link structures, and access methods superimposed on top of the underlying database system. For example, Brown [1995] reports on the use of a persistent object store for managing inverted files. The implementation described there uses many different kinds of blocks for storing headers, directories, location lists, and so forth. Basic lookup operations in that implementation translate into several traversals in the underlying object store.

In this section we present and compare different storage schemes for managing large inverted files in an embedded database system. The schemes we suggest use uniform block structures and exploit the native access methods and features of the database system, as much as possible. To test our schemes, we used a freely available embedded database system called Berkeley DB [Olson et al. 1999], which is widely deployed in many commercial applications.

An embedded database is a library, or toolkit, that provides database support for applications through a well defined programming API. Unlike traditional database systems that are designed to be accessed by applications, embedded

databases are linked (at compile-time or run-time) into an application and act as its persistent storage manager. They provide device-sensitive file allocation, database access methods (such as B-trees and hash indexes), and optimized caching, with optional support for transactions, locking, and recovery. They also have the advantage of much smaller footprints compared to full-fledged client-server database systems.

In the following, we briefly sketch the capabilities of Berkeley DB and propose a B-tree based inverted file storage scheme called the *mixed-list scheme*. We qualitatively and quantitatively compare the mixed-list scheme with two other schemes for storing inverted lists in Berkeley DB databases.

4.1 Rationale and Implementation

Berkeley DB provides a programming library for managing (key, value) pairs, both of which can be arbitrary binary data of any length. It offers four access methods, including B-trees and linear hashing, and supports transactions, locking, and recovery.⁴ We chose to use the B-tree access method since it efficiently supports prefix searches (e.g., retrieve inverted lists for all terms beginning with “pre”) and has higher reference locality than hash-based indexes.

The standard organization of a B-tree based inverted file involves storing the index terms in the B-tree along with pointers to inverted lists that are stored separately. Such an organization, though easy to implement using Berkeley DB, does not fully utilize the capabilities of the database system. Since Berkeley DB efficiently handles arbitrary sized keys and values, it is more efficient to store both the index terms and their inverted lists within the database. This enables us to leverage Berkeley DB’s sophisticated caching schemes while retrieving large inverted lists with a minimum number of disk operations.

Storage schemes. The challenge is to design an efficient scheme for organizing the inverted lists within the B-tree structure. We considered three schemes:

- (1) *Full list:* The key is an index term, and the value is the complete inverted list for that term.
- (2) *Single payload:* Each posting (an index term, location pair) is a separate key.⁵ The value can either be empty or may contain additional information about the posting.
- (3) *Mixed list:* The key is again a posting, that is, an index term and a location. The value contains a number of successive postings in sorted order, however, even those referring to different index terms. The postings in the value field are compressed and in every value field the number of postings is chosen so that the length of the field is approximately the same. Note that in this scheme, the inverted list for a given index term may be spread across multiple (key, value) pairs.

In implementing the *mixed list* storage scheme, any well-known encoding technique [Witten et al. 1999] can be used to pack postings efficiently into the

⁴All these features can be turned off for efficiency, if desired.

⁵Storing the indexing term in the key and a single location in the value is not a viable option as the locations for a given term are not guaranteed to be in sorted order.




Fig. 6. Mixed list storage scheme.

value field, as long as the size of the value field is kept approximately constant. For instance, in our implementation, we used an encoding scheme, adapted from Witten et al. [1999], in which successive index terms in the value field are represented using prefix compression, and successive location identifiers are represented in terms of their numerical differences.⁶

Figure 6 illustrates how the *mixed list* storage scheme, using this encoding, stores inverted lists. For simplicity, in this example we assume that no additional information is maintained along with each posting. In our actual implementation, however, we allocated a 2-byte *payload* field, to store extra posting-level information. The top half of the figure depicts inverted lists for four successive index terms and the bottom half shows how they are stored as (key, value) pairs using the mixed-list scheme. For example, the second (key, value) pair in the figure, stores the set of postings (cat, 311), (cat, 328), (catch, 103), (catcher, 147), and so forth, with the first posting stored in the key and the remaining postings stored in the value. As indicated in the figure, the index terms in the value field are prefix compressed, and location identifiers are represented as differences. For example, the posting (cat, 328) is represented by the sequence of entries 3 (an empty field) 17, where 3 indicates the length of the common prefix between the words for postings (cat, 311) and (cat, 328), the (empty field) indicates that both postings refer to the same word, and 17 is the difference between the locations 328 and 311. Similarly, the posting (catch, 103) is represented by the sequence of entries 3 ch 103, where 3 is the length of the common prefix of cat and catch, ch is the remaining suffix for catch, and 103 is the location.

