

Ginix Generalized Inverted Index for Keyword Search

Dr.D.Devakumari¹, M.Shajitha Begum²

¹ Assistant Professor, PG and Research Department of Computer Science, Government Arts College(Autonomous), Coimbatore, Tamil Nadu, India

² Research Scholar ,Department of Computer Science , L.R.G Government Arts College For Women , Tirupur, Tamil Nadu, India

Abstract: Keyword search has become a ubiquitous method for users to access text data in the face of information explosion. Inverted lists are usually used to index underlying documents to retrieve documents according to a set of keywords efficiently. Since inverted lists are usually large, many compression techniques have been proposed to reduce the storage space and disk I/O time. However, these techniques usually perform decompression operations on the fly, which increases the CPU time. This paper presents a more efficient index structure, the Generalized INverted IndeX (Ginix), which merges consecutive IDs in inverted lists into intervals to save storage space. With this index structure, more efficient algorithms can be devised to perform basic keyword search operations, i.e., the union and the intersection operations, by taking the advantage of intervals. Specifically, these algorithms do not require conversions from interval lists back to ID lists. As a result, keyword search using Ginix can be more efficient than those using traditional inverted indices. The performance of Ginix is also improved by reordering the documents in datasets using two scalable algorithms. Experiments on the performance and scalability of Ginix on real datasets show that Ginix not only requires less storage space, but also improves the keyword search performance, compared with traditional inverted indexes.

Key Words: keyword search; index compression; document reordering

1. Introduction

With the huge amount of new information, keyword search is critical for users to access text datasets. These datasets include textual documents. Users use keyword search to retrieve documents by simply typing in keywords as queries. Current keyword search systems usually use an inverted index, a data structure that maps each word in the dataset to a list

of IDs of documents in which the word appears to efficiently retrieve documents.

To address this problem, this paper presents the Generalized INverted IndeX (Ginix), which is an extension of the traditional inverted index (denoted by InvIndex), to support keyword search. Ginix encodes consecutive IDs in each inverted list of InvIndex into intervals, and adopts efficient algorithms to support keyword search using these interval lists. Ginix dramatically reduces the size of the inverted index, while supporting keyword search without list decompression. Ginix is also compatible with existing d-gap-based compression techniques. As a result, the index size can be further compressed using these methods. Technique of document reordering[3-7], which is to reorder the documents in a dataset and reassign IDs to them according to the new order to make the index achieve better performance, is also used in this paper.

The contributions of this paper are:

- This paper presents an index structure for keyword search, Ginix, which converts inverted lists into interval lists to save storage space.
- Efficient algorithms are given to support basic operations on interval lists, such as union and intersection without decompression.
- Extensive experiments that evaluate the performance of Ginix are conducted. Results show that Ginix not only reduces the index size but also improves the search performance on real datasets.
- The problem of enhancing the performance of Ginix by document reordering is investigated, and two scalable and effective algorithms based on signature sorting and greedy heuristic of Traveling Salesman Problem (TSP)[3] are proposed.

2. Basic Concepts of Ginix

Let $D = \{d_1, d_2, \dots, d_N\}$ be a set of documents. Each document in D includes a set of words, and the

set of all distinct words in D is denoted by W . In the inverted index of D , each word $w \in W$ has an inverted list, denoted by I_w , which is an ordered list of IDs of documents that contain the word with all lists (ID lists and interval lists) sorted in ascending order. For example, Table 1a shows a collection of titles of 7 papers and Table 1b gives its inverted index.

