# Linear Queries in Statistical Databases 

M. D. SCHWARTZ<br>Tektronix, Inc.<br>and<br>D. E. DENNING AND P. J. DENNING<br>Purdue University


#### Abstract

A database is compromised if a user can determine the data elements associated with keys which he did not know previously. If it is possible, compromise can be achieved by posing a finite set of queries over sets of data elements and employing initial information to solve the resulting system of equations. Assuming the allowable queries are linear, that is, weighted sums of data elements, we show how compromise can be achieved and we characterize the maximal initial information permitted of a user in a secure system. When compromise is possible, the initial information and the number of queries required to achieve it is surprisingly small.


Key Words and Phrases: confidentiality, database security, data security, inference, linear query, secure query functions, statistical database.
CR Categories: 3.7

## 1. INTRODUCTION

Statistical databases contain sensitive information about individuals. Their purpose is to provide summary statistics about groups of people, while permitting only authorized access to the records of any one individual. But this objective is difficult to meet, for seemingly innocuous summaries contain vestiges of the original data. By correlating enough summaries, an intruder may infer confidential data. We understand much more about controlling access to individual records than we do about controlling inference.

A query is a function applied to a given subset of the records in the database; this subset is called the query set. Given the form of allowable queries and an assumption about the information known initially by inquirers, we wish to specify whether or not the database is secure. A proof that a database can be compromised amounts to exhibiting a method for constructing a sequence of queries whose responses imply the value associated with a given, arbitrary key. A proof that a database is secure is much more difficult, for it requires us to demonstrate that no such sequence exists.

[^0]Inference has long been a concern of census statisticians, who have developed sophisticated methods for choosing samples for analysis. The principle is to apply queries against a random subfile; deducing an element of the subfile is of little help because the relation between the subfile and the full population is unknown. Unfortunately, many contemporary applications center on small or medium databases where random sampling cannot be used.

Much existing literature on inference applies to counting queries, which return the number of individuals satisfying given characteristics, e.g., "Female astronomers over 40 years of age" [10-12, 15]. These studies have shown that the danger of compromise is greatest when the system responds for small query sets. Some protection is afforded when the system refuses answers for small query sets or for query sets whose differences are small $[1,7,16]$. However, this protection is far from complete [4, 17].

Another class of easily compromised queries is selection queries. Such a query returns a specific value from the query set, such as the largest, the smallest, or the median. These queries are easily compromised even if the database system somehow enforces a restriction that no two query sets overlap by more than one record [2, 3, 6, 14].

The subject of this paper is the class of linear queries. Each of these queries returns a weighted sum of elements in the query set. Our results extend the results of Dobkin, Jones, and Lipton [6], whose queries computed simple (unweighted) sums of elements in the query sets. One aspect of the results is striking: a user who knows as little as the value associated with a single key can often achieve a full compromise of the entire database-with a query sequence whose length is not much longer than the number of stored data elements.

As a running example, we will use the database displayed in Table I, referred to hereafter as the "runner's database." It contains facts about foot racers. Users of the database might be allowed to discover general properties such as average running paces; however, they are not supposed to be able to deduce facts about individual runners from these queries.

## 2. MODEL OF A DATABASE

A statistical database contains a set of records about individuals. Each record contains an identifier field, one or more category fields, and one or more data fields. Some data fields may also be category fields.

The identifier field contains a unique identifier for the record. Reading and writing in single records via identifiers is controlled by an access mechanism. A user presents the symbolic name of a record, the system looks up the identifier corresponding to that name and checks authorization.

A data element is a value stored in a given data field of a given record. The size of the database is the total number, $N$, of data elements in it. In a partitioned database, the data elements are divided into mutually exclusive, collectively exhaustive classes. The $j$ th class, denoted $Y_{j}$, contains $N_{j}$ data elements, and $N$ $=\sum_{j} N_{j}$. The columns of the runner's database are examples of classes. A key $p=(i, j)$ identifies a data element, which is denoted either $x_{p}$ or $x_{i j}$.

