TCS-TR-A-04-01

TCS Technical Report

Combinatorial Item Set Analysis Based on Zero-Suppressed BDDs

by

Shin-ichi Minato and Hiroki Arimura

Division of Computer Science

Report Series A

December 8, 2004

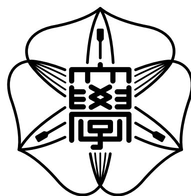

Hokkaido University
Graduate School of
Information Science and Technology

Email: minato@ist.hokudai.ac.jp

Phone: +81-011-706-7682
Fax: +81-011-706-7682

Combinatorial Item Set Analysis Based on Zero-Suppressed BDDs

Shin-ichi Minato
Division of Computer Science
Hokkaido University
North 14, West 9
Sapporo 060-0814, Japan

Hiroki Arimura
Division of Computer Science
Hokkaido University
North 14, West 9
Sapporo 060-0814, Japan

December, 8, 2004

(Abstract) Manipulation of large-scale combinatorial data is one of the important fundamental technique for web information retrieval, integration, and mining. In this paper, we propose a new approach based on BDDs (Binary Decision Diagrams) for database analysis problems. BDDs are graph-based representation of Boolean functions, now widely used in system design and verification area. Here we focus on Zero-suppressed BDDs (ZBDDs), a special type of BDDs, which are suitable for handling large-scale sets of combinations. Using ZBDDs, we can implicitly enumerate combinatorial item set data and efficiently compute set operations over the ZBDDs. We present some encouraging experimental results of frequent item set mining problem for practical benchmark examples, some of which have never been generated by previous method.

1 Introduction

Manipulation of large-scale combinatorial data is one of the fundamental technique for web information retrieval integration, and mining[16]. In particular, frequent item set analysis is important in many tasks that try to find interesting patterns from web documents and databases, such as association rules, correlations, sequences, episodes, classifiers, and clusters. Since the introduction by Agrawal et al.[2], the frequent item set and association rule analysis have been received much attentions from many researchers, and a number of papers have been published about the new algorithms or improvements for solving such mining problems[7, 9, 17].

In this paper, we propose a new approach based on BDDs (Binary Decision Diagrams) for database analysis problems. BDDs are graph-based representation of Boolean functions, now widely used in system design and verification area. Here we focus on Zero-suppressed BDDs (ZBDDs), a special type of BDDs, which are suitable for handling large-scale sets of combinations. Using ZBDDs, we can implicitly enumerate combinatorial item set data and efficiently compute set operations over the ZBDDs.

For a related work, *FP-Tree*[9] is recently received a great deal of attention because it supports fast manipulation of large-scale item set data using compact tree structure on the main memory. Our ZBDD-based method is a similar approach to handle sets of combinations on the main memory, but will be more efficient in the following points:

- ZBDDs are a kind of DAGs for representing item sets, while FP-Trees are tree representation for the same objects. In general, DAGs can be more compact than trees.
- ZBDD-based method provides not only compact data structures but also efficient item set operations written in a simple mathematical set algebra.

We present some encouraging experimental results of frequent item set mining problem for practical benchmark examples, some of which have never been generated by previous method.

Recently, the data mining methods are often discussed in the context of Inductive Databases[4, 12], the integrated processes of knowledge discovery. In this paper, we place the ZBDD-based method as a basis of integrated discovery processes to efficiently execute various operations finding interest patterns and analyzing information involved in large-scale combinatorial item set databases.

2 BDDs and ZBDDs

2.1 BDDs

BDD is a directed graph representation of the Boolean function, as illustrated in Fig. 1(a). It is derived by reducing a binary tree graph representing recursive *Shannon's expansion*, indicated in Fig. 1(b). The following reduction rules yield a *Reduced Ordered BDD (ROBDD)*, which can efficiently represent the Boolean function. (see [5] for details.)

- Delete all redundant nodes whose two edges point to the same node. (Fig. 2(a))
- Share all equivalent sub-graphs. (Fig. 2(b))

ROBDDs provide canonical forms for Boolean functions when the variable order is fixed. Most research on BDDs are based on the above reduction rules. In the following sections, ROBDDs will be referred to as BDDs (or ordinary BDDs) for the sake of simplification.