The inverted index of this sample dataset consists of 18 inverted lists, each of which corresponds to a word. This example shows the lists of 4 most frequent words

Table 1 A sample dataset of 7 paper titles.
(a) Dataset content

ID	Content
1	Keyword querying and ranking in databases
2	Keyword searching and browsing in database
3	Keyword search in relational database
4	Efficient fuzzy type-ahead search
5	Navigation system for product search
6	Keyword search on spatial databases
7	Searching for hidden-web databases

(b) InvIndex		(c) Ginix	
Word	IDs	Word	IDs
Keyword	1,2,3,6	Keyword	[1,3],[6,6]
...
Databases	1,2,3,6,7	Databases	[1,3],[6,7]
Searching	2,7	Searching	[2,2],[7,7]
Search	3,4,5,6	Search	[3,6]
...

i.e., "keyword", "database", "searching", and "search" (word stemming is not considered)

A straightforward way to store an interval in Ginix is to explicitly store both its lower and upper bounds, as is illustrated in Table 1c. However, if an interval $[l, u]$ is a single-element interval, i.e., $l = u$, two integers are still needed to represent the interval. Thus if there are many single-element intervals in the interval list, the space cost will be expensive. The extra overhead for storing the interval lists is reduced by splitting each original interval list into 3 ID lists with one for single-element intervals and the other two for the lower and upper bounds of multi-element intervals. These three lists are denoted as S , L , and U . For example, the interval list $\{[1,1],[3,3],[6,7],[9,9],[12,15]\}$ can be split into 3 ID lists with $S = \{1,3,9\}$ and $L = \{6,12\}$. This reduces the number of integers from 10 to 7. Efficient sequential/sorted access is a basic requirement of keyword search based on the interval lists. Two position indicators, p and q , are used here to indicate the current positions in S and L/U . At the beginning, p and q are all set to 0, indicating that they are all pointing to the first elements in S and L/U . The current interval is found by comparing the two elements S_p and L_q . If S_p is smaller, we return the single-element interval $[S_p, S_p]$ and increment p by 1; if L_q is smaller, return the multi-element interval $[L_q, U_q]$ and increment q by 1.

Given an ID list S containing n IDs and its equivalent interval list R , the three lists, R, S, L , and S, R , used to store R will contain no more than n integers in total. This property of interval lists means that Ginix can be regarded as a compression technique, which is orthogonal to d-gap-based techniques. Moreover, d-gap-based compression algorithms, such as VBE and P For Delta, can still be applied to Ginix, since all the lists in Ginix are ordered lists of IDs.

3. Search Algorithms

A keyword search system usually supports union and the intersection operations on inverted lists. The union operation is a core operation to support OR query semantics in which every document that contains at least one of the query keywords is returned as a result. The intersection operation is

used to support AND query semantics, in which only those documents that contain all the query keywords are returned.

Traditional search algorithms are all based on ID lists.

Specifically, a traditional keyword search system first retrieves the compressed inverted list for each keyword from the disk, then decompresses these lists into ID lists, and then calculates the intersections or unions of these lists in main memory. This method introduces extra computational costs for decompression, and ID list based search methods can be very expensive because ID lists are usually very long.

3.1 Union operation

As in set theory, the union (denoted by U of a set of ID lists, denoted by $S = \{S_1; S_2; \dots; S_n\}$, is another ID list, in which each ID is contained in at least one ID list in S . Thus the union of a set of interval lists can be defined as follows:

Definition 1 Union of Interval Lists Given a set of interval lists, $R = \{R_1; R_2; \dots; R_n\}$, and their equivalent ID lists, $S = \{S_1; S_2; \dots; S_n\}$, the union of R is the equivalent R is the equivalent interval list of $U_{k=1}^n S_k$

For example, consider the following three interval lists: $\{[2,7], [11,13]\}$, $\{[5,7], [12,14]\}$, and $\{[1,3], [6,7], [9,9], [12,15]\}$. Their equivalent ID list are are $\{2; 3; 4; 5; 6; 7; 11; 12; 13\}$, $\{5; 6; 7; 12; 13; 14\}$, and $\{1; 2; 3; 6; 7; 9; 12; 13; 14; 15\}$. The union of these three ID lists is $\{1; 2; 3; 4; 5; 6; 7; 9; 11; 12; 13; 14; 15\}$; thus, the union of the three interval lists is the equivalent interval list of this ID list, i.e., $\{[1, 7] \cup [9, 9] \cup [11, 15]\}$

In this algorithm, the interval lists are first converted into ID lists with the union calculated using the well-known multi-way merge algorithm and the result then converted back into an interval list. This method is called the NAIVEUNION algorithm. Since the goal is to design an algorithm for calculating the union of interval lists without list conversion, this method will be used as a baseline for comparison.