A query set is a collection of data elements identified by a list of keys or by a characteristic formula. (A characteristic formula is a logical formula which is matched against the values in category fields.) A query computes some value for

Table I. Runner's Database

| Keys | $\begin{array}{c\|} 1 \\ \operatorname{Max}^{1} \\ \text { VOX }^{\boldsymbol{a}} \end{array}$ | 2 <br> Train pace (seconds/mile) | $\begin{array}{cc} 3 & 4 \\ \text { Prior } 8 & \text { weeks } \end{array}$ |  | 5 <br> Fastest 1 mile (seconds) | $6$ <br> Fastest 10 miles (seconds) |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  | Total miles | Longest run |  |  |  |
| 1 Smith | 68 | 380 | 680 | 22 | 260 | 3183 | 1948 |
| 2 Jones | 61 | 405 | 530 | 18 | 287 | 3520 | 1950 |
| 3 Burns | 56 | 440 | 460 | 20 | 316 | 3818 | 1940 |
| 4 Cohen | 48 | 485 | 410 | 18 | 368 | 4447 | 1930 |
| 5 Cook | 49 | 470 | 375 | 20 | 360 | 4394 | 1962 |
| 6 Bloom | 54 | 440 | 430 | 20 | 328 | 3980 | 1948 |
| 7 King | 53 | 440 | 405 | 20 | 334 | 4072 | 1943 |
| 8 Frank | 72 | 370 | 705 | 23 | 251 | 2997 | 1947 |

" "Max VOX" is the maximum volume of oxygen processable in milliliters per kilogram body weight per minute.
the query set. To prevent trivial compromises a query program will not respond if the query set contains fewer than $k$ data elements, where $k \geq 2$ is a parameter of the system.

Example 1. The runner's database contains eight records and seven data fields, a total of $N=56$ data elements. The keys are pairs $(i, j)$ where $i$ is a row index and $j$ is a column index. The user would specify symbolic name pairs which would be translated to keys internally; for example (Smith, Max VOX) translates to $(1,1)$ and (Cohen, Year of birth) to (4, 7). The corresponding data elements are $x_{11}=68$ and $x_{47}=1930$. A possible key-specified query is "What is the average train pace for (Smith, Jones, Frank)?" The same query expressed with a characteristic formula is "What is the average train pace among runners whose Max VOX is greater than 60?"

This paper analyzes threats to statistical database security for key-specified query sets. Security for characteristic-specified query sets is treated in other papers [4, 17]. A summary of important results for both types of query sets is given in [5].

## Compromising the Database

A user has identified a key ( $i, j$ ) as soon as he knows the value of its data element $x_{i j}$. A user usually has some initial information about the database: he may know operational details of some query programs, or the values of some data elements. Relative to an assumption about initial information, a database is insecure if it is possible to identify keys not known initially. Identifying previously unknown keys is a compromise, and identifying all the keys is a full compromise. If the database is insecure there will exist a finite sequence of queries $q_{1}, \ldots, q_{m}$, query sets $X_{1}, \ldots, X_{m}$, and responses $v_{1}, \ldots, v_{m}$ for which it is possible to solve the system of equations

$$
\begin{gather*}
v_{1}=q_{1}\left(X_{1}\right) \\
v_{2}=q_{2}\left(X_{2}\right)  \tag{1}\\
\ldots \\
v_{m}=q_{m}\left(X_{m}\right)
\end{gather*}
$$

for some unknown key.
We are interested in characterizing the amount of initial information that can be tolerated in a secure system, and the amount of work of compromise in an insecure system.

Example 2. A regression analysis has yielded a formula to predict the time a runner will take to run a marathon [8]:

$$
\begin{equation*}
q(i)=\sum_{j=1}^{4} a_{j} x_{i j}+319 \text { seconds. } \tag{2}
\end{equation*}
$$

Suppose that this query is available in the database and that the user knows the coefficients $a_{j}$ (e.g. from reading the article). If that user also knows any three of $x_{i 1}, x_{i 2}, x_{i 3}$, and $x_{i 4}$, he can use the system's response to $q(i)$ to solve eq. (2) for the fourth.

## Linear Queries

To pose a query, a user specifies a list of distinct keys $p_{1}, \ldots, p_{k}$; the system determines the query set $\left\{x_{p_{1}}, \ldots, x_{p_{k}}\right\}$ and computes the value of the query function for it. To keep the notation simple, we will show the query set as argument to the query, rather than the list of keys. Thus $q\left(z_{1}, \ldots, z_{k}\right)$ denotes a query over a list of distinct keys that specify the data elements $z_{1}, \ldots, \dot{z}_{k}$. A linear query has the form

$$
\begin{equation*}
q\left(z_{1}, \ldots, z_{k}\right)=\sum_{j=1}^{k} a_{j} z_{j} \tag{3}
\end{equation*}
$$

for some fixed $k \geq 2$ and fixed "query weights" $a_{j}$. This form of query has been called a "weighted sum query" [18].