As shown in Fig. 3, a set of multiple BDDs can be shared each other under the same fixed variable ordering. In this way, we can handle a number of Boolean functions simultaneously in a monolithic memory space.

Using BDDs, we can uniquely and compactly represent many practical Boolean functions including AND, OR, parity, and arithmetic adder functions. Using Bryant's algorithm[5], we can efficiently construct a BDD for the result of a binary logic operation (i.e. AND, OR, XOR), for given a pair of operand BDDs. This algorithm is based on hash table techniques, and the computation time

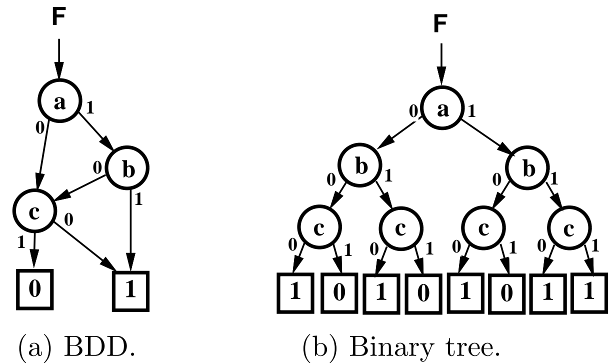

(a) BDD. (b) Binary tree.

Figure 1: BDD and binary tree for $F = (a \wedge b) \vee \overline{c}$.

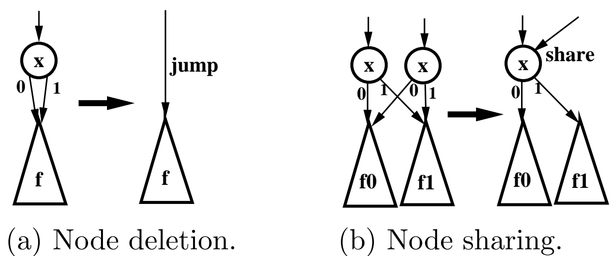

(a) Node deletion. (b) Node sharing.

Figure 2: Reduction rules of ordinary BDDs

is almost linear to the data size unless the data overflows the main memory. (see [14] for details.)

Based on these techniques, a number of BDD packages have been developed in 1990's and widely used for large-scale Boolean function manipulation, especially popular in VLSI CAD area.

2.2 Sets of Combinations and ZBDDs

BDDs are originally developed for handling Boolean function data, however, they can also be used for implicit representation of sets of combinations. Here we call "sets of combinations" for a set of elements each of which is a combination out of n items. This data model often appears in real-life problems, such as combinations of switching devices(ON/OFF), fault combinations, and sets of paths in the networks.

A combination of n items can be represented by an n-bit binary vector, $(x_1 x_2 \ldots x_n)$, where each bit, $x_k \in \{1, 0\}$, expresses whether or not the item is included in the combination. A set of combinations can be represented by a list of the combina-

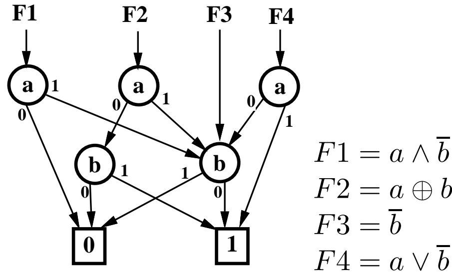

Figure 3: Shared multiple BDDs.

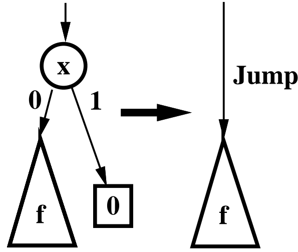

Figure 4: ZBDD reduction rule.