A qualitative comparison of these storage schemes is summarized in Table III. In this table, the symbols “++”, “+”, “+-”, “-”, and “--” denote, in decreasing order, qualitative goodness measures for each scheme relative to the performance metrics.

Index size. The crucial factors determining index size are the number of internal pages (a function of the height of the B-tree) and the number of overflow

⁶These numerical differences are in turn compressed using the ASN.1 Basic Encoding Rules specification [CCITT 1988].

Table III. Comparison of Storage Schemes

Scheme	Index Size	Zig-Zag Joins	Hot Updates
Single payload	--	+	+
Full list	+-	-	-
Mixed list	+-	+-	+-

pages.⁷ In the *single payload* scheme, every posting corresponds to a new key, resulting in rapid growth in the number of internal pages of the database. For large collections, the database size becomes prohibitive, even though Berkeley DB employs prefix compression on keys. At query time, many performance-impeding disk accesses are needed. The situation is significantly better with the *full list* scheme. A database key is only created for every distinct term, and the value field can be well compressed. However, many terms occur only a few times in the collection, whereas others may occur in almost every page. Because of large variations in the size of the value field, many overflow pages are created in the database. In contrast with the *mixed list* scheme, the length of the value field is approximately constant, thus limiting the number of overflow pages. Moreover, the total number of keys (and hence the number of internal pages) can be further reduced by choosing a larger size for the value field. Since the value field can contain postings of different index terms, however, it is not compressed as well as with *full lists*.

Zig-zag joins. The ability to selectively retrieve portions of an inverted list may be very useful when processing conjunctive search queries on an inverted file [Moffat and Zobel 1996]. For example, consider the query *green AND catchflies*. The term *green* occurs on the Web in millions of documents, whereas *catchflies* produces only a couple of dozen hits. A *zig-zag join* [Garcia-Molina et al. 2000] between the inverted lists for *green* and *catchflies* allows us to answer the query without reading out the complete inverted list for *green*. The single payload scheme provides the best support for *zig-zag joins*, since each posting can be retrieved individually. In the *full list* scheme, the entire list must be retrieved to compute the join, whereas with the *mixed list* scheme, access to specific portions of the inverted list is available. For example, in Figure 6, to retrieve locations for *cat* starting at 311, we do not have to read the portion of the list for locations 100–280.

Obviously, the benefit of partial retrieval decreases when only large portions of inverted lists can be accessed individually [Moffat and Zobel 1996]. As we demonstrate in the next section, however, the optimal time for retrieving inverted lists in the *mixed list* scheme is achieved for relatively small sizes of the value field (e.g., 512 bytes in Figure 8). This indicates that with the *mixed list* scheme, a query processor can effectively exploit *zig-zag joins*, reduce the amount of information to be read from disk, and thereby achieve improved performance.

Hot updates. *Hot updates* refers to the ability to modify the index at query time. This is useful when very small changes need to be made to the index

⁷Since values can be of arbitrary length, Berkeley DB uses overflow pages to handle large value fields.




Fig. 7. Varying value field size.

between two successive index rebuilds. For example, Web search services often allow users and organizations to register their home pages. Such additions can be immediately accommodated in the index using the *hot update* facility, without having to defer them until the index is next rebuilt.

In all three schemes, the concurrency control mechanisms of the database can be used to support such *hot updates* while maintaining consistency. The crucial performance factor however, is the length of the inverted list that must be read, modified, and written back, to achieve the update. Since we limit the length of the value field, *hot updates* are faster with *mixed lists* than with *full lists*. The *single payload* scheme provides the best update performance since individual postings can be accessed and modified.

Notice that all three schemes significantly benefit from the fact that the postings are first sorted, and then inserted. Inserting keys into the B-tree in a random order negatively affects the page-fill factor, and expensive tree reorganization is needed. Berkeley DB is optimized for sorted insertions so that high performance and a near- one page fill-factor can be achieved in the initial index construction phase.

In the following section, we present a quantitative comparison of storage and retrieval efficiency for the three storage schemes discussed in this section.

