# Preserving Confidentiality of High-dimensional Tabulated Data: Statistical and Computational Issues 

Adrian Dobra, Alan F. Karr and Ashish P. Sanil<br>National Institute of Statistical Sciences<br>Research Triangle Park, NC 27709-4006, USA

November 13, 2002


#### Abstract

Dissemination of information derived from large contingency tables formed from confidential data is a major problem faced by statistical agencies. In this paper we present solutions to several computational and algorithmic issues that arise in the dissemination of cross-tabulations (marginal sub-tables) from a single underlying table. These include data structures that exploit sparsity and support efficient computation of marginals as well as algorithms such as iterative proportional fitting, and a generalized form of the shuttle algorithm that computes sharp bounds on (small, confidentiality threatening) cells in the full table from arbitrary sets of released marginals. We give examples illustrating the techniques.


Keywords: Branch and bound; Contingency tables; Disclosure limitation; Integer programming; Marginal bounds; Shuttle algorithm.

## 1 Introduction

Statistical agencies, such as the US Census Bureau and Statistics Canada, disseminate immense amounts of tabular information derived from confidential microdata. Protecting this confidentiality (and thereby the privacy of the data subjects) is mandated by law; doing so while releasing as much useful information as possible presents major challenges to statistical agencies.

For the past three years, the National Institute of Statistical Sciences (NISS) has been developing systems for disclosure-limited dissemination of tabular summaries of confidential microdata. Such summaries consist of marginal sub-tables of a large contingency table constructed by "summing out" one or more attributes. The full table of frequency counts of data records with the same (categorical) attribute values is assumed not to be releasable. More important, many sub-tables are also not releasable because either on their own or in conjunction with other sub-tables they provide too much information about the full table. Often, and in this paper, "too much information" means that small count (and therefore high risk [16]) cells can be bounded too accurately on the basis of the released sub-tables.

[^0]Two classes of software systems have been developed [6, 9]. Table servers are "live," responding dynamically to incoming user queries for sub-tables of the full table, and assessing disclosure risk in light of previously answered queries. Table servers can be built at realistic scales, but defensible release rules and operating policies that the user community views as equitable are major impediments to their use in practice. Optimal tabular releases (OTRs), by contrast, are static releases of sets of sub-tables constructed by maximizing the amount of information released, as given by a measure of utility of that information, subject to a constraint on disclosure risk. Common underlying abstractions such as the query space and released and unreleasable sub-tables and frontiers are discussed in [6, 9].

In this paper we describe computational and algorithmic issues that must be confronted in order to build scalable implementations of table servers and OTRs. In the next section we introduce some basic notation that is needed to formally define the bounds problem. In $\S 3$ we describe the infrastructure-data structures and computational techniques-necessary to represent and operate on large (for example, 40-dimensional) contingency tables in the context of safe releases of sets of sub-tables. The common theme is sparsity: real, large tables are very sparse. $\S 4$ focuses on one method for computing sharp integer bounds based on any set of released sub-tables. The approach is based on the underlying hierarchical structure of the categorical data and includes a more general version of the shuttle algorithm [2], which we term the generalized shuttle algorithm. $\S 5$ contains a concluding discussion.

## 2 Terminology and Notation

Consider a $k$-way contingency table $\mathbf{n}=\{n(i)\}_{i \in \mathcal{I}}$ indexed by $\mathcal{I}=\mathcal{I}_{1} \times \cdots \times \mathcal{I}_{k}$, where each $\mathcal{I}_{j}=$ $\left\{1, \ldots, I_{j}\right\}$ is a finite set. With $A=\left\{i_{1}, \ldots, i_{l}\right\}$ an arbitrary subset of $K=\{1, \ldots, k\}$, let $\mathcal{I}_{A}=\mathcal{I}_{i_{1}} \times \cdots \times \mathcal{I}_{i_{l}}$. The $A$-marginal table of counts $\mathbf{n}_{A}=\left\{n_{A}\left(i_{A}\right)\right\}_{i_{A} \in \mathcal{I}_{A}}$ corresponding to $A$ is given by

$$
n_{A}\left(i_{A}\right)=\sum_{i_{K \backslash A} \in \mathcal{I}_{K \backslash A}} n\left(i_{A}, i_{K \backslash A}\right)
$$

We refer to $i$ (and $i_{A}$ ) as the coordinate of the table cell (and marginal table cell) it corresponds to. The corresponding $n(i)$ (and $n_{A}\left(i_{A}\right)$ ) are referred to as the cell count. The grand total of the table is $\mathbf{n}_{\emptyset}$. Also, let $N_{K}$ be the number of cells in $\mathbf{n}$, and similarly, let $N_{A}$ be the number of cells in a marginal subtable $\mathbf{n}_{A}$. Further, let $N_{K}^{+}$and $N_{A}^{+}$be the number of cells in $\mathbf{n}$ and $\mathbf{n}_{A}$ with positive counts (the non-zero cells).