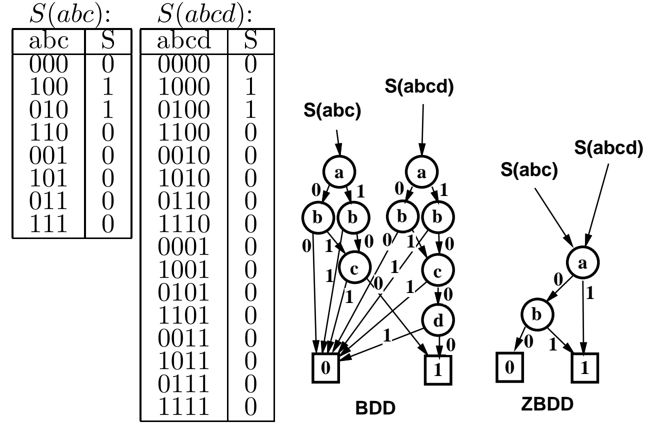

$S(abc)$:

abc	S
000	0
100	1
010	1
110	0
001	0
101	0
011	0
111	0

$S(abcd)$:

abcd	S
0000	0
1000	1
0100	1
1100	0
0010	0
1010	0
0110	0
1110	0
0001	0
1001	0
0101	0
1101	0
0011	0
1011	0
0111	0
1111	0

Figure 5: Example of ZBDD effect.

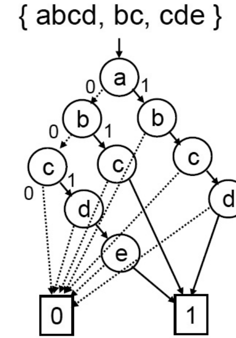

Figure 6: Explicit representation by ZBDD.

tion vectors. In other words, a set of combinations is a subset of the power set of n items.

A set of combinations can be mapped into Boolean space by using n-input variables for each bit of the combination vector. If we choose any one combination vector, a Boolean function determines whether the combination is included in the set of combinations. Such Boolean functions are called *characteristic functions*. The set operations such as union, intersection, and difference can be performed by logic operations on characteristic functions.

By using BDDs for characteristic functions, we can manipulate sets of combinations efficiently. They can be generated and manipulated within a time roughly proportional to the BDD size. When we handle many combinations including similar patterns (sub-combinations), BDDs are greatly reduced by node sharing effect, and sometimes an exponential reduction benefit can be obtained.

Zero-suppressed BDD (ZBDD)[13, 15] is a special type of BDDs for efficient manipulation of sets of combinations. ZBDDs are based on the following special reduction rules.

- Delete all nodes whose 1-edge directly points to the 0-terminal node, and jump through to the 0-edge's destination, as shown in Fig. 4.
- Share equivalent nodes as well as ordinary BDDs.

Notice that we do not delete the nodes whose two edges point to the same node, which used to be deleted by the original rule. The zero-suppressed deletion rule is asymmetric for the two edges, as we do not delete the nodes whose 0-edge points to a terminal node. It is proved that ZBDDs are also gives canonical forms as well as ordinary BDDs under a fixed variable ordering.

Here we summarise the features of ZBDDs.

- In ZBDDs, the nodes of irrelevant items (never chosen in any combination) are automatically

deleted by ZBDD reduction rule. In ordinary BDDs, irrelevant nodes still remain and they may spoil the reduction benefit of sharing nodes. (An example is shown in Fig. 5.)

- ZBDDs are especially effective for representing sparse combinations. For instance, sets of combinations selecting 10 out of 1000 items can be represented by ZBDDs up to 100 times more compact than ordinary BDDs.

- Each path from the root node to the 1-terminal node corresponds to each combination in the set. Namely, the number of such paths in the ZBDD equals to the number of combinations in the set. In ordinary BDDs, this property does not always hold.

- When no equivalent nodes exist in a ZBDD, that is the worst case, the ZBDD structure explicitly stores all items in all combinations, as well as using an explicit linear linked list data structure. An example is shown in Fig. 6. Namely, (the order of) ZBDD size never exceeds the explicit representation. If more nodes are shared, the ZBDD is more compact than linear list. Ordinary BDDs have larger overhead to represent sparser combinations while ZBDDs have no such overhead.