4.2 Experimental Results

The experimental data presented in this section was obtained by building an inverted index over a collection of 2 million Web pages. The collection contains 4.9 million distinct terms with a total of 312 million postings.⁸

Figure 7 illustrates how the choice of the storage scheme affects the size of the inverted file. It shows the variation of index size with value field size, when using the *mixed list* scheme. The dotted line represents the index size when the same database was stored using the *full list* scheme. Note that, since the value field size is not applicable to the *full list* scheme, the graph is just a horizontal

⁸Only one posting was generated for all the occurrences of a term in a page.

Table IV. Mixed-List Scheme Index Sizes

Number of Pages (million)	Input Size (GB)	Index Size (GB)	Index Size (%)
0.1	0.81	0.04	6.17
0.5	4.03	0.24	6.19
2.0	16.11	0.99	6.54
5.0	40.28	2.43	6.33

line. The *single payload* scheme can be viewed as an extreme case of the *mixed* scheme with the value field being empty. Figure 7 shows that both very small and very large value fields have an adverse impact on index size. In the *mixed list* scheme, very small value fields will require a large number of internal database pages (and a potentially taller B-tree index) to accommodate all the postings. On the other hand, very large value fields will cause Berkeley DB to allocate a large number of overflow pages, which leads to a larger index. As indicated in the figure, a value field size of 512 bytes provided the best balance between these two effects. The *full list* scheme results in a moderate number of overflow pages and internal database pages. However, it still requires about 30% more storage space than the optimal *mixed list* inverted file. For all of the examined storage schemes, the time to write the inverted file to disk was roughly proportional to the size of the file.

Table IV shows how the index size (using the *mixed list* scheme) varies with the size of the input collection. The index sizes listed in Table IV include the sum of the sizes of the inverted files and the associated lexicons. The numbers for Table IV were generated by using *mixed lists* with the optimal value field size of 512 bytes derived from Figure 7. Table IV shows that the *mixed list* storage scheme scales very well to large collections. The size of the index is consistently less than 7% of the size of the input HTML text. This compares favorably with the sizes reported for the VLC2 track (which also used crawled Web pages) at TREC-7 [Hawking and Craswell 1998] where the best reported index size was approximately 7.7% of the size of the input HTML. Our index sizes are also comparable to other recently reported sizes for non-Web document collections using compressed inverted files [Anh and Moffat 1998]. Note that exact index sizes are dependent on the type and amount of information maintained along with each posting (e.g., information to handle proximity queries). We believe, however, that the 2-byte payload field used in our implementation can accommodate most posting-level information normally stored in inverted indexes.

Figure 8 illustrates the effect of value field size on inverted list retrieval time. Once again, the dotted horizontal line represents the retrieval time when using the *fixed list* scheme. Figure 8 was produced by generating uniformly distributed query terms and measuring the time⁹ required to retrieve the entire inverted list for each query term. The optimal retrieval performance in the *mixed list* scheme is achieved when the value field size is between 512 and 1024

⁹A warming-up period was allowed before the measurements to fill the database and file system cache.




Fig. 8. Time to retrieve inverted lists.

bytes. Notice that (from Figures 7 and 8) a value field size of 512 bytes results in maximum storage as well as maximum retrieval efficiency for the *mixed list* scheme. Figure 8 also indicates that both the *fixed list* and *mixed list* (with optimal value field size) schemes provide comparable retrieval performance.

Note that Figure 8 only indicates the raw inverted list retrieval performance of the different storage schemes. True query processing performance will be affected by other factors, such as caching (of inverted lists), use of query processing techniques such as *zig-zag joins*, and the distribution of the query terms.

5. COLLECTING GLOBAL STATISTICS

Most text-based retrieval systems use some kind of collection-wide information to increase effectiveness of retrieval [Viles and French 1995]. One popular example is the inverse document frequency (IDF) statistics used in ranking functions. The IDF of a term is the inverse of the number of documents in the collection that contain that term. If query servers have only IDF values over their local collections, then rankings would be skewed in favor of pages from query servers that return few results.

Depending on the particular global statistic, the ranking function, and the nature of the collection, it may or may not be necessary for a statistic to be computed accurately. In some cases, it suffices to estimate the global statistic from the local values at the individual query servers, or from sampling (see related work in Section 6). In this section, however, we analyze the problem of gathering *accurate* collection-wide information (with minimum overhead), for the cases where this is required. We present two techniques that can be used for gathering different types of collection-wide information, though here we focus on the problem of collecting *term-level global statistics*, such as IDF values.¹⁰

¹⁰*Term-level* refers to the fact that any gathered statistic describes only single terms, and not higher level entities such as pages or documents.