The assumption that the query set size is fixed at $k$ elements for all queries is a degenerate case of the restriction that each query set contains at least $k$ elements. When query set size can vary, as for characteristic-specified queries, compromise can be even simpler [4]. Our objective here is to show how easy compromise can be even when the intruder cannot exploit variation in query set size.

A sum query is a linear query with unit weights (all $a_{j}=1$ ). Such queries are used to compute averages. Dobkin, Jones, and Lipton [6] developed upper and lower bounds on the number of sum queries sufficient to compromise under the restriction that no two query sets can overlap by more than $r$ elements. For $N$ data elements with $r=1$ and no initial information, compromise is possible within $2 k-1$ queries if $N \geq N_{2}=k^{2}-k+1$; however, compromise is impossible when $N<N_{1}=\left(k^{2}+k\right) / 2$. Whether or not compromise is possible when $N_{1} \leq N$ $<N_{2}$ was not resolved. Davida et al. tightened these bounds [2].

Kam and Ullman studied the security of a class of key-matching sum queries [13]. Each key is a bit-string of $b$ bits, and there are $2^{b}$ data elements. To make a query, the user specifies the values of some set of $a$ bits, where $a \leq b$; the query program returns the sum of the data elements whose keys match in the given $a$ positions. The database is secure whenever users have no initial information, if $a$
$<b$, and if the ranges of values of the data elements are unknown. Otherwise compromise may be possible.

A partitioned query is a linear query in a partitioned database with the restriction that the $j$ th key specifies a data element from the $j$ th class; there are $k$ classes in all. Equation (2) is an example with four classes.

Our results for linear queries and partitioned queries are outlined in the next sections. The proofs are given in the Appendices.

## 3. SECURITY OF LINEAR QUERIES

We consider queries of the form of eq. (3) applied to a database of $N$ elements. All queries use the same value of $k$ and the same weights $a_{j}$.

We observe first that, knowing one weight $a_{j}$ and one data element $x$, we can compromise the entire database. Let $q_{1}$ and $q_{2}$ be the responses to the two queries

$$
\begin{aligned}
& q_{1}=q\left(z_{1}, \ldots, z_{j-1}, x, z_{j+1}, \ldots, z_{k}\right) \\
& q_{2}=q\left(z_{1}, \ldots, z_{j-1}, y, z_{j+1}, \ldots, z_{k}\right)
\end{aligned}
$$

which differ only in their use of data elements $x$ and $y$ in the $j$ th position. Equation (3) shows that

$$
\begin{equation*}
\left(q_{1}-q_{2}\right)=a_{j}(x-y), \tag{4}
\end{equation*}
$$

which can be solved for $y$. Now the entire database is vulnerable: we pose $q_{1}$ once and a new $q_{2}$ for each of the $N-1$ unknown data elements, effecting a full compromise with $N$ queries and $N-1$ applications of eq. (4). As soon as two data elements, $x$ and $y$, are known, we can also use eq. (4) to solve for any unknown weight.

Example 3. Suppose that the runner's database implements the query

$$
q\left(z_{1}, z_{2}, z_{3}\right)=0.2 z_{1}+0.5 z_{2}+0.3 z_{3} .
$$

A user knows that the weight of the first key is 0.2 and that Smith's Max VOX is 68 ; thus the initial information is

$$
\begin{aligned}
a_{1} & =0.2 \\
x_{11} & =68 .
\end{aligned}
$$

To determine Jones's Max VOX, the user proceeds as follows. He poses the two queries $q\left(x_{11}, x_{12}, x_{13}\right)$ and $q\left(x_{21}, x_{12}, x_{13}\right)$, to which the system will respond, respectively,

$$
\begin{aligned}
& q_{1}=407.6 \\
& q_{2}=406.2 .
\end{aligned}
$$

Equation (4) can be used to solve for $x_{21}$ :

$$
\begin{aligned}
x_{21} & =x_{11}-\left(q_{1}-q_{2}\right) / a_{1} \\
& =68-1.4 / 0.2 \\
& =68-7 \\
& =61
\end{aligned}
$$

which is Smith's Max VOX.