Table 1 shows the most of primitive operations of ZBDDs. In these operations, $\emptyset$, $\mathbf{1}$, *P.top* are executed in a constant time, and the others are almost linear to the size of graph. We can describe various processing on sets of combinations by composing of these primitive operations.

Table 1: Primitive ZBDD operations

"$\emptyset$"	Returns empty set. (0-termial node)
"$\mathbf{1}$"	Returns the set of only null-combination. (1-terminal node)
P.top	Returns the item-ID at the root node of P.
P.offset(v)	Selects the subset of combinations each of which does not include item v.
P.onset(v)	Selects the subset of combinations including item v, and then delete v from each combination.
P.change(v)	Inverts existence of v (add / delete) on each combination.
$P \cup Q$	Returns union set.
$P \cap Q$	Returns intersection set.
$P - Q$	Returns difference set. (in P but not in Q.)
P.count	Counts number of combinations.

Table 2: Statistics of typical benchmark data.

Data name	$\#I$	$\#T$	$avg\|T\|$	$avg\|T\|/\#I$
T40I10D100K	942	100,000	39	4.14%
mushroom	119	8,124	23	19.32%
BMS-WebView-1	497	59,602	2	0.40%
basket	13,103	41,373	9	0.06%

3 ZBDD-based Database Analysis

In this section, we discuss the method of manipulating large-scale item set databases using ZBDDs. Here we consider binary item set databases, each record of which holds a combination of items chosen from a given item list. Such a combination is called a *tuple* (or a *transaction*).

For analyzing those large-scale tuple databases efficiently, basic problems of data mining, such as *frequent item set mining*[3] and *maximum frequent item set mining*[6], are very important and they have been discussed actively in last decade. Recently, graph-based methods, such as FP-Tree[9], are received a great deal of attention, since they can quickly manipulate large-scale tuple data by constructing compact graph structure on the main memory. ZBDD technique is a similar approach to handle sets of combinations on the main memory, so we hope to apply ZBDD-based method effectively in this area.

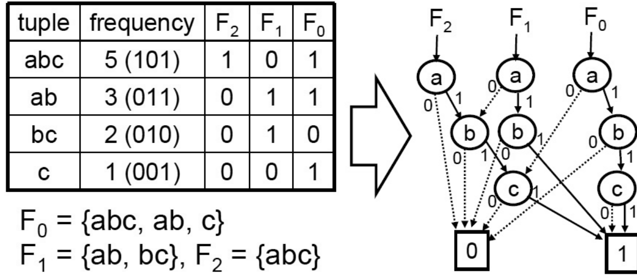

tuple	frequency	F2	F1	F0
abc	5 (101)	1	0	1
ab	3 (011)	0	1	1
bc	2 (010)	0	1	0
c	1 (001)	0	0	1

Figure 7: ZBDD vector for tuple-histogram.

3.1 Property of Practical Databases

Table 2 shows the basic statistics of typical benchmark data[7] often used for data mining/analysis problems. $\#I$ shows the number of items used in the data, $\#T$ is the number of tuples included in the data, $avg|T|$ is the average number of items per tuple, and $avg|T|/\#I$ is the average appearance ratio of each item. From this table, we can observe that the item's appearance ratio is very small in many cases. This is reasonable as considering real-life problems, for example, the number of items in a basket purchased by one customer is usually much less than all the items displayed in a shop. For another example, the number of links from one web page is much less than all the web pages in the network. This observation means that we often handle very sparse combinations in many practical data mining/analysis problems, and in such cases, the ZBDD reduction rule is extremely effective. If the average appearance ratio of each item is 1%, ZBDDs may be more compact than ordinary BDDs up to 100 times. In the literature, there is a first report by Jiang et al.[10] applying BDDs to data mining problems, but the result seems not excellent due to the overhead of ordinary BDDs. We must use ZBDDs in stead of ordinary BDDs for success in many practical data mining/analysis problems.