3.2 Scan-line algorithm

A union algorithm without ID-interval conversion will only use the interval boundaries in the interval lists. Inspired by the scan-line rendering algorithm in computer graphics[9], the boundaries of all intervals

in the interval lists are first sorted into ascending order, with a scan-line moves from the smallest boundary to the largest boundary to calculate the union list. The scan-line movement maintains a reference counter to count the number of intervals that the scan-line is currently hitting. The counter is incremented by 1 when the scan-line hits a lower bound and is decremented by 1 when it hits an upper bound. If the counter increases from 0 to 1 (which means that the scan-line is processing an interval), the current boundary is saved in variable a . When the counter decreases from 1 to 0 (which means that the scan-line will not hit any interval before it hits another lower-bound), the current boundary is saved in variable b and $[a,b]$ is returned as the resulting interval.

The heap-based merge is used on all the interval lists to enumerate all the lower bounds and upper bounds in ascending order. This algorithm is called the SCANLINEUNION algorithm and illustrated in Algorithm 1

4.Document Reordering

Document reordering also improves the performance of Ginix. This section first explains the necessity of document reordering. Then, since finding the best order

Algorithm 1 Get-Interesting-Rules ($D, minsupp, minconf$)

```

1: // frequent subtrees
2:  $F_S = \text{FindFrequentSubtrees}(D, minsupp)$ 
3: ruleSet =  $\emptyset$ 
4: for all  $s \in F_S$  do
5: // rules computed from s
6: tempSet =  $\text{Compute-Rules}(s, minconf)$ 
7: // all rules
8: ruleSet = ruleSet  $\cup$  tempSet
9: end for
10: return ruleSet

```

Function 2 Compute-Rules ($s, minconf$)

```

1: ruleSet =  $\emptyset$ ; blackList =  $\emptyset$ 
2: for all  $c_s$ , subtrees of  $s$  do
3: if  $c_s$  is not a subtree of any element in blackList then
4:  $conf = \text{supp}(s) / \text{supp}(c_s)$ 
5: if  $conf \geq minconf$  then
6: newRule =  $\langle c_s, s, conf, \text{supp}(s) \rangle$ 
7: ruleSet = ruleSet  $\cup$  {newRule}
8: else
9: blackList = blackList  $\cup$   $c_s$ 
10: end if
11: end if
12: end for
13: return ruleSet

```

of documents is NP-hard, a sorting-based method and a sorting-TSP hybrid method are used to find near-optimal solutions.

Algorithm	PROBEISECT+ (\mathcal{R})
Input:	\mathcal{R} A set of interval lists.
Output:	G The resulting interval list.
1:	Sort \mathcal{R} in ascending order of list lengths
2:	for all $r \in R_1$ do Add CASPROBE($r, \mathcal{R} - R_1$) to G
3:	return G
4:	procedure CASPROBE(r, R)
Input:	r_0 A non-empty interval. \mathcal{R} An set of interval lists.
Output:	G The resulting interval list.
5:	$R_1^* \leftarrow \text{CASPROBE}(r_0, R_1)$
6:	for all $r \in R_1^*$ do
7:	$r^* \leftarrow [\max(\text{lb}(r_0), \text{lb}(r)), \min(\text{ub}(r_0), \text{ub}(r))]$
8:	if $n > 1$ then Add CASPROBE($r^*, \mathcal{R} - R_1$) to G
9:	else Add r^* to G
10:	return G
11:	end procedure

4.1 Necessity of document reordering

The time complexities of the search algorithms given in the previous section all depend on the number of intervals in the interval lists instead of the numbers of IDs. For example, the time complexity of the PROBEISECT algorithm is $O(m \cdot n \log m)$ where n denotes the number of interval lists and m denotes the number of intervals in each interval list. Thus, if the interval lists in Ginix contain fewer intervals, the search algorithms will be faster. On the other hand, interval lists containing fewer intervals will require less storage space. Therefore, the search speed and the space cost are both improved by reducing the number of intervals in Ginix.