Consider index sets $C_{1}, \ldots, C_{p}$. Suppose that for each $C_{j}$, we are given a table $\mathbf{n}^{j}$, and that these tables are consistent with each other: if $\mathbf{n}^{j_{1}}$ and $\mathbf{n}^{j_{2}}$ overlap $\left(C_{j_{1}} \cap C_{j_{2}} \neq \emptyset\right)$, then $\mathbf{n}_{C_{j_{1}} \cap C_{j_{2}}}^{j_{1}}=\mathbf{n}_{C_{j_{1}} \cap C_{j_{2}}}^{j_{2}}$, while if $\mathbf{n}^{j_{1}}$ and $\mathbf{n}^{j_{2}}$ do not overlap, then $n_{\emptyset}^{j_{1}}=n_{\emptyset}^{j_{2}}$. Denote by $\mathbf{T}\left(\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}\right)$ the set of all tables—termed feasible-whose $C_{1}, \ldots, C_{p}$-marginals are equal to $\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}$.

Note that $\mathbf{T}\left(\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}\right)$ might be empty even if $\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}$ are consistent with each other. "Good" algorithms (for example, to compute bounds) should first check whether there exists at least one integer table consistent with the constraints, and generate upper and lower bounds only if a feasible table exists.

Let

$$
U\left(i ; \mathbf{T}\left(\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}\right)\right)=\max \left\{\mathbf{x}(i): \mathbf{x} \in \mathbf{T}\left(\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}\right)\right\}
$$

and

$$
L\left(i ; \mathbf{T}\left(\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}\right)\right)=\min \left\{\mathbf{x}(i): \mathbf{x} \in \mathbf{T}\left(\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}\right)\right\} .
$$

be the maximum and minimum values of cell $i$ over all tables in $\mathbf{T}\left(\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}\right)$
We denote by $\mathcal{P}\left(\mathcal{I}_{r}\right)$ the set of all partitions of $\mathcal{I}_{r}$ and by $\mathcal{R D}$ the set of marginal tables that can be obtained by aggregating $\mathbf{n}$ not only across variables, but also across categories within variables. A table $\mathbf{n}^{\prime} \in \mathcal{R D}$ is uniquely determined from $\mathbf{n}$ by choosing $\mathcal{I}_{1}^{\prime} \in \mathcal{P}\left(\mathcal{I}_{1}\right), \ldots, \mathcal{I}_{k}^{\prime} \in \mathcal{P}\left(\mathcal{I}_{k}\right)$ :

$$
\mathbf{n}^{\prime}=\left\{t_{J_{1} \ldots J_{k}}:\left(J_{1}, \ldots, J_{k}\right) \in \mathcal{I}_{1}^{\prime} \times \cdots \times \mathcal{I}_{k}^{\prime}\right\},
$$

where

$$
t_{J_{1} \ldots J_{k}}=\sum_{i_{1} \in J_{1}} \cdots \sum_{i_{k} \in J_{k}} n_{K}\left(i_{1}, \ldots, i_{k}\right) .
$$

Let $\mathbf{T}$ denote the set of cells of all tables in $\mathcal{R D}$. If the set of cell entries in $\mathbf{n}$ that define a "super-cell" $t_{1} \in \mathbf{T}$ is included in the set of cells defining another "super-cell" $t_{2} \in \mathbf{T}$, we write $t_{1} \prec t_{2}$. Define a partial ordering " $\prec$ " on the cells in $\mathbf{T}$ by

$$
t_{J_{1}^{1} \ldots J_{k}^{1}} \prec t_{J_{1}^{2} \ldots J_{k}^{2}} \Leftrightarrow J_{1}^{1} \subseteq J_{1}^{2}, \ldots, J_{k}^{1} \subseteq J_{k}^{2} .
$$

With this partial ordering, $(\mathbf{T}, \prec)$ has a maximal element, namely the grand total of $\mathbf{n}$, and several minimal elements-the cell entries in the initial table $\mathbf{n}$.

If $t_{1}=t_{J_{1}^{1} \ldots J_{k}^{1}}$ and $t_{2}=t_{J_{1}^{2} \ldots J_{k}^{2}}$ are such that $t_{1} \prec t_{2}$ with $J_{r}^{1}=J_{r}^{2}$, for $r=1, \ldots, r_{0}-1, r_{0}+1, \ldots, k$ and $J_{r_{0}}^{1} \neq J_{r_{0}}^{2}$, we define the complement of the cell $t_{1}$ with respect to $t_{2}$ to be the cell $t_{3}=t_{J_{1}^{3} \ldots J_{k}^{3}}$, where for $r=1, \ldots, k$,

$$
J_{r}^{3}=\left\{\begin{array}{cl}
J_{r}^{1}, & \text { if } r \neq r_{0}, \\
J_{r}^{2} \backslash J_{r}^{1}, & \text { if } r=r_{0} .
\end{array}\right.
$$

In this case we write $t_{1} \oplus t_{3}=t_{2}$. The operator " $\oplus$ " is equivalent to joining two blocks of cells in $\mathbf{T}$ to form a third block. The blocks to be joined have to be composed from the same categories in $(k-1)$ dimensions and cannot share any categories in the remaining dimension.