3.2 Tuple-Histograms based on ZBDDs

A *Tuple-histogram* is the table for counting the number of appearance of each tuple in the given database. In practical databases, the same tuple often appears many times. For example, "BMS-WebView-1" in Table 2 includes 59,602 records, and the most frequent tuple appears 1,533 times in the records. The top 10 frequent tuples appears 8,404 times in total. (Shares 14% in the records.)

Here we present a method of representing tuple-histograms by using ZBDDs. Since ZBDDs are representation of sets of combinations, a simple ZBDD distinguishes only existence of each tuple in the databases. In order to represent the numbers of tuple's appearances, we decompose the number into m-digits of ZBDD vector $\{F_0, F_1, \ldots, F_{m-1}\}$ to represent integers up to $(2^m - 1)$, as shown in Fig. 7. Namely, we encode the appearance numbers into binary digital code, as F_0 represents a set of tuples appearing odd times (LSB = 1), F_1 represents a set of tuples whose appearance number's second lowest bit is 1, and similar way we define the set of each digit up to F_{m-1}.

In the example of Fig. 7, The tuple frequencies are decomposed as: $F_0 = \{abc, ab, c\}$, $F_1 = \{ab, bc\}$, $F_2 = \{abc\}$, and then each digit can be represented by a simple ZBDD. The three ZBDDs are shared their sub-graphs each other.

Now we explain the procedure for constructing a ZBDD-based tuple-histgram from given tuple database. We read a tuple data one by one from the database, and accumulate the single tuple data to the histogram. More concretely, we generate a ZBDD of T for a single tuple picked up from the database, and accumulate it to the ZBDD vector. The ZBDD of T can be obtained by starting from "**1**" (a null-combination), and applying "Change" operations several times to join the items in the tuple. Next, we compare T and F_0, and if they have no common parts, we just add T to F_0. If F_0 already contains T, we eliminate T from F_0 and carry up T to F_1. This ripple carry procedure continues until T and F_k have no common part. After finishing accumulations for all data records, the tuple-histogram is completed.

Using the notation F.add(T) for addition of a tuple T to the ZBDD vector F, we describe the procedure of generating tuple-histogram F_T for given database D.

$F_T = \mathbf{0}$
forall $T \in D$ **do**

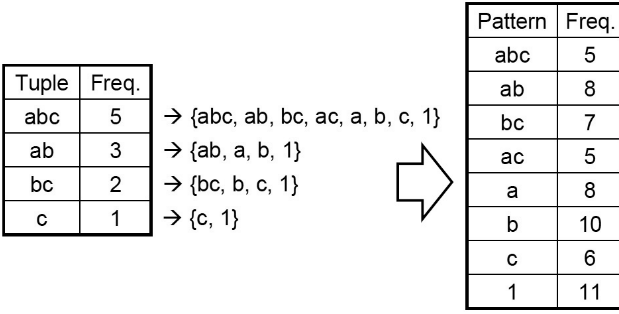

Tuple	Freq.	
abc	5	→ {abc, ab, bc, ac, a, b, c, 1}
ab	3	→ {ab, a, b, 1}
bc	2	→ {bc, b, c, 1}
c	1	→ {c, 1}

Pattern	Freq.
abc	5
ab	8
bc	7
ac	5
a	8
b	10
c	6
1	11

Figure 8: Tuple-histogram to pattern-histogram.

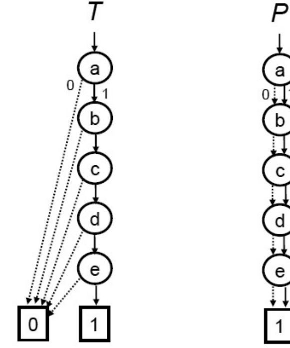

Figure 9: ZBDDs for a tuple and all sub-patterns

$F_T = F_T.\text{add}(T)$
return F_T

When we construct a ZBDD vector of tuple-histogram, the number of ZBDD nodes in each digit is bounded by total appearance of items in all tuples. If there are many partially similar tuples in the database, the sub-graphs of ZBDDs are shared very well, and compact representation is obtained. The bit-width of ZBDD vector is bounded by $\log S_{max}$, where S_{max} is the appearance of most frequent items.