Suppose that A and B are two ID lists with the same number of IDs. A's equivalent interval list will have less intervals than that of B's if A contains more consecutive IDs than B. Thus the order of the documents should be rearranged (or the IDs to the documents should be reassigned) so that the inverted lists contain as many consecutive IDs as possible. For example, if the 4th and 6th records in the dataset in Table 1a are switched, the interval lists for "keyword" and "databases" will become $\{[1,4]\}$ and $\{[1,4],[7,7]\}$. This will save two integers storage space for the interval lists.

There have been many efforts on finding the optimal document ordering that maximizes the frequencies of d-gaps in inverted lists to enhance the performance of existing inverted list compression techniques[3-7]. current problem is a special case of this problem (i.e., to maximize the frequencies of 1-gap). Previous studies of document reordering have all been designed for unstructured long documents (e.g., news and web pages), so methods are needed for structured or short documents, which are the focus of this study.

4.2 SIGSORT: Signature sorting method

The problem of document reordering is equivalent to making similar documents stay near to each other. Silvestri[5] proposed a simple method that sorts web pages in lexicographical order based on their URLs as an acceptable solution to the problem. This method is reasonable because the URLs are usually good indicators of the web page content. However, this method is not applicable to datasets whose URLs do not represent meaningful content (e.g., Wikipedia pages), or even do not have a URL field.

Other fields can also be used to represent the documents. For example, the reordering can use the Conf field (i.e., conference name) in the DBLP dataset. Sorting the documents by this field can also give acceptable results as well. However, a more flexible method is to generate a summary for each document and then sort the documents according to these summaries. Summaries can be generated as follows. First, all the words are sorted in descending order of their frequencies. Then, the top n (e.g., $n = 1000$) most frequent words are chosen as signature vocabulary. For each document, a string, called a signature, is generated by choosing those words belong to the signature vocabulary and sorting them in descending order of their frequencies. The document sorting compares each pair of signatures word-wise instead of comparing them letter-wise. This sort-based algorithm is called the SIGSORT algorithm.

Sorting documents by their signatures is effective because more frequent words are more likely to have consecutive IDs in its inverted list. In addition, since SIGSORT is very simple, it can easily handle large datasets.

SIGSORT is more effective for structured and short text data. Such data has more representative words since more records share the same words than general text data such as long web pages. As a result, each word in the signature vocabulary has a higher clustering power and the signatures are more effective. For general text data, a more effective method should consider more sophisticated summaries based on features other than words, such as categories and statistical topics

4.3 Scale TSP-based method using SIGSORT

Shieh et al.^[3] transformed the problem of finding the optimal ordering to the Traveling Salesman Problem (TSP). They built an undirected graph based on the underlying dataset by considering each document as a vertex and the number of words shared by the two documents as the weight of each edge. Finding an optimal ordering of documents is equivalent to solving the traveling salesman problem on this graph (i.e., to find a cycle on this graph that maximizes the sum of the weights of involved edges).

Finding an optimal cycle for TSP is NP-hard. Shieh et al.^[3] used the Greedy-Nearest-Neighbor (GNN) heuristic, which expands the path by adding a vertex that is closest to the current path, to find near-optimal solutions. The TSP-based method can provide good results for document reordering, but it can not scale to large datasets since solving the TSP using GNN heuristic on a complete graph with n vertexes has a time complexity of $O(n^2)$.

