

# On the Integration of Structured Data and Text: *A Review of the SIRE Architecture*

(An Invited Overview)

O. Frieder, A. Chowdhury, D. Grossman, & M.C. McCabe  
Information Retrieval Laboratory  
Illinois Institute of Technology  
ophir@ir.iit.edu

## 1.0 Introduction

Over the past decade, members of the Information Retrieval Lab have designed, developed, and deployed a variety of information retrieval systems. A central theme for all of our systems was the integration of structured data and text. One of our more recent efforts, SIRE, a Scalable Information Retrieval Engine [Grossman97, Grossman98, Lundquist99] is the focus of this paper. For completeness, we review some of the functionality of SIRE although it is described, in greater detail, in other forums. We describe the architecture of the prototype developed for the National Institutes of Health (NIH) National Center for Complementary and Alternative Medicine (NCCAM) [Frieder00] by some of the members of the laboratory. The version deployed at NCCAM is a more industrialized version of this prototype.

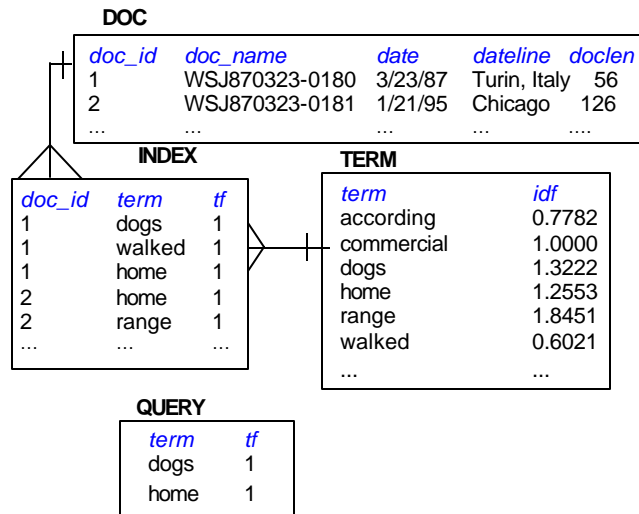
The mainstream approach in the development of information retrieval systems uses a customized inverted index to represent the text. SIRE, on the other hand, is a relational information retrieval approach and uses relations to model an inverted index. Storing the full text in a relational environment integrates the search of unstructured data with the traditional structured data search of Relational Database Management Systems (RDBMS). By using only standard SQL, SIRE leverages investment of the commercial relational database industry, while providing all capabilities of a more traditional information retrieval approach. RDBMS offer a wide variety of functionality such as concurrency control, recovery, security, portability, scalability, and robustness. RDBMS vendors continuously improve these features and incorporate advances made in hardware and software. Thus, an application using an RDBMS is able to keep up with the technology curve with less investment than a custom solution.

SIRE is implemented strictly as a relational database application. Key information retrieval techniques such as leading similarity measures, proximity searching, n-grams, passages, phrase indexing, and relevance feedback are all implemented using standard SQL. By adhering to strictly standard SQL, SIRE is completely portable across platforms and database management systems. Thus far, laboratory members, our industrial collaborators, or our sponsors have implemented SIRE on the NCR DBC-1012, Microsoft SQL Server, Sybase, Oracle, IBM DB2 and SQL/DS database management systems.

SIRE achieves good performance and scalability and is currently in production use in a variety of text search applications both commercially as part of multiple industrial efforts and in various government laboratories including the National Institutes of Health National Center for Complementary and Alternative Medicine. The relational platform offers equivalent capabilities to traditional Information Retrieval (IR) systems while providing a parallel, scalable, maintainable architecture with a unified platform for integrating searches of structured and unstructured data.

## 2.0 SIRE Functionality Overview

In SIRE, relations are used to model an inverted index, see Figure 1[Grossman97]. As shown, three relations, namely DOC, INDEX, and TERM, are used to represent the entire collection. The document (DOC) relation stores all the metadata information about the individual documents. That is, the DOC relation stores the document name, dateline, length, etc. The term (TERM) relation stores all terms in the collection and their corresponding inverted document frequency (idf) score. The example demonstrates only words as terms, but phrases can likewise be stored in the relation. Finally, the index (INDEX) relation itemizes which terms appear in which documents and how many times the term appears within each document.