Once we have generated a ZBDD-based tuple-histogram, it is easy to extract the set of frequent tuples which appears more than α times. By encoding the given threshold α into binary code, we can compose an algorithm of bit-wise arithmetic comparison between α and F_T based on ZBDD operations. After execution of those ZBDD operations, the result of frequent tuples can be obtained as a ZBDD. The computation time is almost linear to total ZBDD size.

3.3 Pattern-Histograms based on ZBDDs

In this paper, a *pattern* means a subset of items included in a tuple. A *pattern-histogram* is the table for counting the number of appearance of each patterns in any tuple in the given database. An example is shown in Fig. 8.

In general, a tuple of k items includes 2^k patterns, so computing a pattern-histogram is much harder than computing a tuple-histogram. In many cases, it is difficult to generate a complete pattern-histogram for a practical size of tuple database. Therefore, conventional methods extract only frequent patterns which appears more than α-times, for a given thresholds α, within a feasible computation time and space[7].

Using the ZBDD-based data structure, we may have more compact representation than previous methods, as a number of similar patterns can be shared in ZBDDs, and in some cases, this makes possible to generate complete pattern-histograms which have never succeeded in previous methods. Figure 9 shows a ZBDD for a tuple $T = abcde$ with five items and a ZBDD representing a set of all 32 patterns $P = \{1, a, b, c, d, e, ab, ac, bc, cd, abc, \ldots, abcde\}$ included in T. Clearly we can see that 2^k patterns in a k-item tuple can be represented by only k nodes of ZBDDs. As well as generating tuple-histograms, we can generate pattern-histograms by accumulating such a single ZBDD P for a set of patterns one by one, Here we summarise the procedure for computing a pattern-histogram F_P from a given database D as follows.

$F_P = \mathbf{0}$
forall $T \in D$ **do**
  $P = T$
  forall $v \in T$ **do**
    $P = P \cup P.\text{onset}(v)$
  $F_P = F_P.\text{add}(P)$
return F_P

Unfortunately, ZBDDs grows larger as repeating accumulations, and eventually may overflow

the memory for some large examples. While tuple-histgrams are bounded by the total items in the tuples, pattern-histograms are not bounded and so many patterns will be generated.

However, if we have succeeded in generating a ZBDD-based pattern-histogram for a given instance, we can enjoy very powerful data processing by using efficient ZBDD operations. It is interesting and important how large-scale instances we can generate complete pattern-histograms. The experimental results will be shown in later section.

In addition, we present an alternative procedure for generating ZBDD-based pattern-histograms. We can generate a pattern-histgram F_P from a complete tuple-histogram F_T.

$F_P = F_T$
forall $v \in F_T$ **do:**
  $F_P = F_P.\text{add}(F_P.\text{onset}(v))$
return F_P

We have not determined which algorithm is faster in practical environments. Anyway, the final form of ZBDD vectors must be the same if the two algorithms are computing for the same instance.

3.4 Utilities of Tuple/Pattern-Histograms

Once we generate tuple-/pattern-histograms using ZBDDs, various operations can be executed efficiently. We show several examples in this section. Suppose that we have obtained $F : \{F_0, F_1, \ldots, F_{m-1}\}$, the ZBDD vector representing a tuple- or pattern-histogram.

- We can efficiently extract a subset of tuples/patterns including a given item or sub-pattern P.

$S = \bigcup F_k$
forall $v \in P$ **do:**
  $S = S.\text{onset}(v).\text{change}(v)$
return S

Inversely, we can extract a subset of tuples/patterns not satisfying the given conditions. It is easily done by computing $\bigcup F_k - S$.

Table 3: Generation of tuple-histograms.

