

An Algorithms for Secure Mining Principle in Secure Database

CH. Nanda Krishna^{1*}, Y. Kalyan Chakravarthi², B.Sobhan Babu³

^{1,2,3}Assistant Professor, Department of Information Technology, VR Siddhartha Engineering College, Vijayawada, A.P, INDIA.

Assistant Professor, Department of Information Technology, Gudlavalleru Engineering College, A.P., India.

Email: nkcherukuri@gmail.com

Abstract—We don't forget the difficulty of discovering organization principles between gadgets in a colossal database of earnings transactions. We reward two new algorithms for solving the concern which can be essentially exclusive from the identified algorithms. Empirical analysis suggests that these algorithms outperform the recognized algorithms through motives starting from three for small problems to greater than an order of magnitude for massive problems. We additionally show how the nice points of the 2 proposed algorithms may also be mixed right into a hybrid algorithm, called Apriori Hybrid. Scale-up experiments show that Apriori Hybrid scales linearly with the quantity of transactions. Apriori Hybrid additionally has fine scale-up properties with admire to the transaction size and the number of gadgets within the database.

Key words: Privacy keeping knowledge Mining; disbursed Computation; standard Item sets; organization principles.

1. 1. INTRODUCTION

Growth in bar-code technology has made it viable for retail corporations to collect and store gigantic quantities of income information, known as the basket knowledge. A document in such data frequently includes the transaction date and the objects bought in the transaction. Effective firms view such databases as important portions of the marketing infrastructure. They're desirous about instituting know-how-driven marketing approaches, managed by way of database technology, that enable marketers to enhance and implement customized marketing applications and procedures [S]. The difficulty of mining association principles over basket knowledge was introduced in [4]. An instance of such a rule maybe that 98% of purchasers that purchase Tires and auto components additionally get car offerings achieved. Discovering all such ideas is valuable for crossmarketing and hooked up mailing functions. Other applications incorporate catalog design, add-on earnings, store layout, and patron segmentation situated on shopping patterns. The databases worried in these applications are very big. It's relevant, therefore, to have fast algorithms for this project.

The following is a proper statement of the drawback [4]: Let $Z = (i_1, i_2, \dots, i_m)$ be a suite of literals, known as items. Let \mathcal{T} be a suite of transactions, the place every transaction T is a collection of objects such that $T \subset Z$. Related to each transaction is a detailed identifier, known as its TID. We say that a transaction T involves X , a suite of some gadgets in Z , if $X \subset T$. In the absence of this type of relied on third celebration, it is needed to plan a protocol that the avid gamers can run on their possess so as to arrive at the required output y . The sort of protocol is considered flawlessly comfortable if no player can learn from his view of the protocol greater than what he would have learnt in the idealized setting where the computation is implemented via a relied on 1/3 social gathering. Yao [32] was once the first to advocate a popular answer for this concern within the case of two avid gamers. Different ordinary solutions, for the multi-occasion case, have been later proposed in [3], [5], [15].

In our obstacle, the inputs are the partial databases, and the desired output is the list of association principles that preserve in the unified database with help and self belief no smaller than the given thresholds s and c , respectively. As the above recounted generic solutions rely upon a description of the perform f as a Boolean circuit, they can be applied most effective to small inputs and features which can be realizable through easy circuits. In additional complicated settings, such as ours, other approaches are required for engaging in this computation. In such instances, some relaxations of the proposal of ideal safety probably inevitable when looking for useful protocols, furnished that the excess understanding is deemed benign (see examples of such protocols in e.G. [18], [28], [29], [31], [34]).

Kantarcioglu and Clifton studied that drawback in [18] and devised a protocol for its answer. The predominant part of the protocol is a sub-protocol for the cozy computation of the union of personal subsets that are held by way of the exceptional players. (The confidential subset of a given participant, as we give an explanation for under, entails the itemsets that are s -universal in his partial database.) That is probably the most pricey part of the protocol and its implementation relies upon cryptographic primitives equivalent to commutative encryption, oblivious transfer, and

hash functions. This is also the one part in the protocol where the players could extract from their view of the protocol knowledge on other databases, beyond what's implied via the ultimate output and their own enter.