**Figure 1: The Relational Structure of the Scalable IR Engine**

Storing the full text in a relational environment integrates the search of unstructured data with the traditional RDBMS search of structured data. The retrieval accuracy of SIRE, as measured by the metrics *precision* and *recall*, is comparable with the state-of-the-art systems because leading scoring measures are implemented in standard SQL within SIRE. An example of such SQL is shown in SQL-1, using the cosine similarity measure.

```

SQL-1:  SELECT      d.doc_id, d.doc_name, SUM((i.tf * t.idf * q.tf * t.idf)/d.doclen)
          FROM        INDEX i, DOC d, QUERY q, TERM t
          WHERE       d.doc_id = i.doc_id      AND
                    q.term = i.term           AND
                    t.term = q.term
          GROUP BY    d.doc_id, d.doc_name
          ORDER BY 3 DESC;

```

Modifications to the SUM() element permits implementation of most leading similarity measures. For instance, with the additional computation and storage of some document statistics, (log of the average term frequency), some collection statistics (average document length and the number of documents) and term statistics (document frequency), the following measures, as listed under SQL-2 and SQL-3, are implemented. In the pivoted normalization measure, the constant 0.20 is the pivot value proposed in [Singhal96]. In the OKAPI measure, the constants are the values described in [Robertson98].

```

SQL-2: Pivoted normalization measure
          SUM(((1 + LOG(i.tf)) / ((d.LogAvgTF) * (AvgDocLen + (0.20 * d.doclen))))
            * (t.idf * ((1 + LOG(q.tf)) / (q.LogAvgTF))))

```

```

SQL-3: OKAPI Probabilistic measure
          SUM(LOG((((NumDocs - t.df) + 0.5) / (t.df + 0.5))
            * ((2.2*i.tf) / (0.3 + ((0.9 * d.doclen)/AvgDocLen) + i.tf))))

```

Techniques such as relevance feedback, proximity, phrases, and n-grams are common in traditional information retrieval and are therefore also available in the SIRE system. For brevity, we only discuss

term	doc_id	tf-idf
dogs	{( 1	1.07 ),
	( 7	3.02 ),
	( 19	2.09 ),
	( 44	1.76 ),
	( 45	1.99 ),
	( 63	2.65 )}

(A) Term Processing

doc_id	term	tf-idf
1	{( dogs	1.07 ),
	( walked	3.33 ),
	( home	2.23 ),
	( cats	1.58 ),
	( climbed	1.36 ),
	( trees	2.11 )}

(B) Relevance Feedback Processing

Figure 2. Clustered Indexes

relevance feedback, as it is the more commonly used information retrieval technique. To implement relevance feedback, the SQL shown is SQL-1 is executed and the best (using  $N \times \text{idf}$  ranking) terms from the top-ranked documents are added to the original query terms (inserted into the Query table) and the query is rerun [Lundquist99].

The above discussion focused on the vector space model. However, some applications rely not only on the vector space model but either on just the Boolean model or on a combination of both models. In the Boolean model, one operator, namely TAND (threshold AND), requires a certain number of the specified terms from within the entire query to be present in a document to qualify a document as relevant. This feature can be quite complex to implement in traditional information retrieval systems. However, with SQL, TAND is easily achieved by adding a `HAVING COUNT(*) >= threshold_number_of_terms`.

```
SQL-4:  SELECT      d.doc_id, d.doc_name
        FROM        INDEX i, DOC d, QUERY q
        WHERE       d.doc_id = i.doc_id AND
                   q.term = i.term
        GROUP BY   d.doc_id, d.doc_name
        HAVING COUNT(*) >= threshold_number_of_terms;
```

### 3.0 Performance Enhancements

The primary performance enhancements supported in SIRE are:

- Clustered Indexes
- Thresholding
  - Index Thresholding
  - Query Thresholding

#### 3.1 Clustered Indexes

Physically implementing the relational structure of SIRE as illustrated in Figure 1 introduces redundancy due to the repetition of the `doc_id` attribute in the RDBMS indexing of the INDEX relation. Furthermore, for relevance feedback processing, the same duplication occurs in the indexing of the `term` attribute. This redundancy increases the storage requirements, hence the I/O processing, and therefore the total processing time. In the actual SIRE implementation, clustered indexes are used, as shown in Figure 2, to reduce the I/O. As seen, the redundancy is eliminated. As an aside, also note that instead of the term frequency value, the computed `tf-idf` value is stored. This reduces the computational time as repeated `tf * idf` computations are eliminated.

#### 3.2 Thresholding

Thresholding is a technique where terms in certain frequency ranges are eliminated from the search – either removing them from the query or, equivalently, eliminating them from the index, saving space and

Query Threshold Percentage	Relevant Documents Retrieved	Execution Time (CPU Cycles)	New Relevant Documents per CPU Cycle
10	831	475	1.75
20	1505	950	1.58
25	1601	1238	1.29
33	1657	1736	0.95
50	1856	5215	0.36
75	2119	13181	0.16
100	2138	33989	0.06

**Figure 3. Relevant Retrieved and Computation Time as a Function of Query Threshold**

processing speed [Grossman95]. Index Thresholding addresses the possibility that the large term weights of the *infrequently* occurring terms may be artificially inflating the relevance ranking scores of documents. Our experiments with TREC data indicate that eliminating terms occurring in fewer than 75 documents (0.014% of the collection) improves the precision/recall by 24%, reduces storage required by 26% and improves runtime.