## 3 Data Structures and Algorithms for Large, Sparse Tables

The simplest computer representation of a contingency table is as a multi-dimensional array of nonnegative integers. (Even though it is unwieldy to use multi-dimensional arrays in most programming languages, this structure can readily be implemented as a one-dimensional array by representing the multidimensional coordinates as mixed-radix integers [10].) However, the number of cells in a table increases exponentially with the table dimension, and so even moderate sized tables (10-20 dimensional) can have unmanageably large numbers of cells. For instance, a 14 -dimensional table derived from the Current Public Survey (CPS) data from 1993, which we have used as a test case, has 4.5 billion cells! The time and space requirements of processing such tables using a naive representation are clearly prohibitive. Fortunately, tables of real data are extremely sparse-the 14-dimensional table mentioned above with 4.5 billion cells has merely 76,000 cells with non-zero counts.

In this section we describe useful data structures and algorithms that exploit the extreme sparsity of the tables to enable us to perform certain operations on the tables. We first present a hash-table based structure for tables and outline algorithms for building the table from microdata and for generating marginal subtables (§3.1) followed by an illustration (§3.2) of how the basic structures and minor extensions can be used to perform a variant of Iterative Proportional Fitting (IPF) [1]. We note that AD-Trees [13] and certain Online Analytical Processing (OLAP) technologies [8] provide complementary methods for handling tables.

### 3.1 Building Tables and Marginal Tables

It is clear that any viable representation for large, sparse tables ( $N_{K} \gg N_{K}^{+}$) must exploit the table's sparseness. The most natural strategy is to store the location (coordinate) and content (count) of only the non-zero cells. This leads to an enormous saving of space, since we only have to store $N_{K}^{+}$items instead of $N_{K}$ items. However, the naive storage of a list of (coordinate, count) pairs suffers from the disadvantage that retrieving the count of an arbitrary cell entails a search through the list of cells. Hence, this query, which could be processed in $O(1)$ time with the multidimensional array representation, now takes $O\left(N_{K}^{+}\right)$time (or $O\left(\log N_{K}^{+}\right)$if the list is sorted). Fortunately, we can achieve the space-saving of a list with the fast access time of an array by using a hash table.

Hash tables are well-known data structures that support efficient storage and retrieval for sets of (key, value) pairs-(coordinate, count) pairs in our case. In a hash table, the data are stored in an array. The essential component is a hash function that maps every key (coordinate) to a location in the array and allows us almost instantly to access data corresponding to a given key. There are many design issues pertaining to construction of a good hash function, selection of a good array size, and several implementation tricks that could be employed. We do not cover the details here since this is a widely-used and well-studied data structure [3,10], and has implementations incorporated as standards in computer languages like Java (part of the "Java collections framework" [15]) and C++ (part of the standard library [14]).

Three essential aspects are relevant to us. First, tables and marginal subtables are stored in hash tables where the key is a $K$-variate generalized coordinate, $\left(x_{i}, \ldots, x_{K}\right)$ with $x_{j} \in \mathcal{I}_{j} \cup\{0\}$ and the value is the corresponding cell count. (We call it a generalized coordinate since it can represent cells in both the full table and any marginal subtable, with the key for a cell in a marginal table $\mathbf{n}_{A}$ being a $K$-variate coordinate with $x_{j}=0$ if $j \notin A$ ). Second, the hash function is defined over the set of all generalized coordinates.

Third, the following operations are supported efficiently:

1. AddCount(Tab, coord): Increments by one the count for the cell with coordinate coord in table or marginal table Tab.
2. GetCount(Tab, coord): Retreives the cell count for the cell with coordinate coord in table or marginal table Tab.
3. GetMarginalCoord(M, coord): Returns the generalized coordinate in the marginal subtable M corresponding to coord in the full table (if $\mathrm{M} \equiv \mathbf{n}_{A}$, then this can be computed by setting $x_{j}=0$ for $j \notin A$ ).

We assume the the full table fits into the computer's main memory. This is a reasonable assumption for modern computers in cases such as survey data, where the number of subjects, $\mathbf{n}_{\varnothing}$, is seldom more than an order of $10^{5}$ and since $N_{K}^{+}<\mathbf{n}_{\varnothing}$ and, usually, $N_{K}^{+} \ll \mathbf{n}_{\emptyset}$. It is clear that the hash table representation of the contingency table can be easily constructed by reading through the microdata records sequentially and incrementally updating the table using AddCount.

Given the hash table representation, computation of an arbitrary marginal sub-table is easy:

```
Pseudocode 3.1 ComputeMARGINAL(Table,Marginal)
    for each coord in Table do
        marginalCoord = GETMARGINALCOORD(Marginal,coord)
        AdDCount(Marginal, marginalCoord)
    end for
```


### 3.2 Example: IPF

We illustrate here how the hash table representation can be used to develop efficient algorithms to process large, sparse tables, using Iterative Proportional Fitting (IPF) as an example. Intuitively, IPF finds the "best" reconstruction of the full table consistent with a set of marginal tables of the original full table. More precisely, IPF calculates the maximum likelihood estimate for a log-linear model defined on variables on the full table whose minimal sufficient statistics are the given by the set of marginal tables [1].