Even as such leakage of understanding renders the protocol no longer perfectly relaxed, the perimeter of the surplus information is explicitly bounded in [18] and it is argued there that such knowledge leakage is innocuous, whence acceptable from a useful point of view.

Herein we endorse an alternative protocol for the comfortable computation of the union of personal subsets. The proposed protocol improves upon that in [18] in phrases of simplicity and effectivity as well as privateness. In unique, our protocol does not depend upon commutative encryption and oblivious transfer (what simplifies it vastly and contributes towards much reduced communiqué and computational costs). Even as our answer continues to be no longer flawlessly comfy, it leaks excess expertise simplest to a small number (three) of possible coalitions, unlike the protocol of [18] that discloses know-how also to a couple single players. Furthermore, we declare that the excess information 2 that our protocol may leak is much less touchy than the surplus knowledge leaked by the protocol of [18].

The protocol that we recommend here computes a parameterized loved ones of services, which we call threshold services, wherein the 2 extreme circumstances correspond to the problems of computing the union and intersection of personal subsets. These are in fact common-intent protocols that can be utilized in different contexts as well. A different challenge of secure multiparty computation that we solve here as a part of our discussion is the set inclusion problem; specifically, the obstacle where Alice holds a private subset of some floor set, and Bob holds an element within the ground set, they usually want to verify whether Bob's detail is within Alice's subset, with out revealing to either of them expertise in regards to the other occasion's input beyond the above described inclusion.

1.1 Main issue Decomposition and Paper institution

The challenge of discovering all association principles can be decomposed into two subproblems [4]:

1. In finding all units of gadgets (itemseis) which have transaction aid above minimal support. The support.
2. For an item set is the quantity of transactions that incorporate the item set. Item sets with minimal help are called massive itemsets, and all others small item sets. In section 2, we give new algorithms, Apriori and AprioriTid, for fixing this drawback. Use the gigantic itemsets to generate the favored rules. Right here is a straightforward algorithm for this mission. For each gigantic itemset 1, to find all non-empty subsets of 1. For every such subset a, output a rule of the form $a \Rightarrow (1 - a)$ if the ratio of aid(1) to help(a) is at least minconf We need to do not forget all subsets of 1 to generate principles with multiple consequents. Due to lack of area, we do not talk about this subproblem further, but refer the reader to [5] for a speedy algorithm.

In part three, we exhibit the relative performance of the proposed Apriori and AprioriTid algorithms towards the AIS [4] and SETM [13] algorithms. To make the paper self-contained, we incorporate an overview of the AIS and SETM algorithms in this section. We also describe how the Apriori and AprioriTid algorithms will also be combined into a hybrid algorithm, AprioriHybrid, and exhibit the scale up homes of this algorithm. We conclude by using declaring some related open issues in section four.

2. DISCOVERING MASSIVE ITEMSETS

Algorithms for discovering big itemsets make a couple of passes over the data. Within the first move, we rely the aid of character objects and investigate which of them are significant, i.E. Have minimumsupport. In every subsequent move, we start with a seed set of itemsets located to be enormous in the prior cross. We use this seed set for producing new potentially huge itemsets, known as candidate itemsets, and count the precise support for these candidate itemsets during the go over the data. At the end of the cross, we check which of the candidate itemsets are truly gigantic, they usually grow to be the seed for the next cross. This process continues unless no new massive itemsets are determined.

The Apriori and AprioriTid algorithms we recommend range essentially from the AIS [4] and SETM [13] algorithms in phrases of which candidate itemsets are counted in a move and in the way that these candidates are generated. In each the AIS and SETM algorithms, candidate itemsets are generated on-the-fly for the period of the go as data is being read. Particularly, after studying a transaction, it's decided which of the itemsets found enormous in the earlier cross are present within the transaction. New candidate itemsets are generated via extending these colossal itemsets with different objects in the transaction. However, as we will be able to see, the disadvantage is that this results in unnecessarily producing and counting too many candidate itemsets that prove to be small.