Table V. Comparing Strategies

	Phase	Statistician Load	Memory Usage	Parallelism
ME	merging	+–	+	+–
FL	flushing	–	–	++

5.1 Design

Some authors suggest computing global statistics at query time. This would require an extra round of communication among the query servers to exchange local statistics. This communication adversely impacts query processing performance, especially for large collections spread over many servers. Since query response times are critical, we advocate precomputing and storing statistics at the query servers during index creation.

Our approach is based on using a dedicated server, known as the *statistician*, for computing statistics. Having a dedicated *statistician* allows most computation to be done in parallel with other indexing activities. It also minimizes the number of conversations among servers, since indexers exchange statistical data with only one *statistician*. Local information is sent to the *statistician* at various stages of index creation, and the *statistician* returns global statistics to the indexers in the merging phase. Indexers then store the global statistics in the local lexicons. A lexicon consists of entries of the form (*term*, *term-id*, *local-statistics*, *global-statistics*), where the terms stored in a lexicon are only those terms occurring in the associated inverted file (Section 2).

In order to avoid extra disk I/O, local information is sent to the *statistician* only when it is already in memory. We have identified two phases in which this occurs: *flushing*—when sorted runs are written to disk, and *merging*—when sorted runs are merged to form inverted lists and the lexicon. Sending information in these two phases leads to two different strategies with various tradeoffs, which are discussed in the next section. We note here only that by sending information to the *statistician* in these phases without additional I/O’s, a huge fraction of the statistic collection is eliminated.

Sending information to the *statistician* is further optimized by summarizing the postings. In both identified phases, postings occur in at least partially sorted order, that is, multiple postings for a term pass through memory in groups. Groups are condensed into (*term*, *local aggregated information*) pairs which are sent to the *statistician*. For example, if an indexer holds 10,000 pages that contain the term “cat,” instead of sending 10,000 individual postings to the *statistician*, the indexer can count the postings as they pass through memory in a group and send the summary (“cat”, 10000) to the *statistician*. The *statistician* receives local counts from all indexers, and aggregates these values to produce the global document frequency for “cat”. This technique greatly reduces network overhead in collecting statistics.

5.2 Statistic Gathering Strategies

Here we describe and compare the two strategies mentioned above for sending information to the *statistician*. Table V qualitatively summarizes their




Fig. 9. ME strategy.




Fig. 10. FL strategy.

characteristics (using the same notation as in Table III). The column titled “Parallelism,” refers to the degree of parallelism possible within each strategy.

ME Strategy (sending local information during merging). Summaries for each term are aggregated as inverted lists are created in memory, and sent to the *statistician*. The *statistician* receives parallel sorted streams of $(term, local-aggregate-information)$ values from each indexer and merges these streams by term, aggregating the sub-aggregates for each term to produce global statistics. The statistics are then sent back to the indexers in sorted term order. This approach is entirely stream based, and does not require in-memory or on-disk data structures at the *statistician* or indexer to store intermediate results. Using streams, however, means that the progress of each indexer is synchronized with that of the *statistician*, which in turn causes indexers to be synchronized with each other. As a result, the slowest indexer in the group becomes the bottleneck, holding back the progress of faster indexers. Figure 9 illustrates the ME strategy for collecting document frequency statistics for each term. Note that the bottom lexicon does not include statistics for “rat” because the term is not present in the local collection.

FL Strategy (sending local information during flushing). As sorted runs are flushed to disk, postings are summarized and the summaries sent to the *statistician*. Since sorted runs are accessed sequentially during processing, the *statistician* receives streams of summaries in globally *unsorted* order. To compute statistics from the unsorted streams, the *statistician* keeps an in-memory hash table of all terms and their statistics, and updates the statistics as summaries for a term are received. At the end of the processing phase, the *statistician* sorts the statistics in memory and sends them back to the indexers. Figure 10 illustrates the FL strategy for collecting document frequency statistics.




Fig. 11. Overhead of statistics collection.