In Query Thresholding, the most *frequent* terms are removed from the query since they are the worst discriminators. In Figure 3, we show that a query threshold of 75% retrieves approximately the same number (99.1%) of relevant documents (2119) as the same query where all 100% of the terms are used (2138). By using a 75% query threshold, a 61.3% reduction in runtime came at a cost of only 19 (0.9%) relevant documents. As shown, the rate of new document detection significantly degrades as the query threshold increases. Adjusting the query threshold provides the user with a means to significantly reduce the runtime without significantly reducing the accuracy.

## 4.0 System Software Architecture

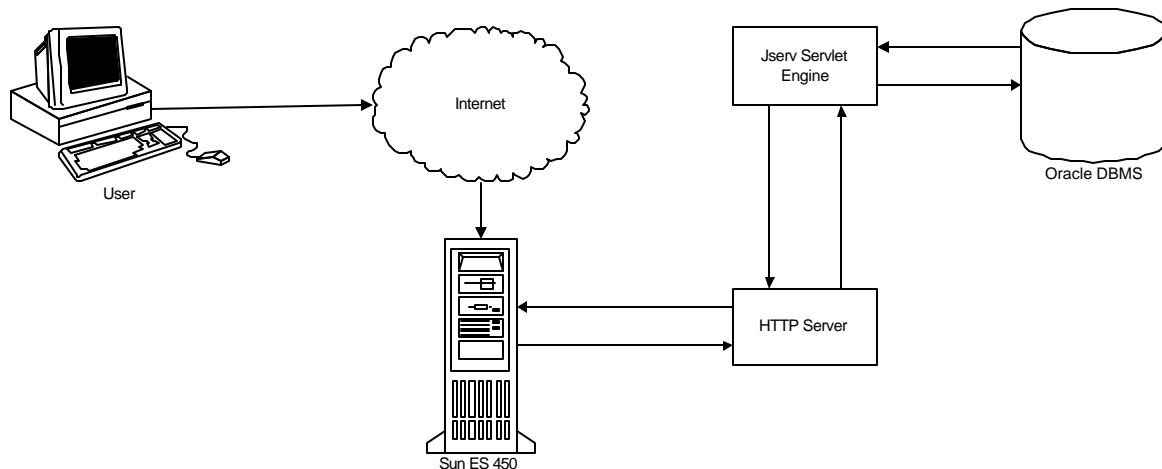
The SIRE software architecture comprises of four parts:

- System Software: Operating Systems, Disk Striping
- Relational Database Management Systems
- Java servlets
- Web server

The system software deals with the operating system layout, namely, the distribution of disk resources, optimization of I/O throughput, memory usage, disk striping to reduce I/O contention, etc. The second component is the use of an RDBMS to store structured and unstructured information. The third component is the query software architecture, caching algorithms, and other optimizations for efficiency, while the fourth is a web server providing the UI (User Interface) to the client via a traditional web interface.

### 4.1 System Software Architecture

The SIRE system uses Solaris 2.7 as its operating system providing the interface to the underlying hardware. The system is running on a Sun ES-450 with four 300MHz processors and 4 Gigabytes of system memory. The system has approximately 300 gigabytes of disk space connected to three SCSI controllers to distribute the I/O load for each controller. To reduce I/O contention for large data files, we use Disk Suit 4.2 to build logical disks that stripe the file system over several controllers and disks in 16K blocks. The disk layout for the data files is described in the next section.



**Figure 4. Client Request Flow Chart**

## 4.2 RDBMS Software Architecture

The SIRE system uses an RDBMS to store data. We currently use Oracle 8.0.5 as the RDBMS. The SIRE system uses indexes to speed up the retrieval of information for queries. We use logical disks to reduce I/O contention when running multiple queries at the same time. The indexes for each table are striped across two controllers and over two physical disks. The data files are striped across two controllers and four disks. We model the vector space model to search unstructured text and provide relevance rankings. We search structured information with standard RDBMS techniques.