This reasoning suggests that security, if it is possible at all, must require keeping all the weights secret. But a more complex argument, given in Appendix A, shows that we can pose a set of queries which yield a set of equations in the unknowns $a_{j}$. If these equations are linearly independent, we can solve for the $a_{j}$ and then use the known value of one data element to compromise the rest of the database. In other words, knowing only one data element and none of the weights often suffices for a compromise. The exact statement of the result is:

Compromising Weighted Sums. Suppose that $x_{p}$ is known for some key p. With no more than $k(k+1)$ queries it may be possible to determine $\dot{x}_{p 1}, \ldots, x_{p k}$ for $k$ additional keys $p 1, \ldots, p k$. The remaining $N-k-1$ data elements can be deduced with an additional $N-k-1$ queries using the method of eq. (4). Thus full compromise is possible within $N+k^{2}-1$ queries.

The technique will be illustrated in an example below. These arguments lead to the conclusion that security depends on denying the inquirer any initial information about the database, for example, by restricting his access to records about which he knows nothing. In Appendix B we prove that such a restriction would work by showing the impossibility of solving any system of equations for any weight or any data element:
Secure Weighted Sums. If the inquirer knows no data element and none of the weights, the database is secure under his queries.

Unfortunately there is no way to be absolutely sure that a user knows nothing about the records to which his access has been restricted. Even if such an idea were enforceable, the presence of other types of queries can invalidate this result. The unweighted-sum compromise of Dobkin, Jones, and Lipton, for example, requires no initial knowledge [6].
The next example illustrates the method of compromising a database using weighted-sum queries when one data element is known.
Example 4. Let $y_{i}=x_{i 1}$ denote the Max VOX for runner $i$. Let $k=2$ and suppose $y_{1}=68$ is known. The database implements the query

$$
q\left(z_{1}, z_{2}\right)=2 z_{1}+3 z_{2}
$$

but the weights $\left(a_{1}, a_{2}\right)=(2,3)$ are unknown. Form three sets of two queries:

$$
\begin{array}{ll}
q_{10}=q\left(y_{2}, y_{3}\right)=290 & \mathcal{q}_{11}=q\left(y_{3}, y_{2}\right)=295 \\
q_{20}=q\left(y_{1}, y_{3}\right)=304 & q_{21}=q\left(y_{3}, y_{1}\right)=316 \\
q_{30}=q\left(y_{1}, y_{2}\right)=319 & q_{32}=q\left(y_{2}, y_{1}\right)=326
\end{array}
$$

The numbers show the system's responses. Let $A=a_{1}+a_{2}$ and observe that

$$
\begin{aligned}
& q_{10}+q_{11}=A\left(y_{2}+y_{3}\right)=585=Q_{1} \\
& q_{20}+q_{21}=A\left(y_{1}+y_{3}\right)=620=Q_{2} \\
& q_{31}+q_{31}=A\left(y_{1}+y_{2}\right)=645=Q_{3}
\end{aligned}
$$

whereupon

$$
\begin{aligned}
A=Q_{1} /\left(y_{2}+y_{3}\right) & =Q_{2} /\left(y_{1}+y_{3}\right)=Q_{3} /\left(y_{1}+y_{2}\right) \\
& \text { ACM Transactions on Database Systems, Vol. 4, No. 2, June } 1979 .
\end{aligned}
$$

This can be manipulated into two equations in unknowns $y_{1}$ and $y_{2}$ :

$$
\begin{aligned}
-Q_{2} y_{2}+\left(Q_{1}-Q_{2}\right) y_{3} & =-Q_{1} y_{1} \\
Q_{2} y_{2}-Q_{3} y_{3} & =\left(Q_{3}-Q_{2}\right) y_{1} .
\end{aligned}
$$

For the given numbers

$$
\begin{aligned}
-620 y_{2}-35 y_{3} & =-(585)(68) \\
620 y_{2}-645 y_{3} & =(25)(68) .
\end{aligned}
$$

It is easily verified that $y_{2}=61$ and $y_{3}=56$ is the solution. We may solve for $a_{1}$ and $a_{2}$ using the equations

$$
\begin{aligned}
& a_{1} y_{2}+a_{2} y_{3}=q_{10} \\
& a_{1} y_{1}+a_{2} y_{3}=q_{20} .
\end{aligned}
$$

To find any other $y_{i}$ we solve $q\left(y_{1}, y_{i}\right)=136+3 y_{i}$ once the system's response is known.