5.3 Experiments

To demonstrate the performance and scalability of the collection strategies, we ran the index-builder and merging processes on our testbed, using a hardware configuration consisting of four indexers.¹¹ We experimented with four different collection sizes—100000, 500000, 1000000, and 2000000 pages, respectively. The results are shown in Figure 11, where we can see the *relative overhead* (defined as $\frac{T_2 - T_1}{T_1}$ where T_2 is the time for full index creation with statistics collection and T_1 is the time for full index creation with *no* statistics collection) for both strategies. In general, experiments show that the FL strategy outperforms ME, although they seem to converge as the collection size becomes large. Furthermore, as the collection size grows, the relative overheads of both strategies decrease.

Comparison of strategies. At first glance ME might be expected to outperform FL. Since the *statistician* receives many summary streams in FL, but only one from each indexer in ME, it performs more comparison and aggregation in FL than in ME. As mentioned earlier, however, merging progress in ME is synchronized among the servers. Hence, a good portion of computation done at the *statistician* cannot be done in parallel with merging activities at the indexer.

In FL, on the other hand, the indexer simply writes summaries to the network and continues with its other work. The *statistician* can then asynchronously process the information from the network buffer in parallel. Not all work can be done in parallel, however, since the *statistician* consumes summaries at a slower rate than the indexer writes them to network, and the network buffer generally cannot hold all the summaries from a sorted run. Hence there is still nontrivial waiting at the indexer during flushing as summaries are sent to the *statistician*.

¹¹All indexers had the specifications listed in Section 2.




Fig. 12. Varying lexicon buffer size.

Enhancing parallelism. In the ME strategy, synchronization occurs when an indexer creates a lexicon entry and summary for a term, sends the summary to the *statistician*, and then waits for the global statistic to be returned so that the lexicon entry can be completed. To reduce the effect of synchronization, the merging process can, instead, write lexicon entries to a lexicon buffer, and a separate process will wait for global statistics and include them in the entries. In this way, the first process need not block while waiting, and both processes can operate in parallel.

Figure 12 shows the effect of lexicon buffer size on merging performance over a collection of a million pages. Because lexicon entries are created faster than global statistics are returned on all indexers but the slowest, the lexicon buffer often becomes full. When this occurs, the process creating lexicon entries must block until the current state changes. Because larger lexicon buffers reduce the possibility of saturation, we see that as expected, initial increases in size result in large performance gains. As lexicon buffer size becomes very large, however, performance slowly deteriorates due to memory contention. Although the entire buffer need not be present in memory at any one time, the lexicon buffer is accessed cyclically; therefore LRU replacement and the fast rate at which lexicon entries are created cause buffer pages to cycle rapidly through memory, swapping out other, non buffer, pages.

Sub-linear growth of overhead. The constant decrease of the ME and FL relative overhead in Figure 11 is because of the fact that the number of distinct terms in a page collection is a sub-linear function of collection size. The overhead incurred by gathering statistics grows linearly with the number of terms in the collection, while the cost of index creation grows linear-logarithmically with the size of the collection. As a result, overhead of statistic collection displays sub-linear growth with respect to index creation time. This prediction is consistent with our experimental results.

The decreasing relative overhead for FL, however, is subject to the constraint that the hashtable can fit in memory. Considering that a collection of a billion

pages would require a hash table of roughly 5–6 GB in size,¹² this constraint may become a problem for very large collections. While a memory of 6 GB is not completely unreasonable, a simple alternative, using only commodity hardware, would be to run several *statisticians* in parallel, and partition the terms alphabetically between *statisticians*. In this way, each *statistician* can collect a moderately sized set of global statistics. We have not yet implemented this option in our system.

6. RELATED WORK

There has been recent interest, motivated by the Web, in designing scalable techniques to speed up inverted index construction using distributed architectures. Ribeiro-Neto et al. [1999] describe three techniques to efficiently build an inverted index using a distributed architecture. However, they focus on building global (partitioning index by term), however, rather than local (partitioning by collection), inverted files. Furthermore, they do not address issues such as global statistics collection and optimization of the indexing process on each individual node.

Our technique for structuring the index engine as a pipeline has much in common with pipelined query execution strategies employed in relational database systems [Garcia-Molina et al. 2000]. In Chakrabarti and Muthukrishnan [1996], the authors present a variety of algorithms for resource scheduling, with applications to scheduling pipeline stages.