The Apriori and AprioriTid algorithms generate the candidate itemsets to depend in a go by means of utilising handiest the itemsets located massive in the earlier cross - without seeing that the transactions in the database. The elemental instinct is that any subset of a large itemset have to be colossal. For that reason, the candidate itemsets having okay objects can be generated via joining gigantic itemsets having okay - 1 items, and

deleting those who include any subset that's not big. This approach outcome in iteration of a a lot smaller number of candidate itemsets.

The AprioriTid algorithm has the further property that the database is not used at inquisitive about counting the aid of candidate itemsets after the primary go. Rather, an encoding of the candidate itemsets used in the prior pass is employed for this cause.

In later passes, the size of this encoding can emerge as a lot smaller than the database, as a consequence saving a lot reading effort. We will be able to give an explanation for these elements in additional detail once we describe the algorithms. Notation We assume that items in each transaction are stored sorted of their lexicographic order. It is simple to adapt these algorithms to the case the place the database 2, is stored normalized and each database file is a <TID, item> pair, where TID is the identifier of the corresponding transaction.

We call the quantity of items in an itemset its measurement, and phone an itemset of measurement okay a k-itemset. Items within an itemset are saved in lexicographic order. We use the notation $c[1] . C[2] C[k]$ to represent a kitemset c along with items $c[1], c[2], . . . C[k]$, where $c[1] < c[2] < . . . < c[k]$. If $c = X + Y$ and Y

is an m-itemset, we also name Y an m-eziension of X. Related to each itemset is a count subject to retailer the aid for this itemset. The count discipline is initialized to zero when the itemset is first created.

We summarize in table 1 the notation used within the algorithms. The set $i^?K$ is utilized by AprioriTid and can be further discussed after we describe this algorithm.

k-item set	An itemset having k items
Lk	Set of large k-items& Lk (those with minium support). Each member of this set haa two fields: i) itemset and ii) support count.
Ck	Set of candidate k-item&s ck (potentiahyally large itemsets). Each member of this set has two fields: i) itemset and ii) support count.
\bar{C}_k	Set of candidate k-itemsets when the TIDs of the generating transactions are kept associated with the candidates.

1.3 A running example

Let D be a database of N = 18 itemsets over a set of L = 5 items, $A = \{1, 2, 3, 4, 5\}$. It is partitioned between M = 3 players, and the corresponding partial databases are:

$D1 = \{12, 12345, 124, 1245, 14, 145, 235, 24, 24\}$

$D2 = \{1234, 134, 23, 234, 2345\}$

$D3 = \{1234, 124, 134, 23\}$.

For example, D1 includes N1 = 9 transactions, the third of which (in lexicographic order) consists of 3 items — 1, 2 and 4. Setting $s = 1/3$, an itemset is s-frequent in D if it is supported by at least $6 = sN$ of its transactions. In this case,

$F1s = \{1, 2, 3, 4\}$

$F2s = \{12, 14, 23, 24, 34\}$

$F3s = \{124\}$

$F4s = F5s = \emptyset$,

and $Fs = F1s \cup F2s \cup F3s$

s . For example, the itemset 34 is indeed globally s-frequent since it is contained in 7 transactions of D. However, it is locally s-frequent only in D2 and D3. 3 In the first round of the FDM algorithm, the three players compute the sets $C1,m s$ of all 1-itemsets that are locally frequent at their partial databases:

$C1,1 s = \{1, 2, 4, 5\}$,

$C1,2 s = \{1, 2, 3, 4\}$,

$C1,3 s = \{1, 2, 3, 4\}$.

Hence, $C1s = \{1, 2, 3, 4, 5\}$. Consequently, all 1-itemsets have to be checked for being globally frequent; that check reveals that the subset of globally s-frequent 1-itemsets is $F1s = \{1, 2, 3, 4\}$.In the second round, the candidate itemsets are:

$C2,1s = \{12, 14, 24\}$

$C2,2s = \{13, 14, 23, 24, 34\}$

$C2,3s = \{12, 13, 14, 23, 24, 34\}$.

(Note that 15, 25, 45 are locally s -frequent at $D1$ but they are not included in $C2,1s$ since 5 was already found to be globally infrequent.) Hence, $C2s = \{12, 13, 14, 23, 24, 34\}$.