## 4. PARTITIONED QUERIES

We consider weighted-sum queries with the restriction that each weight $a_{j}$ is associated with a separate class $Y_{j}$ of data elements, and that the $j$ th key in a query's specification list must refer to the class $Y_{j}$. We denote the number of data elements in $Y_{j}$ by $N_{j}$, so that $N=\Sigma_{j} N_{j}$. The columns of the runner's database are examples of classes.

We observe first that the method of eq. (4) can be used to compromise all of class $Y_{j}$ if the weight $a_{j}$ and some element $x$ of that class both are known. In all, $N_{j}$ queries are required for this. If we know all the weights and one data element in each class, we can compromise the entire database with $N$ queries.

Suppose that we know all $k$ of the weights and data elements from only $k-1$ classes. We can determine an element from the unknown class (say $Y_{j}$ ) with one query using the known elements from the other classes (say $z_{1}, \ldots, z_{j-1}, z_{j+1}$, $\ldots, z_{k}$ ):

$$
\begin{equation*}
q\left(z_{1}, \ldots, z_{k}\right)=a_{1} z_{1}+\cdots+a_{k} z_{k} \tag{5}
\end{equation*}
$$

After the system responds with the value of this query, $z_{j}$ is the only unknown quantity in this equation. It follows that full compromise is possible with $N+1$ queries when all weights are known and data elements for all but one class are known.

Similarly, if we know one element from each class, and all but one of the weights, we can apply eq. (5) using the known $z_{1}, \ldots, z_{k}$; the missing weight is the only unknown quantity. Full compromise is again possible with $N+1$ queries.
The method illustrated in Example 4 was able to effect a compromise when none of the weights was known and the entire database was a single class. This method depends on permuting the same keys among different positions in the query specification lists. Partitioning the databases prevents these permutations and makes compromise more difficult. In Appendix C we prove that: If one knows all the weights, one data element from each of $k-2$ classes, and no data element
from the other two classes, then full compromise of the $k-2$ known classes is possible. The two unknown classes are secure.
The proof shows that no system of equations constructed contains enough information to solve for data elements in the unknown classes. The proof of this result shows further that every unknown class is secure provided that at least two classes are unknown initially.
These results show that keeping the weights secret is essential to security especially if users may know data elements from some of the classes. However, if other types of queries are possible in the same database, compromise may be possible without prior knowledge. For example, DeMillo, Dobkin, and Lipton have shown that $0\left(k^{2}\right)$ selection queries (median, maximum, minimum, etc.) can determine an element from an unknown class [3]; such queries can provide the initial information needed to compromise with partitioned queries (see also [17]).

## 6. CONCLUSIONS

Our results have shown that the maximum tolerable information to prevent compromise using weighted-sum or partitioned-weighted-sum queries is quite low. Security under weighted-sum queries can be assured only if users have no prior knowledge of any data element; the weights must be kept secret. Security in partitioned databases can be assured only if at least two of the weights or at least two of the classes are kept secret. It is not always possible to keep weights secret, and it may be impossible to enforce the restriction that the user knows no data element.

When the inquirer has sufficient initial information for a compromise, his task is easy. At worst, he needs to pose one query for each unknown data element in the (class of the) database which he is compromising, plus possibly a small additional but fixed number of queries. If the database supports other types of queries such as selection queries, the user may be able to deduce unknown data elements with small cost, before using algorithms of this paper for a compromise.

Compromise of statistical databases is usually possible, and it is cheap.

## APPENDIX A. COMPROMISING LINEAR QUERIES

We wish to show that knowing one data element, say $z_{1}$, may be sufficient to compromise a database under linear queries. The compromise proceeds in three parts. First, $k(k+1)$ queries are used to derive $k+1$ equations whose solution yields $k$ additional data elements $z_{2}, \ldots, z_{k+1}$. Then the $k+1$ known data elements are used to derive $k$ more equations whose solution is the weights $a_{1}, \ldots, a_{k}$. Finally, the method of eq. (4) is used to deduce the remaining $N-$ $(k+1)$ unknown data elements. Full compromise is achievable with $N-(k+1)$ $+k(k+1)=N+k^{2}-1$ queries in all.