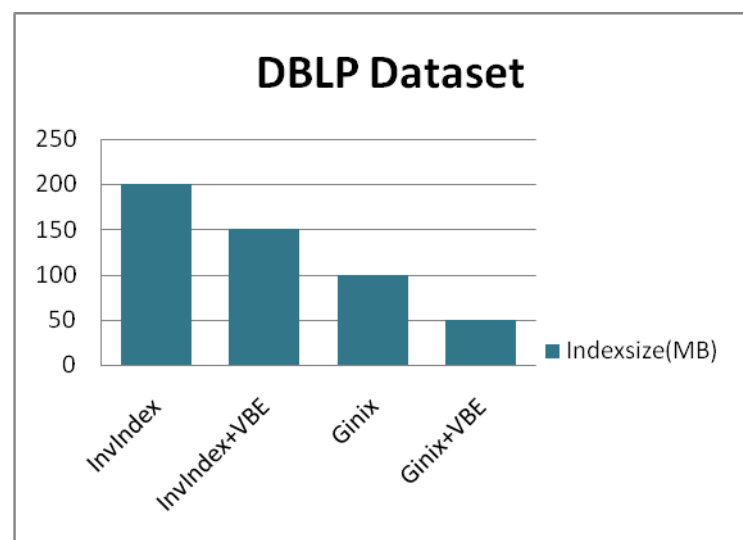
SIGSORT can be used to scale the TSP-based method to larger datasets, such as DBLP and PubMed datasets. First, all the documents are sorted according to their signatures using SIGSORT. Then, when the current path is expanded, the nearest vertex (document) is found within only a small set of candidates. Instead of the entire datasets, the candidate set for each document is the k consequent documents in the order obtained by SIGSORT. This method is called the SIGSORTTSP algorithm, which is more efficient than traditional TSP methods and which can be slightly better than pure SIGSORT for finding near-optimal solutions for the document reordering problem.

5.Experiments

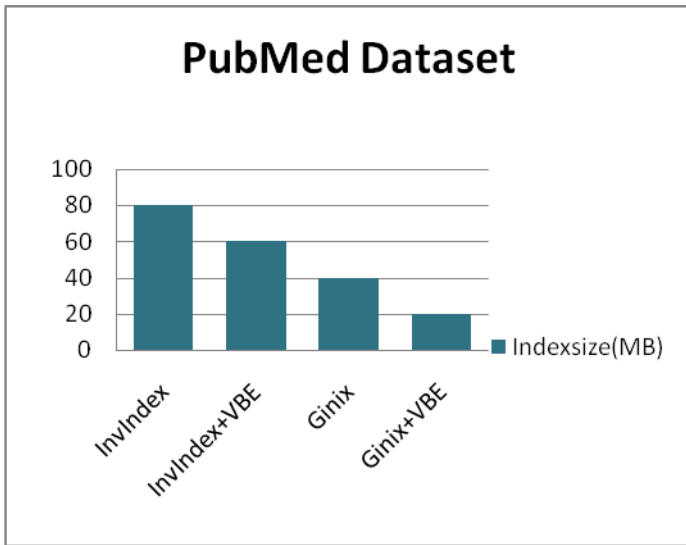
The performance and scalability of Ginix was evaluated by experiments on a Linux server with an Intel Xeon 2.50 GHz CPU and 16 GB RAM. Two datasets were used in the experiments, DBLP^[10] and PubMed^[11]. The DBLP dataset is a bibliography database on computer science that contains more than 1.4 million publications. The Title, Authors, Year, Conf (i.e., conference name), and URL of each publication were concatenated as a document with indexes built for these documents. PubMed is an extension of the MEDLINE database that contains citations, abstracts, and some full text articles on life sciences and biomedical topics. This study used 1.4 million articles with the Title, Journal Issue, and Journal Title attributes as the dataset. Ginix was implemented in C++ using the gcc compiler and /O3 flag.

5.1 Index size

Figure 1 shows the index sizes using different compression techniques. The widely-adopted VBE is used to evaluate the present technique of converting consecutive IDs to intervals in Ginix. Figure 1 compares the original inverted index (denoted by InvIndex), the inverted index compressed by VBE (denoted by InvIndex+VBE), the present inverted index (denoted by Ginix), and the present inverted index compressed by VBE (denoted by Ginix+VBE) for both the DBLP and PubMed datasets. The results show that the Ginix



(a)DBLP Data set



(b) PubMed Set

compression is much better than that of VBE. The Ginix+VBE result has the smallest index size.