Then, after verifying global frequency, we are left with $F2s = \{12, 14, 23, 24, 34\}$.

In the third round, the candidate itemsets are:

$C3,1s = \{124\}$, $C3,2s = \{234\}$, $C3,3s = \{124\}$.

So, $C3s = \{124, 234\}$ and, then, $F3s = \{124\}$. There are no more frequent itemsets.

```

-----
1) L1 = {large 1-itemsets};
2) for ( k = 2; Lk-1 # 0; k++ ) do begin
3) ck = applyGenerate(); // New candidates
4) forall transactions T E D do begin
5) C* = Support(T, ck); // Candidatea Contained in T
6) forall candidates c E C* do
7) c.count++;
8) end
9) Lk = {C E ck | C.count > minsup}
10) end
11) Answer = Union Lk;
-----

```

Figure 1: Algorithm Apriori

2.1.2 Subset operate

Candidate itemsets C_k are stored in a hash-tree. A node of the hash-tree either includes a list of itemsets (a leaf node) or a hash desk (an interior node). In an inside node, every bucket of the hash desk aspects to one more node. The foundation of the hash-tree is outlined to be at depth 1. An inside node at depth d features to nodes at depth $d+1$. Itemsets are stored in the leaves. Once we add an itemset c , we start from the root and go down the tree unless we attain a leaf. At an interior node at depth d , we come to a decision which department to follow

by way of applying a hash perform to the d th item of the itemset. All nodes are initially created as leaf nodes.

When the quantity of itemsets in a leaf node exceeds a distinct threshold, the leaf node is transformed to an inside node. Starting from the basis node, the subset perform finds the entire candidates contained in a transaction t as follows. If we are at a leaf, we discover which of the itemsets in the leaf are contained in t and add references to them to the answer set. If we are at an inside node and we've got reached it with the aid of hashing the item i , we hash on each and every item that comes after i in t and recursively apply this process to the node in the corresponding bucket. For the foundation node, we hash on each item in t . To see why the subset operate returns the favored set of references, do not forget what happens at the root node. For any itemset c contained in transaction t , the first item of c have got to be in t . On the root, by way of hashing on each object in t , we be certain that we most effective ignore itemsets that begin with an item no longer in t . An identical arguments follow at decrease depths. The one additional aspect is that, given that the gadgets in any itemset are ordered, if we attain the current node with the aid of hashing the object i , we handiest ought to keep in mind the items in t that occur after i .

2.2 Algorithm AprioriTid

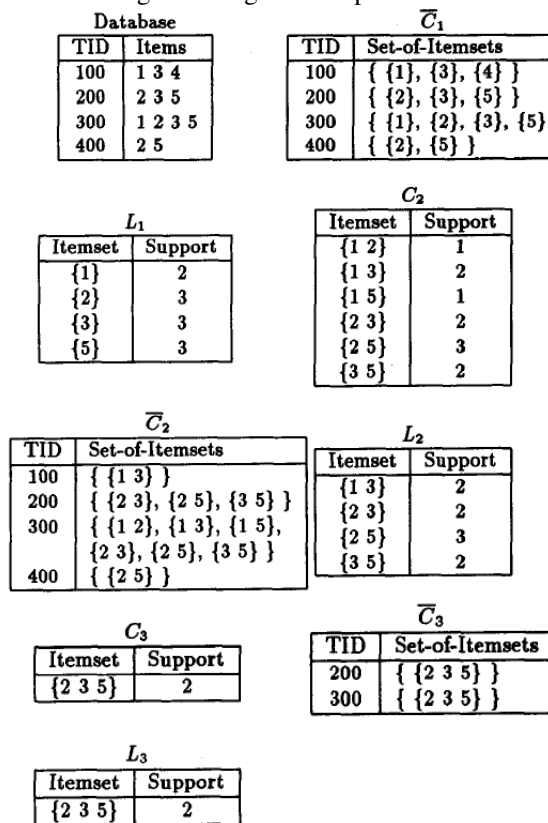
The AprioriTid algorithm, proven in figure 2, also makes use of the apriori-gen perform (given in part 2.1.1) to examine the candidate itemsets before the move starts. The interesting characteristic of this algorithm is that the database V isn't used for counting aid after the primary go. Instead, the set ck is used for this reason. Each and every member of the set ck is of the shape $\langle TID, \& \rangle$, the place each x_i is a possibly gigantic okay-object present within the transaction with identifier TID . For $ok = 1$, ck corresponds to the database V , despite the fact that conceptually each object is replaced with the aid of the itemset i . For $ok > 1$, ck is generated by the algorithm (step 10). The member of ck similar to transaction t is $\langle t.TID, c \in C_k \text{ IC contained in } t \rangle$. If a transaction does no longer incorporate any candidate ic -itemset, then ck will no longer have an entry for this transaction. As a consequence, the quantity of entries in ck could also be smaller than the quantity of transactions in the database, particularly for big values of k . Furthermore, for big values of ic , each and every entry is also smaller than the corresponding transaction considering that very few candidates could also be contained within the transaction. However, for small values for L , every entry may be larger than the corresponding transaction considering an entry in C_k involves all candidate ok -itemsets contained in the transaction.