Let $Z=\left(z_{1}, \ldots, z_{k+1}\right)$ be a list of data clements for which $z_{1}$ is known. Let $Z_{i}$ be $Z$ with $z_{i}$ deleted, and let $Z_{i j}$ denote the $j$ th one-step cyclic permutation of $Z_{i}$. (Note that $Z_{i 0}=Z_{i k}$.) Let $q_{i j}$ denote the response to the query $q\left(Z_{i j}\right)$, where $i=1$, $\ldots, k+1$ and $j=0, \ldots, k-1$. Note that $k(k+1)$ queries are needed to determine all the $q_{i j}$.

Each element of $Z_{i}$ is permuted into each of the $k$ positions by the queries $q\left(Z_{i 0}\right), \ldots, q\left(Z_{i, k-1}\right)$. If $S=z_{1}+z_{2}+\cdots+z_{k+1}$ then the sum of elements in $Z_{i}$ is $S-z_{i}$ and the sum of these permuted queries is

$$
Q_{i}=\sum_{j=0}^{k-1} q_{i j}=\sum_{j=0}^{k-1} a_{j}\left(S-z_{i}\right)=A\left(S-z_{i}\right)
$$

where $A=a_{1}+\cdots+a_{k}$. It follows that $Q_{i} /\left(S-z_{i}\right)=A=Q_{i+1} /\left(S-z_{i+1}\right)$. This gives a set of $k$ equations in the $k$ unknowns $z_{2}, \ldots, z_{k+1}$ :

$$
S\left(Q_{i}-Q_{i+1}\right)+Q_{i+1} z_{i}-Q_{i} z_{i+1}=0, \quad i=1, \ldots, k .
$$

If these equations are linearly independent, their solution will yield the values of $z_{2}, \ldots, z_{k+1}$. (If these equations are not linearly independent, nothing can be concluded about the security of the database, because another choice of $z_{2}, \ldots, z_{k+1}$ might produce a solvable set of equations.) We can use the known values of $z_{1}, \ldots, z_{k}$ to invert the equations corresponding to the queries $q\left(Z_{i 0}\right)$, $\ldots, q\left(Z_{i k}\right)$ and find the unknown coefficients:

$$
\left[\begin{array}{c}
Z_{1} \\
Z_{2} \\
\vdots \\
Z_{k}
\end{array}\right]\left[\begin{array}{c}
a_{1} \\
a_{2} \\
\vdots \\
a_{k}
\end{array}\right]=\left[\begin{array}{c}
q_{10} \\
q_{20} \\
\vdots \\
q_{k 0}
\end{array}\right]
$$

Now the method of eq. (4) can be applied to find the remaining data elements.

## APPENDIX B. SECURE LINEAR QUERIES

We suppose that the user has no initial knowledge about any data element. Our objective is proving that it is impossible for him to find enough linearly independent equations to solve for any unknown data element or query weight.

Consider any $m$ queries $q_{1}, \ldots, \boldsymbol{q}_{m}$, among which appear the data elements $Z$ $=\left\{z_{1}, \ldots, z_{s}\right\}$. We assume $m \geq k+s$ since we can always pose more queries without impairing our ability to compromise. We will apply the implicit function theorem (IFT) [9], which states that a set of functions $q_{1}, \ldots, q_{m}$ over variables $x_{1}, \ldots, x_{n}$ are mutually independent if the $\operatorname{det}(M) \neq 0$, where

$$
M=\left[\begin{array}{ccc}
\frac{\partial q_{1}}{\partial x_{1}} & \cdots & \frac{\partial q_{1}}{\partial x_{n}} \\
\vdots & & \vdots \\
\frac{\partial q_{m}}{\partial x_{1}} & \cdots & \frac{\partial q_{m}}{\partial x_{n}}
\end{array}\right]
$$

For our problem, the queries $q_{1}, \ldots, q_{m}$ depend on the $n=k+s$ quantities $a_{1}$, $\ldots, a_{k}$ and $z_{1}, \ldots, z_{s}$. The matrix is

$$
M=\left[\begin{array}{ccc:ccc}
z_{11} & \cdots & z_{1 k} & b_{11} & \cdots & b_{1 s} \\
& \cdots & & & \cdots & \\
z_{m 1} & \cdots & z_{m k} & b_{m 1} & \cdots & b_{m s}
\end{array}\right]
$$

where $z_{i j}$ is the member of $Z$ appearing in the $j$ th position of the $i$ th query, and

$$
b_{i t}= \begin{cases}a_{j,}, & \text { if } z_{i j}=z_{t} \\ 0, & \text { otherwise }\end{cases}
$$