The algorithm involves iteratively refining estimated cell values for the table. Each iteration consists of stepping through the list of marginal tables and scaling the current cell estimates to make the current table estimate consistent with the marginal table. (To illustrate, for a two-way table, entries are alternately rescaled row-wise to make the row sums "correct" and then column-wise to make the column sums "correct.") Specifically, consider cell $i$ with $\hat{n}(i)$ as its current estimate. Let $n(+)$ be the cell count in the marginal table currently under consideration to which $i$ contributes to, and let $\hat{n}(+)$ be the sum of all current estimates of cells like $i$ that contribute to the marginal total. Then $\hat{n}(i)$ is adjusted as $\hat{n}(i) \leftarrow[n(+) / \hat{n}(+)] \hat{n}(i)$. After adjusting all cells for a given marginal, the current table estimate will be consistent with the marginal under consideration [1].

The IPF variant we present here only considers non-zero cells-this is equivalent to the original IPF with structural or known zeroes. Implementing it involves a minor extension the basic data structure. The (key, value) pair in the hash table now consists of key $=$ coordinate as before and value $=$ (count,fit) $=(n, \hat{n})$, where $n \equiv n(i), \hat{n} \equiv \hat{n}(i)$ for the table, and $n \equiv n(+), \hat{n} \equiv \hat{n}(+)$ for the marginal table. We define procedures AddFit(Tab, coord) and GetFit analogous to AddCount and GetCount in §3.1. We also define a procedure UpDATEFIT(Tab, coord, fit) that sets the "fit" component to the value of fit.

The IPF employs two subroutines MAKEMARGINALSUMS(Table,Marginal) (Pseudocode 3.2) and AdjustTable(Table,Marginal) (Pseudocode 3.3) that compute the $\hat{n}(+)$ s and the $\hat{n}(i)$ s respectively.

The IPF algorithm can be implemented as outlined in Pseudocode 3.4. This version trades off time efficiency for memory space savings by recomputing marginal tables every time they are required. Since this strategy maintains only two tables in memory at any given time, it allows us to handle arbitrarily large

```
Pseudocode 3.2 MAKEMARGINALSUMS(Table,Marginal)
    for each coord in Table do
        marginalCoord = GETMARGINALCOORD(MarginalTable,coord)
        AdDFIT(Marginal, marginalCoord)
    end for
```

```
Pseudocode 3.3 ADJUSTTABLE(Table,Marginal)
    for each coord in Table do
        marginalCoord = GeTMARGINALCoord(Marginal,coord)
        marginalFit = GetFIT(Marginal,marginalCoord)
        tableFit = GetFit(Table,coord)
        marginalCount = GETCouNT(Marginal,marginalCoord)
        tableFit \leftarrowtableFit*(marginalCount/marginalFit)
        UpDATEFIT(Table,coord)
    end for
```

sets of marginals. It is also possible to precompute the marginal tables and avoid the ComputeMarginal in the inner loop if both the set of marginal tables and the full table can fit into main memory.

```
Pseudocode 3.4 IPF(Table,ListOfMarginalNames)
    repeat
        for each Marginal in ListOfMarginalNames do
            ComputEMARGINAL(Table,Marginal)
            MAKEMARGINALSumS(Table,Marginal)
            AdJustTABLE(Table,Marginal)
        end for
    until Estimates converge
```


## 4 The Generalized Shuttle Algorithm

The fundamental idea behind the "shuttle" algorithm is that the upper and lower bounds for the cells in a table T, based on knowledge of an arbitrary set of marginal sub-tables, are interlinked. Our method builds on [2], which treats the problem of a 3-way table given the three 2 -way marginals, and sequentially improves the bounds for cells of interest until no further adjustment can be made.

Denote by $L(t)$ and $U(t)$ the current lower and upper bounds for the "super-cell" $t \in \mathbf{T}$. (See §2) for details of the notation.) Let $L(\mathbf{T})=\{L(t): t \in \mathbf{T}\}$ and $U(\mathbf{T})=\{U(t): t \in \mathbf{T}\}$. If $\mathcal{N} \subset \mathbf{T}$ is the set of cells in $\mathbf{x} \in \mathbf{T}\left(\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}\right)$, then $L(\mathcal{N})$ and $U(\mathcal{N})$ are the bounds arrays to be determined. Every $t \in \mathbf{T}$ has a value $V(t)$ assigned to it. If $t$ corresponds to an entry in a fixed marginal, we "know" the value $V(t)$ of that entry, so we set both current bounds of $t$ to $V(t)$.