In part 2.2.1, we give the information constructions used to put into effect the algorithm. See [5] for a proof of correctness and a discussion of buffer administration.

```

1) L1 = {large l-item&s};
2) G = database V;
3) for ( k = 2; Lkel # 8; k++ ) do begin
4) Ck = apriori-gen(Lk-1); // New candidates
5) Ek = B;
6) forall entries t E Ek-1 do begin
7) // determine candidate itemsets in Ck contained
// in the transaction with identifier L.TID
Ct = {c E Ck | (c - c[k]) E 'kset-of-itemsets A
(c - c[k - 1]) E &set-of-itemsets};
8) forall candidates c E Ct do
9) c.count++;
10) if (C, # 0) then ck += < t.TID, Ct >;
11) end
12) Lk = {c E ck | c.count 2 minsup}
13) end
14) Answer = uk Lk;
    
```

Figure 2: Algorithm AprioriTid



3. EFFICIENCY:

To determine the relative performance of the algorithms for locating massive sets, we carried out a number of experiments on an IBM RS/SOOO 530H computing device with a CPU clock fee of 33 MHz, sixty four MB of foremost reminiscence, and strolling AIX three.2. The info resided in the AIX file process and was stored on a 2GB SCSI three.5” drive, with measured sequential throughput of about 2 MB/2nd. We first give an outline of the AIS [4] and SETM [13] algorithms in opposition to which we examine the efficiency of the Apriori and AprioriTid algorithms. We then describe the unreal datasets used in the efficiency analysis and exhibit the performance

outcome. Sooner or later, we describe how the exceptional performance aspects of Apriori and AprioriTid may also be combined into an AprioriHybrid algorithm and demonstrate its scale-up houses.

3.1 The AIS Algorithm

Candidate itemsets are generated and counted on the fly because the database is scanned. After studying a transaction, it is decided which of the itemsets that had been found to be large within the prior pass are contained in this transaction. New candidate itemsets are generated by way of extending these gigantic itemsets with different gadgets within the transaction. A gigantic itemset I is increased with only these items which might be giant and arise later within the lexicographic ordering of gadgets than any of the objects in I . The candidates generated from a transaction are delivered to the set of candidate itemsets maintained for the cross, or the counts of the corresponding entries are elevated in the event that they were created through an prior transaction. See [4] for additional important points of the AIS algorithm.

3.2 The SETM Algorithm

The SETM algorithm [13] was encouraged by the wish to make use of SQL to compute large itemsets. Like AIS, the SETM algorithm also generates candidates onthefly established on transactions read from the database. It as a result generates and counts every candidate itemset that the AIS algorithm generates. However, to use the common SQL become a member of operation for candidate iteration, SETM separates candidate iteration from counting.

It saves a replica of the candidate itemset along side the TID of the generating transaction in a sequential structure. At the finish of the cross, the aid count of candidate itemsets depends upon sorting and aggregating this sequential constitution.