5.2 Search performance

The performance of keyword search algorithms was compared using synthetic queries. Each dataset had 9 query workloads, each containing 1000 k-word queries, where $k = 2, 3, \dots, 10$. The keywords in each query were drawn according to their frequencies, in other words, if a keyword appears more frequently in the dataset, it is more likely to be drawn as a query keyword. The memory-based algorithms have their indexes in main memory without VBE compression. IDUNION, IDISECT-HEAP and IDISECT PROBE denote the three algorithms for union and intersection operations on InvIndex.

The results show that:

- SCANLINEUNION and TWINHEAPISECT, the two merge-based algorithms, are 30% and 20% faster than IDUNION and IDHEAPISECT.
- PROBEISECT runs 2 times faster than IDPROBESET, so interval list intersection is more efficient than ID list intersection.

- PROBEISECT+, the improved probe-based interval list intersection algorithm, runs faster than

PROBEISECT since many unnecessary probe operations are avoided. However, when there are many keywords in the query, the computation savings are not significant since the result list is already short.

Note that the naive probe-based intersection algorithm is very inefficient compared with the other probe-based intersection methods. As a result, it was omitted in these two figures for clarity.

Disk-based search algorithms introduce additional time to load the lists of inverted index (or Ginix) into main memory during query processing. In addition, if VBE is used on indexes, additional decompression operations must be performed, thus the overall query time gets longer compared with memory-based algorithms. Figure 4 shows the query processing times for probe-based intersections of InvIndex, InvIndex+VBE, Ginix, and Ginix+VBE on the DBLP and PubMed datasets. These four indexes are denoted as "I", "IV", "G", and "GV" in the figure

The results show that:

- **IO time:** The IO time of Ginix is approximately 30% shorter than that of InvIndex because the interval lists are shorter than the ID lists.
- **Decompression time:** Since the computational cost of VBE is proportional to the list length, the decompression time of Ginix+VBE is also approximately 30% shorter than that of InvIndex+VBE.
- **Search time:** Since the current algorithms take advantage of the intervals, the search time of Ginix is nearly 2x faster than that of InvIndex.

In summary, the overall performance of Ginix is much higher than that of InvIndex, with or without VBE compression.

5.3 Impact of document reordering

The impact of document reordering was evaluated for the DBLP dataset. The experiments considered four reordering methods: (1) RAND, which randomly shuffles the dataset; (2) CONF, which sorts the records according to the values of the Conf attribute; (3) SIGSORT, which uses the top 1000 most frequently occurring words as signature words; and (4) SIGSORTTSP, which uses 100 consequent records in the sorted list obtained by SIGSORT as the candidate set for each record ($k = 100$) and uses GNN heuristics to solve the TSP. The original InvIndex is used as a baseline. The method in Shieh et al.[3] was not evaluated because it can not scale to handle large datasets like DBLP. The index sizes and average query times are illustrated in Table 2.

The results in Table 2 show that:

- The size of Ginix is smaller than that of InvIndex, even when the records are ordered randomly.
- Sorting records according to their Conf values provides a good ordering, with which the index size is 128.9 MB and the average query time is 0.88 ms.
- Ginix can achieve the best performance in terms of both the index size and the average query time when reordering the records using SIGSORTTSP. Similar results were found for the PubMed dataset.

5.4 Scalability

The scalability of Ginix was evaluated using different numbers of records in the DBLP dataset. The index sizes for InvIndex and Ginix without VBE and the search speeds of the SCANLINEUNION+, TWINHEAPISECT, PROBEISECT, and PROBEISECT+

Table 2 Impact of document reordering (DBLP).

	Size (MB)	Time (ms)
BASELINE	200.30	1.47
RAND	172.90	1.80
CONF	128.90	0.88
SIGSORT	130.30	0.68
SIGSORTTSP	124.80	0.62

Overall query processing time of performing probe-based intersections.

6 Related Work

Keyword search is widely used by users to access text data with many studies in recent years. Keyword search is not only convenient for document collections but also for accessing structured or semi-structured data, such as relational databases and XML documents[12-19].