Let $r=\operatorname{rank}(M)$. We observe that $r<k+s$ because column $k+p(p=1, \ldots, s)$ can be written as a linear combination of the others:

$$
\sum_{j=1}^{k} a_{j} z_{i j}-\sum_{\substack{t=1 \\ t \neq p}}^{s} z_{t} b_{i t}=b_{i p} z_{p}, \quad i=1, \ldots, m
$$

Note that $z_{i j}=z_{t}$ if and only if $b_{i t}=a_{j}$, which implies that the nonzero terms in the second sum cancel all but $b_{i p} z_{p}$ in the first. Since $r<k+s, \operatorname{det}(M)=0$, which implies that some of $q_{1}, \ldots, q_{m}$ are redundant.

Consider an $r \times r$ submatrix $N$ of $M$ formed by taking any $r$ linearly independent rows, and any $r$ linearly independent columns except column $k+p$. Renumber the queries so that $q_{1}, \ldots, q_{r}$ correspond to the rows of $N$. By construction $\operatorname{det}(N)$ $\neq 0$, but any determinant of $r+1$ rows and columns of $M$ must be zero. The IFT shows that $q_{1}, \ldots, q_{r}$ are mutually independent and that each query among $q_{r+1}$, $\ldots, q_{k+s}$ can be expressed in terms of $q_{1}, \ldots, q_{r}$ [9].
We will show that $x_{p}$ is not a function of any of the queries $q_{1}, \ldots, q_{r}$. To do this we treat $x_{p}$ as a new function of the same quantities ( $a_{j}$ or $z_{t}$ ) on which $q_{1}$, $\ldots, q_{r}$ depend according to the matrix $N$; we add the appropriate $r+1$ st row and column of $N$ :

$$
N^{\prime}=\left[\begin{array}{cc:c} 
& \frac{\partial q_{1}}{\partial x_{p}} \\
N & \vdots \\
& & \frac{\partial q_{r}}{\partial x_{p}} \\
\hdashline 0 & \cdots & 0
\end{array}\right]
$$

It is not hard to see that $\operatorname{det}\left(N^{\prime}\right)=(-1)^{r} \operatorname{det}(N) \neq 0$, which implies that $x_{p}$ is not a function of any quantity on which $q_{1}, \ldots, q_{r}$ depends and, therefore, that $x_{p}$ is not a function of any of $q_{1}, \ldots, q_{r}$.

## APPENDIX C. SECURE PARTITIONED QUERIES

We wish to show that if all weights are known but no element from at least two classes is known, then the unknown classes are secure. The classes from which we know at least one element are vulnerable to compromise via eq. (4).

It is sufficient to prove security when exactly two classes are unknown since additional unknown classes only make compromise more difficult. Let $Y_{1}$ and $Y_{2}$ be the two unknown classes, and let $y_{11}, \ldots, y_{1 N_{1}}$ be the elements of $Y_{1}$ and $y_{21}$, $\ldots, y_{2 N_{2}}$ be the elements of $Y_{2}$. Because it is assumed that data from the known classes are functionally independent of data in the unknown classes, $Y_{3}, \ldots, Y_{k}$ are irrelevant to compromising $Y_{1}$ and $Y_{2}$. Therefore, we need only to prove that
queries of the form

$$
q\left(z_{1} z_{2}\right)=a_{1} z_{1}+a_{2} z_{2}
$$

cannot reveal any element of $Y_{1}$ or $Y_{2}$.
Suppose that queries $q\left(z_{i 1}, z_{i 2}\right)$ have been posed for $i=1, \ldots, m$, yielding responses $q_{1}, \ldots, q_{m}$. These queries can be represented by the matrix equation

$$
\left[\begin{array}{cc:ccc}
a_{11} & \cdots & a_{1 N_{1}} & b_{11} & \cdots
\end{array} b_{1 N_{2}}\right]\left[\begin{array}{c}
y_{11} \\
\vdots \\
\cdots
\end{array}: \begin{array}{ccc} 
\\
a_{m 1} & \cdots & a_{m N_{1}} \\
y_{1 N_{1}} & b_{m 1} & \cdots
\end{array}\right]=b_{m N_{2}}\left[\begin{array}{c}
q_{1} \\
\vdots \\
q_{m}
\end{array}\right]
$$

where

$$
\begin{aligned}
& \alpha_{i j}= \begin{cases}a_{1}, & \text { if } z_{i 1}=y_{1 j} \\
0, & \text { otherwise }\end{cases} \\
& b_{i j}= \begin{cases}a_{2}, & \text { if } z_{i 2}=y_{2 j} \\
0, & \text { otherwise. }\end{cases}
\end{aligned}
$$