Data name	$\#T$	$total\|T\|$	|ZBDD|	Time(s)
T10I4D100K	100,000	1,010,228	552,429	43.2
T40I10D100K	100,000	3,960,507	3,396,395	895.0
chess	3,196	118,252	40,028	1.4
connect	67,557	2,904,951	309,075	58.1
mushroom	8,124	186,852	8,006	1.5
pumsb	49,046	3,629,404	1,750,883	188.5
pumsb_star	49,046	2,475,947	1,324,502	123.6
BMS-POS	515,597	3,367,020	1,350,970	895.0
BMS-WebView-1	59,602	149,639	46,148	18.3
BMS-WebView-2	77,512	358,278	198,471	138.0
accidents	340,183	11,500,870	3,877,333	107.0

After extracting a subset, we can quickly count a number of tuples/patterns by using a primitive ZBDD operation S.count. The computation time is linearly bounded by ZBDD size, not depending on the amount of tuple/pattern counts.

- For given α, we can extract all frequent tuples/patterns appearing more than α times. Computation time is almost linear to the ZBDD size. Repeating this procedure with different α's, we can determine the threshold α_m to pick up the top m frequent tuples/patterns. After generating ZBDD-based histograms, it is quite easy to extract frequent sets with different α's, while previous methods need almost recomputing again for each α.

- From ZBDD-based histograms, we can efficiently calculate indexes, such as *Support* and *Confidence*, which are often used in probabilistic/statistic analysis and machine learning area.

A feature of ZBDD-based method is to construct powerful data structure on the main memory, and we can interactively execute various queries to the database. Moreover, it is very interesting that the queries can be specified by mathematical set operations.

Table 4: Pattern-histogram for "mushroom".

ZBDD nodes	Time(s)	#Pattern
513,762	214.0	(> 2G)

Table 5: FP-Tree-based method for "mushroom".

Threshold α	Time(s)	#Pattern
81	22.60	91,273,269
40	67.96	295,117,613
16	244.06	1,176,182,553
8	494.05	1,983,493,667
4	891.31	(> 2G)
1	1,322.48	(> 2G)

4 Experimental Results

For evaluation of our method, we conducted experiments to construct ZBDD-based tuple- and pattern-histograms for typical benchmark examples[8] used in data mining/analysis problems.

We used a Pentium-4 PC, 800MHz, 512MB of main memory, with SuSE Linux 9. We can deal with up to 10,000,000 nodes of ZBDDs in this machine. Table 3 shows the results of generating tuple-histograms. In this table, $\#T$ shows the number of tuples, $total|T|$ is the total of tuple sizes (total appearances of items), and |ZBDD| is the number of ZBDD nodes for the tuple-histograms. We can see that tuple-histograms can be constructed for all instances in a feasible time and space. The ZBDD sizes are almost same or less than $total|T|$.

On the other hand, we succeeded in constructing the complete pattern-histogram within a given memory space, only for "mushroom", which is a relatively small instance (Table 4). In this case, the pattern-histogram requires about 65 times more ZBDD nodes than the tuple-histogram for the same data. In other words, if the tuple-histogram can be generated in a space tens or more times smaller than main memory size, it may be possible to generate a complete pattern-histogram in the main memory.

The "mushroom" pattern-histogram is implicitly representing at least 2,000,000,000 patterns. (The counter overflows the range of 32-bit integers.) The number of ZBDD nodes are 4,000 times smaller than number of patterns. This shows a great benefit of compression ratio obtained by ZBDDs.

To compare with a previous method, we applied a frequent pattern mining program based on FP-Tree[9], to extract the set of frequent patterns appearing more than α times, for the same example "mushroom". The results are shown in Table 5. For smaller α's, more patterns are extract, and the computation time increases up to hundreds or thousands of seconds. Notice that the FP-Tree-based method only extracts a set of frequent patterns for a given α, but does not generates a complete histogram for all possible patterns. Our ZBDD-based method generates a complete histogram for all patterns in 214 seconds, and it corresponds to computing frequent pattern sets for all α's at once. Consequently, ZBDD-based method is especially effective for handling the sets of huge number of patterns.