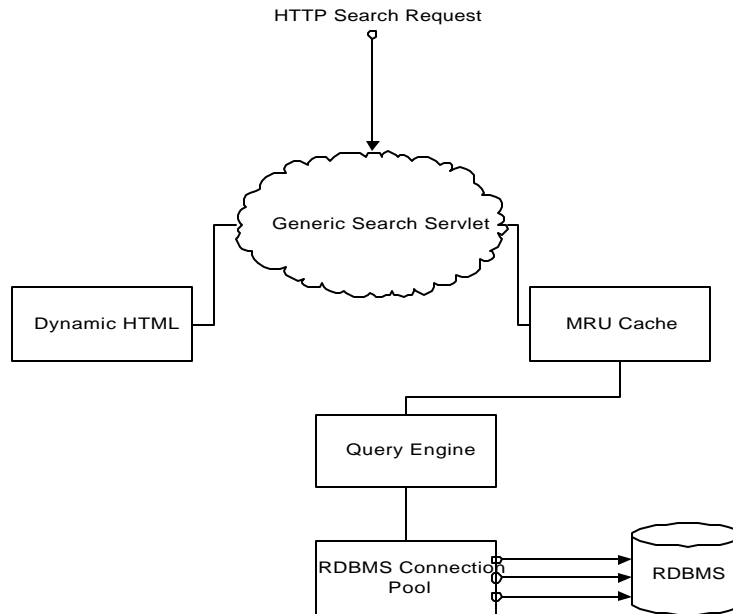
## 4.3 Query Software Architecture

The overall client request flow is illustrated in Figure 4. A user contacts the server via a web browser. The web server provides static information via standard HTML pages. Information requests are static html pages with forms. The form is filled out by the client and sent back to the server via an http get message. The HTTP server passes the message to the Servlet Engine. The Servlet Engine builds a SQL request to the RDBMS. The servlet then queries the database for the requested data and gets the results. The servlet builds a result page and passes it back to the client's browser for display.

As requests come into the servlet, the servlet builds SQL statements to get the needed information from the RDBMS. That request is passed on to a MRU (Most Recently Used) cache engine. The MRU cache keeps the most recently requested queries in memory. If the current request was previously made and its results are still in the cache, the results are returned to the servlet to be processed into an html results page. If the results are not found the query is passed down to the query engine. The query engine keeps a cache of database connections to speed up the request time. A pre-cached connection is pulled from the cache and the request is sent to the database. The results are returned to the query engine. The engine passes the results up to the MRU cache. The MRU cache stores the results in memory and returns a copy to the servlet. If the same request is made again, the results can be returned right from the MRU cache without going to the RDBMS. After the servlet has received the results, a dynamic HTML page is created. The formatted results are then returned to the client's browser for display. This query processing is illustrated in Figure 5.

## 5.0 Parallelism

We have investigated the use of a parallel database engine to support the SIRE backend. Experimental evaluation results presented in [Lundquist99] demonstrated that nearly a 92% parallel efficiency was achieved using an NCR Teradata database machine with 24 nodes. That is, roughly a 22-fold speed-up was observed as compared to a single node system.



**Figure 5. Servlet Flow Chart**

## 6.0 Conclusions

We provided an overview of the SIRE architecture. SIRE is implemented as an application of an RDBMS. This offers a single platform for the integration of searches across structured and unstructured data. In addition, it has the advantages of portability, scalability and the leverage of advances of the RDBMS industry. Currently, SIRE is deployed and in daily use in both commercial and governmental sites.

## Acknowledgments

This entire effort would not be possible without the contributions of the entire Information Retrieval Laboratory members. Special additional thanks are extended to Nazli Goharian for her commentary that improved the correctness and clarity of this paper. This work was supported in part by grants from the National Science Foundation, the National Institutes of Health, the US Army, and the IIT Research Institute.

## References

- [Frieder00] Frieder, O., et. al., "On the Development of the NIH Complementary and Alternative Medicine Digital Library," [www.ir.iit.edu/publications/2000/NCCAM.pdf](http://www.ir.iit.edu/publications/2000/NCCAM.pdf). Also available off of the CAM Citation Index link <http://nccam.nih.gov/nccam/databases.html>.
- [Grossman95] Grossman, D., D. Holmes, and O. Frieder, "A Parallel DBMS Approach to IR in TREC-3," *Overview of the Third Text Retrieval Conference*, NIST Special Publication 500-225, April 1995.
- [Grossman97] Grossman, D., O. Frieder, D. Holmes, and D. Roberts, "Integrating Structured Data and Text: A Relational Approach," *JASIS*, 48(2), February 1997.
- [Grossman98] Grossman, D. and O. Frieder, *Information Retrieval: Algorithms and Heuristics*, Kluwer Academic Publishers, ISBN 0-7923-8271-4, 1998.
- [Lundquist99] Lundquist, C., O. Frieder, D. Holmes, and D. Grossman, "A Parallel Relational DBMS Approach to Relevance Feedback in IR," *JASIS*, 50(5), April 1999
- [Robertson98] Robertson, S., S. Walker and M. Beaulieu, "Okapi at TREC-7: automatic ad hoc, filtering, VLC and interactive," *Overview of the Seventh Text Retrieval Conference*, 1998.
- [Singhal96] Singhal, A., C. Buckley and M. Mitra, "Pivoted Document Length Normalization," *Proceedings of Nineteenth SIGIR*, 1996.