SETM remembers the TIDs of the generating transactions with the candidate itemsets. To avert wanting a subset operation, it makes use of this information to verify the large itemsets contained in the transaction read. Z_k is obtained by using deleting these candidates that would not have minimum aid. Assuming that the database is sorted in TID order, SETM can with no trouble in finding the significant itemsets contained in a transaction within the subsequent go with the aid of sorting & on TID. In truth, it wants to visit each member of & simplest as soon as in the TID order, and the candidate generation can be performed utilising the relational merge-join operation The drawback of this method is by and large due to the scale of candidate sets ck . For each and every candidate itemset, the candidate set now has as many entries as the quantity of transactions wherein the candidate itemset is gift. In addition, after we are able to rely the aid for candidate itemsets on the end of the move, i is within the mistaken order and needs to be sorted on itemsets. After counting and pruning out small candidate itemsets that should not have minimal help, the resulting set &! Wants yet another sort on TID before it may be used for generating candidates in the next go. Three.Three generation of artificial information We generated synthetic transactions to valuate the efficiency of the algorithms over a giant variety of knowledge characteristics. These transactions mimic the transactions in the retailing atmosphere. Our mannequin of the “real” world is that individuals tend to purchase sets of objects together.

Each and every such set is potentially a maximal giant itemset. An instance of any such set possibly sheets, pillow case, comforter, and ruffles. Nonetheless, some people may purchase only probably the most objects from such a set. For illustration, some people might buy most effective sheets and pillow case, and some handiest sheets. A transaction may just include a couple of huge itemset. For instance, a client would place an order for a dress and jacket when ordering sheets and pillow circumstances, the place the dress and jacket together form an extra tremendous itemset. Transaction sizes are as a rule clustered round an average and a few transactions have many objects. Traditional sizes of gigantic itemsets are also clustered round an average, with a couple of big itemsets having a huge quantity of gadgets, To create a dataset, our artificial data iteration application takes the parameters shown in desk 2.

$ D $	Number of transactions
$ T $	Average size of the transactions
$ I $	Average size of the maximal potentially large itemsets
$ L $	Number of maximal potentially large itemsets
N	Number of items

We first check the size of the following transaction. The scale is picked from a Poisson distribution with mean p equal to $|TI|$. Be aware that if each object is chosen with the identical probability p , and there are N objects, the anticipated number of gadgets in a transaction is given by means of a binomial distribution with parameters N and p , and is approximated by way of a Poisson distribution with imply Np .

We then assign objects to the transaction. Each and every transaction is assigned a sequence of probably giant itemsets. If the giant itemset available does not slot in the transaction, the itemset is put within the transaction

anyway in 1/2 the circumstances, and the itemset is moved to the subsequent transaction the leisure of the circumstances.

Giant itemsets are chosen from a collection I of such itemsets. The quantity of itemsets in ‘T is ready to ILj. There’s an inverse relationship between IL1 and the typical help for potentially giant itemsets. An itemset in T is generated by means of first determining the dimension of the itemset from a Poisson distribution with imply ~1 equal to III. Items in the first itemset are chosen randomly. To mannequin the phenomenon that big itemsets more often than not have fashioned gadgets, some fraction of objects in subsequent itemsets are chosen from the earlier itemset generated. We use an exponentially disbursed random variable with mean equal to the correlation level to make a decision this fraction for each itemset. The rest objects are picked at random. Within the datasets used within the experiments, the correlation degree used to be set to zero.5. We ran some experiments with the correlation level set to zero.25 and zero. Seventy five but did not in finding a lot difference in the nature of our performance outcome.

Name	T	I	D	Size in Megabytes
T5.I2.D100K	5	2	100K	2.4
T10.I2.D100K	10	2	100K	4.4
T10.I4.D100K	10	4	100K	
T20.I2.D100K	20	2	100K	8.4
T20.I4.D100K	20	4	100K	
T20.I6.D100K	20	6	100K	

3.4 Relative Performance

Figure 4 shows the execution times for the six synthetic datasets given in Table 3 for decreasing values of minimum support. As the minimum support decreases, the execution times of all the algorithms increase because of increases in the total number of candidate and large itemsets.