There has been prior work on using relational or object-oriented data stores to manage and process inverted files [Blair 1988; Brown et al. 1994; Gorssman and Driscoll 1992]. Brown et al. [1994] describe the architecture and performance of an IR system that uses a persistent object store to manage inverted files. Their results show that using an off-the-shelf data management facility improves the performance of an information retrieval system, primarily because of intelligent caching and device-sensitive file allocation. We experienced similar performance improvements, for the same reasons, by employing an embedded database system. Our storage format differs greatly from theirs because we utilize a single B-tree for storing all the inverted lists in uniformly structured blocks.

As with the *mixed list* scheme presented in this paper, the “self-indexing” inverted list structures described in Moffat and Zobel [1996] (skipped-list) and Anh and Moffat [1998] (random-access lists) also provides selective access to portions of an inverted list. Moffat and Zobel [1996] propose dividing the inverted list into blocks, each containing a fixed number of postings. To extract postings from within a block, the block needs to be sequentially decoded, however, a separate skip list (also compressed) is used to navigate from one block to another. A given posting is, therefore, located by first navigating the skip list to identify the block in which it might be contained and then decoding that block. Anh and Moffat [1998] use fixed-length blocks (constant number of bits),

¹²From collection statistics described in Melnik et al. [2000], a billion pages will contain roughly 310 million distinct terms. If each term uses 20 bytes of storage, this will result in a hashtable of about 5.77 GB.

rather than blocks with a constant number of postings, and develop a scheme to efficiently encode and decode such blocks.

Conceptually, the blocks used in these structures correspond to the value fields of the *mixed list* scheme. The *synchronization points* used in Moffat and Zobel [1996] correspond to the key fields of the *mixed list* scheme. There are a couple of significant differences, however. First, the value fields are organized as the leaf nodes of a prefix-compressed B-tree,¹³ whereas in Moffat and Zobel [1996] and Anh and Moffat [1998], all blocks belonging to an inverted list are organized sequentially. Second, a given value field can include postings from more than one inverted list, whereas that is not the case with the structures described in Moffat and Zobel [1996] and Anh and Moffat [1998]. Hence, the synchronization points in Moffat and Zobel [1996] contain just document identifiers, whereas keys also contain index terms.

De Kretser et al. [1998] present different types of global statistics that may be used for query processing. Their focus, however, is on the types and uses of these statistics rather than on the actual process of collecting them. Viles [1994] and Viles and French [1995] discuss maintenance of global statistics in a distributed text index, specifically addressing the challenges that arise from incremental updates. Their work is complementary to our strategies for gathering statistics during index construction.

Global statistics are also important in meta-search environments [Lawrence and Giles 1998; Craswell et al. 1999; Gravano et al. 1997], where ranked results from several (possibly autonomous) search servers must be merged to produce a global ranking. For such environments, as an alternative to accurate global statistics, Craswell et al. [1999] suggest the use of “reference statistics”—estimates of the true statistics derived from sampling, or from the statistics of a different document collection. The use of reference statistics, however, was recommended mainly for systems in which collection of accurate statistics was not feasible or was extremely expensive (for instance, when merging results from *autonomous* search services). In our system, collection of accurate statistics is feasible since all the individual query servers and indexers are under our control.

A great deal of work has been done on several other issues relevant to inverted-index based information retrieval, that has not been discussed in this paper. Such issues include index compression [Moffat and Bell 1995; Anh and Moffat 1998; Witten et al. 1999], incremental updates [Brown et al. 1994; Jeong and Omiecinski 1995; Tomasic et al. 1994; Witten et al. 1999; Zobel et al. 1992], and distributed query performance [Tomasic and Garcia-Molina 1993a; Tomasic and Garcia-Molina 1993b].

7. CONCLUSIONS

In this paper we addressed the problem of efficiently constructing inverted indexes over large collections of Web pages. We proposed a new pipelining technique to speed up index construction and demonstrated how to identify the

¹³Note that all the leaf nodes of the B-tree are also sequentially linked into a list as part of the B-tree implementation.

right buffer sizes for maximum performance. We showed that for large collection sizes, the pipelining technique can speed up index construction by several hours. We proposed and compared different schemes for storing and managing inverted files using an embedded database system. We showed that an intelligent scheme for packing inverted lists in the storage structures of the database can provide performance and storage efficiency comparable to tailored inverted file implementations. Finally, we identified the key characteristics of methods for efficiently collecting global statistics from distributed inverted indexes. We proposed two such methods and compared and analyzed the tradeoffs between them.