5 Conclusion

In this paper, we presented a new method of using ZBDDs for database analysis problems. Our work is just starting now, and we have many future works to be considered, such as ZBDD variable ordering problem for reducing graph size, and more efficient implementation of ZBDD set operations.

We expect that it would be too memory-consuming to construct ZBDDs of the complete pattern-histograms for the large-scale benchmarks, besides "mushroom". However, hopefully we will be able to handle those practical size of databases by using well-known improvement techniques, such as preprocessing of pruning not frequent items and patterns, or handling only maximum item set data[6], etc. ZBDD-based method will be useful as a fundamental techniques for various processing of database analysis, and will be utilized for web information retrieval and integration.

Acknowledgment

This research is partially supported by Grant-in-Aid from Ministry of Education, Culture, Sports, Science and Technology Japan.

References

[1] Akers, S. B., Binary decision diagrams, IEEE Trans. Comput., C-27, 6 (1978), 509–516.

[2] R. Agrawal, T. Imielinski, and A. N. Swami, Mining Association rules between sets of items in large databases, In P. Buneman and S. Jajodia, edtors, *Proc. of the 1993 ACM SIGMOD International Conference on Management of Data*, Vol. 22(2) of SIGMOD Record, pp. 207–216, ACM Press, 1993.

[3] R. Agrawal, H. Mannila, R. Srikant, H. Toivonen and A. I. Verkamo, Fast Discovery of Association Rules, In *Advances in Knowledge Discovery and Data Mining*, MIT Press, 307–328, 1996.

[4] J.-F. Boulicaut, Proc. 2nd International Workshop on Knowledge Discovery in Inductive Databases (KDID'03), Cavtat-Dubrovnik, 2003.

[5] Bryant, R. E., Graph-based algorithms for Boolean function manipulation, IEEE Trans. Comput., C-35, 8 (1986), 677–691.

[6] D. Burdick, M. Calimlim, J. Gehrke, MAFIA: A Maximal Frequent Itemset Algorithm for Transactional Databases, In Proc. ICDE 2001, 443–452, 2001.

[7] B. Goethals, "Survey on Frequent Pattern Mining", Manuscript, 2003. `http://www.cs.helsinki.fi/u/goethals/publications/survey.ps`

[8] B. Goethals, M. Javeed Zaki (Eds.), Frequent Itemset Mining Dataset Repository, Frequent Itemset Mining Implementations (FIMI'03), 2003. `http://fimi.cs.helsinki.fi/data/`

[9] J. Han, J. Pei, Y. Yin, R. Mao, Mining Frequent Patterns without Candidate Generation: A Frequent-Pattern Tree Approach, Data Mining and Knowledge Discovery, 8(1), 53–87, 2004.

[10] L. Jiang, M. Inaba, and H. Imai: A BDD-based Method for Mining Association Rules, in Proceedings of 55th National Convention of IPSJ, Vol. 3, pp. 397-398, Sept. 1997, IPSJ.

[11] Lee, C. Y., Representation of switching circuits by binary-decision programs, Bell Sys. Tech. Jour., 38 (1959), 985–999.

[12] H. Mannila, H. Toivonen, Multiple Uses of Frequent Sets and Condensed Representations, In *Proc. KDD*, 189–194, 1996.

[13] S. Minato: Zero-suppressed BDDs for set manipulation in combinatorial problems, In Proc. 30th ACM/IEEE Design Automation Conf. (DAC-93), (1993), 272–277.

[14] S. Minato: "Binary Decision Diagrams and Applications for VLSI CAD", Kluwer Academic Publishers, November 1996.

[15] S. Minato, Zero-suppressed BDDs and Their Applications, International Journal on Software Tools for Technology Transfer (STTT), Springer, Vol. 3, No. 2, pp. 156–170, May 2001.

[16] Ricardo Baeza-Yates, Berthier Ribiero-Neto, "Modern Information Retrieval", Addison Wesley, 1999.

[17] M. J. Zaki, Scalable Algorithms for Association Mining, IEEE Trans. Knowl. Data Eng. 12(2), 372–390, 2000.