Let $\mathbf{T}_{0}$ be the set of cells in $\mathbf{T}$ for which the lower bound is currently equal to the upper bound. If we fix all the cells in $\mathcal{N}$ at a certain value, then all the remaining cells in $\mathbf{T}$ will also be fixed: $\mathcal{N} \subset \mathbf{T}_{0}$ implies $\mathbf{T}=\mathbf{T}_{0}$. Consider $\mathbf{M} \subset \mathbf{T}$ to be the set of cells in the fixed marginals $\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}$. When the iterative procedure described below starts, $\mathbf{T}_{0}$ will contain only the cells in the fixed marginals, i.e., $\mathbf{T}_{0}=\mathbf{M}$. For the remaining cells in $\mathbf{T}$, we set $L(t)=0$ and $U(t)=\mathbf{n}_{\varnothing}$. Denote by $L_{0}(\mathbf{T})$ and $U_{0}(\mathbf{T})$ this initial set of upper and lower bounds induced by $\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}$.

As the algorithm progresses, the current bounds $L(\mathbf{T})$ and $U(\mathbf{T})$ are improved (lower bounds increase and upper bounds decrease), and more and more cells are added to $\mathbf{T}_{0}$. When the bounds associated with $t$ become equal, $t$ is added to $\mathbf{T}_{0}$, and is assigned value $V(t)=L(t)=V(t)$. We state the bounds problem in a new equivalent form: Find sharp integer bounds for the cells in $\mathbf{T}$ if the values of some cells $\mathbf{T}_{0} \subset \mathbf{T}$ are fixed.

Let $\mathcal{Q}=\mathcal{Q}(\mathbf{T})$ denote the triplets of cells $\mathcal{Q}(\mathbf{T})=\left\{\left(t_{1}, t_{2}, t_{3}\right) \in \mathbf{T} \times \mathbf{T} \times \mathbf{T}: t_{1} \oplus t_{3}=t_{2}\right\}$ that represent the cell dependencies to be satisfied. Let $S\left[L_{0}(\mathbf{T}), U_{0}(\mathbf{T})\right]$ be the set of integer tables consistent with $L_{0}(\mathbf{T})$ and $U_{0}(\mathbf{T})$. It is easy to see that

$$
\left\{V(\mathcal{N}): V(\mathbf{T}) \in S\left[L_{0}(\mathbf{T}), U_{0}(\mathbf{T})\right]\right\}=\mathbf{T}\left(\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}\right) .
$$

To improve the current bounds, we go sequentially through all dependencies in $\mathcal{Q}$ and update upper and lower bounds in the following way. Consider a triplet $\left(t_{1}, t_{2}, t_{3}\right) \in \mathcal{Q}$ with $t_{1} \prec t_{2}$ and $t_{3} \prec t_{2}$. If $t_{1}, t_{2}, t_{3} \in \mathbf{T}_{0}$, we check whether we came across an inconsistency. The procedure stops if $V\left(t_{1}\right)+V\left(t_{3}\right) \neq$ $V\left(t_{2}\right)$. Assume that $t_{1}, t_{3} \in \mathbf{T}_{0}$, and $t_{2} \notin \mathbf{T}_{0}$. Then $t_{2}$ can only take one value, namely $V\left(t_{1}\right)+V\left(t_{3}\right)$. If $V\left(t_{1}\right)+V\left(t_{3}\right) \notin\left[L\left(t_{2}\right), U\left(t_{2}\right)\right]$, we have encountered an inconsistency and exit the procedure. Otherwise we set $V\left(t_{2}\right)=L\left(t_{2}\right)=U\left(t_{2}\right)=V\left(t_{1}\right)+V\left(t_{3}\right)$, and include $t_{2}$ in the set $\mathbf{T}_{0}$ of cells having a fixed value. Similarly, if $t_{1}, t_{2} \in \mathbf{T}_{0}$ and $t_{3} \notin \mathbf{T}_{0}, t_{3}$ can only be equal to $V\left(t_{2}\right)-V\left(t_{1}\right)$. If $V\left(t_{2}\right)-V\left(t_{1}\right) \notin\left[L\left(t_{3}\right), U\left(t_{3}\right)\right]$, we again discovered an inconsistency. Otherwise, we set $V\left(t_{3}\right)=L\left(t_{3}\right)=U\left(t_{3}\right)=V\left(t_{2}\right)-V\left(t_{1}\right)$ and $\mathbf{T}_{0}=\mathbf{T}_{0} \cup\left\{t_{3}\right\}$. In the case that $t_{2}, t_{3} \in \mathbf{T}_{0}$ and $t_{1} \notin \mathbf{T}_{0}$ is handled analogously.