In the future, we intend to extend our testbed to incorporate distributed query processing and explore algorithms and caching strategies for efficiently executing queries. We also intend to experiment with indexing and querying over larger collections and integration of our text-indexing system with indexes on the link structure of the Web.

APPENDIX

A. PROOF OF OPTIMALITY OF EQUISIZE BUFFERS

We are given $T_p = \max\{\sum_{i=1}^N l_i, \sum_{i=1}^N p_i, \sum_{i=1}^N f_i\}$. If loading or flushing is the bottleneck, T_p is either λB_{total} or φB_{total} , and has the same value for all distributions of B_i including an equisize distribution. If processing is the critical phase, $T_p = \sum_{i=1}^N (\delta B_i + \sum B_i \log B_i)$. Under the constraint that $\sum_{i=1}^N B_i = B_{total}$, the absolute minimum of T_p is reached when $B_i = \frac{B_{total}}{N}$ for each i , i.e., when all buffers have equal sizes. This global extremum can be easily determined using standard analysis techniques such as Lagrange multipliers.

REFERENCES

- ANH, V. N. AND MOFFAT, A. 1998. Compressed inverted files with reduced decoding overheads. In *Proceedings of the 21st International Conference on Research and Development in Information Retrieval* (August), 290–297.
- BLAIR, D. C. 1988. An extended relational document retrieval model. *Information Processing and Management* 24, 3, 349–371.
- BROWN, E. W. 1995. Fast evaluation of structured queries for information retrieval. In *Proceedings of ACM Conference on Research and Development in Information Retrieval (SIGIR)*, ACM Press, New York, NY, 30–38.
- BROWN, E. W., CALLAN, J. P., AND CROFT, W. B. 1994. Fast incremental indexing for full-text information retrieval. In *Proceedings of the 20th International Conference on Very Large Databases* (September), 192–202.
- BROWN, E. W., CALLAN, J. P., CROFT, W. B., AND MOSS, J. E. B. 1994. Supporting full-text information retrieval with a persistent object store. In *4th International Conference on Extending Database Technology* (March), 365–378.
- CCITT. 1988. Recommendation X.209, Specification of Basic Encoding Rules for Abstract Syntax Notation one (ASN. 1).
- CHAKRABARTI, S. AND MUTHUKRISHNAN, S. 1996. Resource scheduling for parallel database and scientific applications. In *8th ACM Symposium on Parallel Algorithms and Architectures* (June), ACM Press, New York, NY, 329–335.
- CHO, J. AND GARCIA-MOLINA, H. 2000. The evolution of the web and implications for an incremental crawler. To appear in the 26th International Conference on Very Large Databases.