Inverted indexes are widely used to efficiently answer keyword queries in most modern keyword search systems, with techniques designed to compress the inverted indexes[20]. Most techniques first convert each ID in an inverted list to the difference between it and the preceding ID, called the d-gaps, and then encode the list using integer compression algorithms[1, 20-24]. Variable- Byte Encoding is widely used in systems since it is simple and provides fast decoding[1].

Other studies have focused on how to improve the compression ratio of inverted index using document reordering[4, 6, 7]. Here, if the document IDs are reassigned so that similar documents are close to each other, then there are more small d-gaps in the converted lists and the overall compression ratio is improved.

The interval tree[25] is widely used to directly Scalability for DBLP. calculate the unions and intersections of sets of intervals. However, interval trees are not good for keyword search because: (1) an interval tree is needed for each word, which increases the index size; (2) interval trees can not be easily compressed; and (3) interval trees can not support multi-waymerging and probing, which are important for accelerating calculations.

7. SCOPE FOR FUTURE ENHANCEMENT

Users can use any information system continuously if it is evaluating continuously time to time. For this reason this system is designed as flexible as possible to future enhancements. In future, well defined project metrics and statistical methods to evaluate

staff performance and to do better scheduling of project can be included.

In future ,I will study how to incrementally Update mined TARs when the original XML datasets change and how to further optimize our mining algorithm. Moreover for the moment I deal with a (substantial) fragment of XQuery , I would like to find the exact fragment of XQuery which lends itself to translation into intentional queries.

8. CONCLUSION

Mine all frequent association rules without imposing any a-priori restriction on the structure and the content of the rules. Store mined information in XML format. Use extracted knowledge to gain information about the original datasets. We have developed a C++ prototype that has been used to test the effectiveness of our proposal. We have not discussed the updatability Of both the document storing TARs and their index.

Reference Books

1. R.Agrawal and R.Srikant. Fast, "Algorithms for mining association rules in large databases". In Proc. of the 20th Int. Conf. on Very Large Databases', Morgan Kaufmann Publishers Inc., 1994.
2. T. Asai, K. Abe, S. Kawasoe, H. Arimura, H. Sakamoto, and S. Arikawa, "Efficient substructure discovery from large semi-structured data". In Proc. of the SIAM Int. Conf. on Data Mining, 2002.
3. T. Asai, H. Arimura, T. Uno, and S. Nakano, "Discovering frequent substructures in large unordered trees". In Technical Report DOI-TR216, Department of Informatics, Kyushu University, 2003.
4. E. Baralis, P. Garza, E. Quintarelli, and L. Tanca, "Answering xml queries by means of data summaries". ACM Transactions on Information Systems, 25(3):10, 2007.
5. D. Barbosa, L. Mignet, and P. Veltri. Studying the xml web, "Gathering statistics from an xml sample". World Wide Web, 8(4):413-438, 2005.

6. Ding, J. Attenberg, and T. Suel, Scalable techniques for document identifier assignment in inverted indexes. in Proc. of the 19th International Conference on World Wide Web, Raleigh, North Carolina, USA, 2010, pp. 311-320. Hao Wu et al.: Ginix: Generalized Inverted Index for Keyword Search 87 [8]

7. <http://www.i.kyushuu.ac>

8. <http://www.ieee.org>

9. <http://www.computer.org/publications/dlib>



Dr. D. Devakumari has received M. Phil degree from Manonmaniam Sundaranar University in 2003 and Ph.D from Mother Teresa Womens' University in 2013. Currently she is working as

Assistant Professor in the PG and Research Department of Computer Science, Government Arts College (Autonomous), Coimbatore, India. Her research papers have been published in International journals including Inderscience, Springer etc. She has presented papers in National and International Conferences. Her research interests include Data Pre-processing and Pattern Recognition.



Mrs. M. Shajitha begum has received B.SC(IT) degree from Cheran's Womens College and M.SC(IT) from K.S.G Arts and Science College .Pursuing her M.Phil degree from L.R.G Government College for Women.