Now we examine the situation when at least two of the cells $t_{1}, t_{2}, t_{3}$ do not have a fixed value. For each of the three cells not having a fixed value, we update its upper and lower bounds so that the new bounds satisfy the dependency $t_{1} \oplus t_{3}=t_{2}$. Suppose that $t_{1} \notin \mathbf{T}_{0}$. If $U\left(t_{2}\right)-L\left(t_{3}\right)<L\left(t_{1}\right)$ or if $L\left(t_{2}\right)-U\left(t_{3}\right)>U\left(t_{1}\right)$, an inconsistency is detected and the procedure stops. Otherwise, the updated bounds for $t_{1}$ will be $U\left(t_{1}\right)=$ $\min \left\{U\left(t_{1}\right), U\left(t_{2}\right)-L\left(t_{3}\right)\right\}$ and $L\left(t_{1}\right)=\max \left\{L\left(t_{1}\right), L\left(t_{2}\right)-U\left(t_{3}\right)\right\}$. If $t_{3} \notin \mathbf{T}_{0}$, we update $L\left(t_{3}\right)$ and $U\left(t_{3}\right)$ in the same way. Finally, assume that $t_{2} \notin \mathbf{T}_{0}$. If $U\left(t_{1}\right)+U\left(t_{3}\right)<L\left(t_{2}\right)$ or if $L\left(t_{1}\right)+L\left(t_{3}\right)>U\left(t_{2}\right)$, we stop the algorithm. Otherwise, we set $U\left(t_{2}\right)=\min \left\{U\left(t_{2}\right), U\left(t_{1}\right)+U\left(t_{3}\right)\right\}$ and $L\left(t_{2}\right)=\max \left\{L\left(t_{2}\right), L\left(t_{1}\right)+L\left(t_{3}\right)\right\}$. After updating the bounds of $t \in \mathbf{T}$, we check whether the new upper bound is equal to the new lower bound. If so, we add $t$ to $\mathbf{T}_{0}$ and set $V(t)=L(t)=U(t)$.

We continue cycling through dependencies in $\mathcal{Q}$ until the upper bounds no longer decrease, the lower bounds no longer increase and no new cells are added to $\mathbf{T}_{0}$. The procedure terminates in a finite number of steps: either an inconsistency is detected or bounds cannot be improved.

If an inconsistency is detected, $S\left[L_{0}(\mathbf{T}), U_{0}(\mathbf{T})\right]$ is empty and no bounds are generated. However, $S\left[L_{0}(\mathbf{T}), U_{0}(\mathbf{T})\right]$ could still be empty even if the shuttle procedure did not come across any inconsistencies and has converged to bounds $L_{s}(\mathbf{T})$ and $U_{s}(\mathbf{T})$. These two bounds arrays define the same feasible set of
tables as the arrays $L_{0}(\mathbf{T})$ and $U_{0}(\mathbf{T})$ we started with, namely $S\left[L_{s}(\mathbf{T}), U_{s}(\mathbf{T})\right]=S\left[L_{0}(\mathbf{T}), U_{0}(\mathbf{T})\right]$. The fact that, for any $t \in \mathbf{T}$, we have $L_{0}(t) \leq L_{s}(t) \leq U_{s}(t) \leq U_{0}(t)$, could make one think that $S\left[L_{s}(\mathbf{T}), U_{s}(\mathbf{T})\right]$ might be strictly included in $S\left[L_{0}(\mathbf{T}), U_{0}(\mathbf{T})\right]$. Nevertheless, the dependencies in $\mathcal{Q}(\mathbf{T})$ imply the equality of the two sets of feasible integer tables [4].

### 4.1 Convergence Properties

The bounds produced by the generalized shuttle algorithm (provided it did not stop due to an inconsistency) are valid in the sense that the sharp bounds for each cell $t_{0} \in \mathbf{T}$ lie inside the interval defined by the shuttle bounds for $t_{0}$. However, there are two cases when the shuttle bounds are sharp: dichotomous $k$-dimensional tables with all $(k-1)$-dimensional marginals fixed, and when fixed the marginals are the minimal sufficient statistics of a decomposable log-linear model. In both instances, explicit formulas for the bounds exist, and employing the generalized shuttle algorithm is equivalent to using these formulas. Computing the bounds directly is more efficient, but it is insightful to examine how the algorithm works in these two cases.

### 4.2 Dichotomous $k$-way Tables with Fixed $(k-1)$-way Marginals

Consider a $k$-way table $\mathbf{n}=\{n(i)\}_{i \in \mathcal{I}}$ for which $\mathcal{I}_{1}=\cdots=\mathcal{I}_{k}=\{1,2\}$. Collapsing $\mathbf{n}$ across categories is equivalent to collapsing $\mathbf{n}$ across variables, thus the set $\mathbf{T}$ associated with the dichotomous table $\mathbf{n}$ is the set of cells in every marginal of $\mathbf{n}$. Assume that $(k-1)$-dimensional marginals of $\mathbf{n}$ are fixed. This implies that every lower-dimensional marginal of $\mathbf{n}$ will also be known. The only cells in $\mathbf{T}$ that are unknown are those in the original table.

The ( $k-1$ )-dimensional marginals of $\mathbf{n}$ are the minimal sufficient statistics of the log-linear model of no $(k-1)$-order interaction. This log-linear model has only one degree of freedom because $\mathbf{n}$ is dichotomous [7]. Consequently, we can uniquely express the count in any cell as a function of one single fixed cell alone-only one more quantity is needed in order to determine the entries for the full table.