- CRASWELL, N., HAWKING, D., AND THISTLEWALTE, P. 1999. Merging results from isolated search engines. In *Proceedings of the 10th Australasian Database Conference* (January).
- DE KRETSER, O., MOFFAT, A., SHIMMIN, T., AND ZOBEL, J. 1998. Methodologies for distributed information retrieval. In *Proceedings of the 18th International Conference on Distributed Computing Systems*.
- FALOUTSOS, C. AND CHRISTODOULAKIS, S. 1984. Signature files: An access method for documents and its analytical performance evaluation. *ACM Transactions on Office Information Systems* 2, 4 (October), ACM Press, New York, NY, 267–288.
- GARCIA-MOLINA, H., ULLMAN, J., AND WIDOM, J. 2000. *Database System Implementation*. Prentice-Hall, Eaglewood Cliffs, NJ.
- GORSSMAN, D. A. AND DRISCOLL, J. R. 1992. Structuring text within a relation system. In *Proceedings of the 3rd International Conference on Database and Expert System Applications* (September), 72–77.
- GRAVANO, L., CHANG, K., GARCIA-MOLINA, H., LAGOZE, C., AND PAEPCKE, A. 1997. STARTS—stanford protocol for internet retrieval and search. <http://www-db.stanford.edu/gravano/starts.html>.
- HAWKING, D. AND CRASWELL, N. 1998. Overview of TREC-7 very large collection track. In *Proceedings of the Seventh Text Retrieval Conference* (November), 91–104.
- HIRAI, J., GARCIA-MOLINA, H., AND PAEPCKE, A., RAGHAVAN, S. 2000. WebBase: A repository of web pages. In *Proceedings of the 9th International World Wide Web Conference* (May), 277–293.
- INKTOMI. 2000. Inktomi WebMap. <http://www.inktomi.com/webmap/>.
- JEONG, B.-S. AND OMIECINSKI, E. 1995. Inverted file partitioning schemes in multiple disk systems. *IEEE Transactions on Parallel and Distributed Systems* 6, 2 (February), IEE Computer Society Press, Los Alamitos, CA, 142–153.
- LAWRENCE, S. AND GILES, C. L. 1998. Inquirus, the NECI meta search engine. In *Proceedings of the 7th International World Wide Web Conference*.
- LAWRENCE, S. AND GILES, C. L. 1999. Accessibility of information on the web. *Nature* 400, 107–109.
- MANBER, U. AND MYERS, G. 1990. Suffix arrays: A new method for on-line string searches. In *Proceedings of the 1st ACM-SIAM Symposium on Discrete Algorithms*, ACM Press, New York, NY, 319–327.
- MARTIN, P., MACLEOD, I. A., AND NORDIN, B. 1986. A design of a distributed full text retrieval system. In *Proceedings of the ACM Conference on Research and Development in Information Retrieval* (September), ACM Press, New York, NY, 131–137.
- MELNIK, S., GARCIA-MOLINA, H., YANG, B., AND RAGHAVAN, S. 2000. Building a distributed full-text index for the web. Technical Report SIDL-WP-2000-0140 (July), Stanford Digital Library Project, Computer Science Dept., Stanford University. Available at www.diglib.stanford.edu/cgi-bin/get/SIDL-WP-2000-0140.
- MOFFAT, A. AND BELL, T. 1995. In situ generation of compressed inverted files. *Journal of the American Society for Information Science* 46, 7, 537–550.
- MOFFAT, A. AND ZOBEL, J. 1996. Self-indexing inverted files for fast text retrieval. *ACM Transactions on Information Systems* 14, 4 (October), ACM Press, New York, NY, 349–379.
- OLSON, M., BOSTIC, K., AND SELTZER, M. 1999. Berkeley DB. In *Proceedings of the 1999 Summer Usenix Technical Conference* (June).
- RIBEIRO-NETO, B. AND BARBOSA, R. 1998. Query performance for tightly coupled distributed digital libraries. In *Proceedings of the 3rd ACM Conference on Digital Libraries* (June), ACM Press, New York, NY, 182–190.
- RIBEIRO-NETO, B., MOURA, E. S., NEUBERT, M. S., AND ZIVIANI, N. 1999. Efficient distributed algorithms to build inverted files. In *Proceedings of the 22nd ACM Conference on Research and Development in Information Retrieval* (August), ACM Press, New York, NY, 105–112.
- SALTON, G. 1989. *Information Retrieval: Data Structures and Algorithms*. Addison-Wesley, Reading, Massachusetts.
- TOMASIC, A. AND GARCIA-MOLINA, H. 1993a. Performance of inverted indices in shared-nothing distributed text document information retrieval systems. In *Proceedings of the 2nd International Conference on Parallel and Distributed Information Systems* (January), 8–17.
- TOMASIC, A. AND GARCIA-MOLINA, H. 1993b. Query processing and inverted indices in shared-nothing document information retrieval systems. *VLDB Journal* 2, 3, 243–275.

- TOMASIC, A., GARCIA-MOLINA, H., AND SHOENS, K. 1994. Incremental update of inverted list for text document retrieval. In *Proceedings of the 1994 ACM SIGMOD International Conference on Management of Data* (May), ACM Press, New York, NY, 289–300.
- VILES, C. L. 1994. Maintaining state in a distributed information retrieval system. In *32nd Southeast Conference of the ACM*, ACM Press, New York, NY, 157–161.
- VILES, C. L. AND FRENCH, J. C. 1995. Dissemination of collection wide information in a distributed information retrieval system. In *Proceedings of the 18th International ACM Conference on Research and Development in Information Retrieval* (July), ACM Press, New York, NY, 12–20.
- WITTEN, I. H., MOFFAT, A., AND BELL, T. C. 1999. *Managing Gigabytes: Compressing and Indexing Documents and Images* (2nd ed.). Morgan Kaufman Publishing, San Francisco.
- ZOBEL, J., MOFFAT, A., AND SACKS-DAVIS, R. 1992. An efficient indexing technique for full-text database systems. In *18th International Conference on Very Large Databases* (August 1992), pp. 352–362.

Received April 2001; revised May 2001; accepted July 2001