For SETM, we have only plotted the execution times for the dataset T5.I2.D100K in Figure 4. The execution times for SETM for the two datasets with an average transaction size of 10 are given in Table 4. We did not plot the execution times in Table 4 on the corresponding graphs because they are too large compared to the execution times of the other algorithms. For the three datasets with transaction sizes of 20, SETM took too long to execute and we aborted those runs as the trends were clear. Clearly, Apriori beats SETM by more than an order of magnitude for large datasets.




Figure 4: Execution times

Algorithm	Minimum Support				
	2.0%	1.5%	1.0%	0.75%	0.5%
Dataset T10.I2.D100K					
SETM	74	161	838	1262	1878
Apriori	4.4	5.3	11.0	14.5	15.3
Dataset T10.I4.D100K					
SETM	41	91	659	929	1639
Apriori	3.8	4.8	11.2	17.4	19.3

4. CONCLUSIONS AND FUTURE WORK

We offered two new algorithms, Apriori and AprioriTid, for locating all big association rules between items in a tremendous database of transactions. We in comparison these algorithms to the earlier known algorithms, the AIS [4] and SETM [13] algorithms. We offered experimental results, showing that the proposed algorithms continually outperform AIS and SETM. The performance hole elevated with the hindrance dimension, and ranged from a element of three for small issues to greater than an order of magnitude for huge issues.

One of the main ingredients in our proposed protocol is a novel secure multi-party protocol for computing the union (or intersection) of private subsets that each of the interacting players hold. Another ingredient is a

protocol that tests the inclusion of an element held by one player in a subset held by another. Those protocols exploit the fact that the underlying problem is of interest only when the number of players is greater than two.

REFERENCES:

1. D. S. Associates. The new direct marketing. Business One Irwin, Illinois, 1990.
2. R. Brachman et al. Integrated support for data archeology. In AAAI-93 Workshop on Knowledge Discovery in Databases, July 1993.
3. L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone. Classification and Regression Trees. Wadsworth, Belmont, 1984.
4. P. Cheeseman et al. Autoclass: A Bayesian classification system. In 5th Int'l Conf. on Machine Learning. Morgan Kaufman, June 1988.
5. D. H. Fisher. Knowledge acquisition via incremental conceptual clustering. Machine Learning, 2(2), 1987.
6. J. Han, Y. Cai, and N. Cercone. Knowledge discovery in databases: An attribute oriented approach. In Proc. of the VLDB Conference, pages 547-559, Vancouver, British Columbia, Canada, 1992.
7. M. Holsheimer and A. Siebes. Data mining: The search for knowledge in databases. Technical Report CS-R9406, CWI, Netherlands, 1994.
8. M. Ho&ma and A. Swami. Set-oriented mining of association rules. Research Report RJ 9567, IBM Almaden Research Center, San Jose, California, October 1993.
9. R. Krishnamurthy and T. Imielinski. Practitioner problems in need of database research: Research directions in knowledge discovery. SIGMOD RECORD, 20(3):76-78, September 1991.
10. P. Langley, H. Simon, G. Bradshaw, and J. Zytkow. Scientific Discovery: Computational Explorations of the Creative Process. MIT Press, 1987.
11. H. Mannila and K.-J. Raiha. Dependency inference. In Proc. of the VLDB Conference, pages 155-158, Brighton, England, 1987.
12. H. Mannila, H. Toivonen, and A. I. Verkamo. Efficient algorithms for discovering association rules. In KDD-94: AAAI Workshop on Knowledge Discovery in Databases, July 1994.
13. S. Muggleton and C. Feng. Efficient induction of logic programs. In S. Muggleton, editor, Inductive Logic Programming. Academic Press, 1992.
14. J. Pearl. Probabilistic reasoning in intelligent systems: Networks of plausible inference, 1992.
15. G. Piatetsky-Shapiro. Discovery, analysis, and presentation of strong rules. In G. Piatetskyshapiro, editor, Knowledge Discovery in Databases. AAAI/MIT Press, 1991.
16. G. Piatetsky-Shapiro, editor. Knowledge Discovery in Databases. AAAI/MIT Press, 1991.
17. J. R. Quinlan. C4.5: Programs for Machine Learning. Morgan Kaufman, 1993.