Let $n^{*}$ be the unknown count in the $(1, \ldots, 1)$ cell. In Proposition 1 [4] we give an explicit formula for computing the count in an arbitrary cell based on $n^{*}$ and on the set of fixed marginals.

Proposition 1. Consider an index $i^{0} \in \mathcal{I}$. Let $\left\{q_{1}, \ldots, q_{l}\right\} \subset K$ be such that, for $r \in K$,

$$
i_{r}^{0}=\left\{\begin{array}{lll}
1, & \text { if } & r \in K \backslash\left\{q_{1}, \ldots, q_{l}\right\},  \tag{1}\\
2, & \text { if } & r \in\left\{q_{1}, \ldots, q_{l}\right\} .
\end{array}\right.
$$

For $s=1, \ldots, l$, let $C_{s}=K \backslash\left\{q_{s}\right\}$. Then

$$
\begin{equation*}
n\left(i^{0}\right)=(-1)^{l} \cdot n^{*}-\sum_{s=0}^{l-1}(-1)^{l+s} \cdot n_{C_{(l-s)}}\left(1, \ldots, 1, i_{q_{(l-s)}+1}^{0}, \ldots, i_{k}^{0}\right) . \tag{2}
\end{equation*}
$$

The upper and lower bounds can therefore be obtained by imposing the non-negativity constraints $n\left(i^{0}\right) \geq 0, i^{0} \in \mathcal{I}$, in these relations. For example, the sharp lower bound for the cell $(1, \ldots, 1)$ is

$$
\begin{equation*}
\max \left\{\sum_{s=0}^{l-1}(-1)^{s} \cdot n_{C_{(l-s)}}\left(1, \ldots, 1, i_{q_{(l-s)}+1}^{0}, \ldots, i_{k}^{0}\right): l \text { even }\right\}, \tag{3}
\end{equation*}
$$

where $i^{0}$ is as in (1), and the corresponding upper bound is given by

$$
\begin{equation*}
\min \left\{\sum_{s=0}^{l-1}(-1)^{s} \cdot n_{C_{(l-s)}}\left(1, \ldots, 1, i_{q(l-s)+1}^{0}, \ldots, i_{k}^{0}\right): l \text { odd }\right\} . \tag{4}
\end{equation*}
$$

The generalized shuttle algorithm converges to the bounds in (3) and (4) [4]. Moreover, one can obtain all feasible tables consistent with the $(k-1)$-dimensional marginals of $\mathbf{n}$ by replacing every possible value $n^{*}$ that the cell $(1,1, \ldots, 1)$ can take in (2) applied for all cells $i^{0}$. In particular, all the cells in $\mathbf{n}$ can take the same number of values-the difference between the upper and lower bounds is constant for all cells, and is equal with the number of feasible integer tables consistent with the ( $k-1$ )-dimensional marginals of $\mathbf{n}$.

### 4.3 The Decomposable Case

Log-linear models are a common way of representing and studying contingency tables with fixed marginals. In particular, a graphical log-linear model corresponds to conditional independence relationships that can be summarized by means of an independence graph [12]. Decomposable log-linear models [11] are a sub-class of graphical models with closed form structure and special properties that lead to explicit formulas for computing bounds [5]. The complete proof of the next theorem appears in [4].

Theorem 1. Let $\mathbf{n}^{1}=\mathbf{n}_{C_{1}}^{1}, \ldots, \mathbf{n}^{p}=\mathbf{n}_{C_{p}}^{1}$ be a set of tables that are consistent with each other. Assume that $C_{1}, \ldots, C_{p}$ are the minimal sufficient statistics of a decomposable log-linear model. Let $S_{2}, \ldots, S_{p}$ be the set of separators associated with $C_{1}, \ldots, C_{p}$ in the corresponding independence graph. Then:
(i) $\mathbf{T}\left(\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}\right)$ contains at least one table.
(ii) The sharp upper bound for $i \in \mathcal{I}$ given $\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}$ is

$$
U\left(i ; \mathbf{T}\left(\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}\right)\right)=\min \left\{n^{j}\left(i_{C_{j}}\right): j=1, \ldots, p\right\}
$$

(iii) The sharp lower bound for $i \in \mathcal{I}$ given $\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}$ is

$$
L\left(i ; \mathbf{T}\left(\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}\right)\right)=\max \left\{\sum_{j=1}^{p} n^{j}\left(i_{C_{j}}\right)-\sum_{j=2}^{p} n_{S_{j}}^{j}\left(i_{S_{j}}\right), 0\right\} .
$$

Therefore, when the fixed marginals define a decomposable graphical model, pairwise consistency of marginals implies the existence of a feasible table. In this case, the generalized shuttle algorithm converges to the bounds in Theorem 1 [4].