Now let $A_{j}$ represent the $j$ th column in the left part of the $p$ by $\left(N_{1}+N_{2}\right)$ coefficient matrix, and $B_{j}$ represent the $j$ th column of the right part. Note that

$$
\frac{1}{a_{1}} \sum_{j=1}^{N_{1}} A_{j}-\frac{1}{a_{2}} \sum_{j=1}^{N_{2}} B_{j}=0 .
$$

This means that the columns are not independent and that the rank of the coefficient matrix is less than $N_{1}+N_{2}$, i.e., the system of equations cannot be solved for any of the $N_{1}+N_{2}$ unknowns.

## REFERENCES

1. Chin, F.Y. Security in statistical databases for queries with small counts. ACM Trans. Database Syst. 3, 1 (March 1978), 92-104.
2. Davida, G.I., et al. Data base security. TR-CS-76-14, Dept. E.E. and Comptr. Sci., U. of Wisconsin, Madison, Wis., July 1976
3. DeMillo, R.A., Dobkin, D., and Lipton, R.J. Even data bases that lie can be compromised. IEEE Trans. on Software Engrg. SE-4, 1 (Jan. 1977), 73-75.
4. Denning, D.E., Denning, P.J., and Schwartz, M.D. The tracker: A threat to statistical database security. ACM Trans. Database Syst: 4, 1 (March 1979), 76-96.
5. Denning, D.E. Are statistical data bases secure? Proc. AFIPS 1978 NCC, Vol. 47, AFIPS Press, Montvale, N.J., pp. 525-530.
6. Dobkin, D., Jones, A.K., and Lipton, K.J. Secure databases: Protection against user inference. ACM Trans. Database Syst. 4, 1 (March 1979), 97-106.
7. Fellegi, I.P., and Phillips, J.L. Statistical confidentiality: Some theory and applications to data dissemination. Anals Econ. and Social Measmt. 3, 2 (April 1974), 399-409.
8. Foster, C., and Daniels, J. Running by the numbers. Runner's World (July 1975), 14-17.
9. Garsoux, J. Analyse Mathematique. Dunod, Paris, 1968, p. 444.
10. Haq, M.I. Security in a statistical data base. Proc. Amer. Soc. Inform. Sci. 11 (1974), 33-39.
11. HAQ, M.I. Insuring individual's privacy from statistical data base users. Proc. AFIPS 1975 NCC, Vol. 44, AFIPS Press, Montvale, N.J., pp. 941-946.

ACM Transactions on Database Systems, Vol. 4, No. 2, June 1979.
12. Hoffman, L.J., and Miller, W.F. Getting a personal dossier from a statistical data bank. Datamation 16, 5 (May 1970), 74-75.
13. Kam, J.B., and Ullman, J.D. A model of statistical databases and their security. ACM Trans. Database Syst. 2, 1 (March 1977), 1-10.
14. Reiss, S.P. Medians and data base security. Comptr. Sci. Dept., Brown U., Providence, R.I., Oct. 1977.
15. Schlörer, J. Indentification and retrieval of personal records from a statistical data bank. Methods of Inform. in Medicine 14, 1 (1975), 7-13.
16. Schlörer, J. Confidentiality of statistical records: A threat monitoring scheme for on-line dialogue. Methods of Inform. in Medicine 15, 1 (1976), 36-42.
17. Schwartz, M.D. Inference from statistical data bases. Ph.D. Th., Comptr. Sci. Dept., Purdue U., W. Lafayette, Ind., Aug. 1977.
18. Schwartz, M.D., Denning, D.E., and Denning, P.J. Securing data bases under linear queries. IFIP Congress 77, North-Holland Pub. Co., Amsterdam, 1977, pp. 395-398.

Received January 1977; revised October 1978


[^0]:    Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.
    This work was supported in part by the National Science Foundation under and by Purdue David Ross Grant 7994-56-13985.
    Authors' addresses: M.D. Schwartz, Tektronix, Inc., P.O. Box 500, Beaverton, OR 97077; D.E. Denning and P. J. Denning, Department of Computer Science, Purdue University, West Lafayette, IN 47907. © 1979 ACM 0362-5915/79/0600-0156 \$00.75