### 4.4 Finding a Feasible Integer Table

As noted above, the generalized shuttle algorithm can converge to bounds $L_{s}(\mathbf{T})$ and $U_{s}(\mathbf{T})$ even if there does not exist an integer table consistent with the fixed marginals $\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}$. Since bounds make no sense when no table exists, we augment our procedure by proposing a method that will actually determine a feasible table in $\mathbf{T}\left(\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}\right)$, provided one exists.

This is done by sequentially choosing possible values for cells in $\mathcal{N}$. Once a new cell has been fixed, bounds for all the cells in $\mathbf{T}$ are updated using the shuttle algorithm. If the shuttle procedure did not stop because of an inconsistency, we pick a value for another cell. Otherwise, we choose a new value for the cell fixed at the current step. If all the cells in $\mathcal{N}$ have been fixed and the shuttle procedure completed successfully, which indicates that the dependencies in $\mathcal{Q}(\mathbf{T})$ are satisfied, we have determined a table in $\mathbf{T}\left(\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}\right)$. We denote by $T_{0}^{(0)}=T_{0}$ the set of cells fixed by the shuttle procedure when computing the bounds $L_{s}(\mathbf{T})$ and $U_{s}(\mathbf{T})$. Also, initialize the bounds arrays $L^{(0)}(\mathbf{T})=L_{s}(\mathbf{T})$ and $U^{(0)}(\mathbf{T})=U_{s}(\mathbf{T})$. Here is the procedure in pseudo-code:

Step 1. Set $l=1$.
Step 2. Check whether $\mathcal{N}^{l}=\mathcal{N} \cup\left(\mathbf{T} \backslash \mathbf{T}_{0}^{(l-1)}\right)$ is empty. If so, a feasible table has been determined; stop the algorithm.

Step 3. Select a cell $t^{l} \in \mathcal{N}^{l}$.
Step 4. FOR every integer $v_{l} \in\left[L^{(l-1)}\left(t_{l}\right), U^{(l-1)}\left(t_{l}\right)\right]$ DO

- Initialize new bound arrays $L^{(l)}(\mathbf{T})=L^{(l-1)}(\mathbf{T})$ and $U^{(l)}(\mathbf{T})=U^{(l-1)}(\mathbf{T})$.
- Set $V\left(t_{l}\right)=L^{(l)}\left(t_{l}\right)=U^{(l)}\left(t_{l}\right)=v_{l}$ and put $\mathbf{T}^{(l)}=\mathbf{T}^{(l-1)} \cup\left\{t_{l}\right\}$.
- Run the generalized shuttle algorithm to update $L^{(l)}(\mathbf{T}), U^{(l)}(\mathbf{T})$ and $\mathbf{T}^{(l)}$.
- If the generalized shuttle algorithm did not stop because of an inconsistency, set $l=l+1$ and go to Step 2.


## END FOR.

Step 5. The algorithm stops; there does not exist a feasible integer table in $\mathbf{T}\left(\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}\right)$.
Instead of "blindly" searching the entire space of all possible combinations of values for the cells in $\mathcal{N}$ defined by the bounds arrays $L_{s}(\mathcal{N})$ and $U_{s}(\mathcal{N})$, the procedure reduces the search space continuously as the algorithm progresses because the bounds are updated at each step. Several heuristics exist for substantially increasing the speed of this search procedure [4]. Moreover, if we do not stop the algorithm at Step 2 once the first feasible table is generated, then all the integer tables in $\mathbf{T}\left(\mathbf{n}^{1}, \ldots, \mathbf{n}^{p}\right)$ will be identified.

### 4.5 Calculating Sharp Bounds

We can now obtain sharp bounds for any cell $t_{0} \in \mathcal{T}$. Let $L_{s}(\mathbf{T})$ and $U_{s}(\mathbf{T})$ be the bounds produced by the generalized shuttle algorithm. Then

$$
L_{s}\left(t_{0}\right) \leq L_{t_{0}} \leq U_{t_{0}} \leq U_{s}\left(t_{0}\right)
$$

We need to "adjust" $L_{s}\left(t_{0}\right)$ and $U_{s}\left(t_{0}\right)$ to the sharp bounds $L_{t_{0}}$ and $U_{t_{0}}$ by making use of the simple fact that, for any integer $v \in\left[L_{s}\left(t_{0}\right), L_{t_{0}}\right) \cup\left(U_{t_{0}}, U_{s}\left(t_{0}\right)\right]$, there does not exist an integer array $V(\mathbf{T}) \in S\left[L_{s}(\mathbf{T}), U_{s}(\mathbf{T})\right]$ such that $V\left(t_{0}\right)=v$. To determine $L_{t_{0}}$, we start with $v=L_{s}\left(t_{0}\right)$ and attempt to determine an array $V(\mathbf{T}) \in S\left[L_{s}(\mathbf{T}), U_{s}(\mathbf{T})\right]$ with $V\left(t_{0}\right)=v$. If such an array exists, $L_{t_{0}}=v$. Otherwise, we do the same thing for $v=L_{s}\left(t_{0}\right)+1$, and so on. Here is the complete algorithm in pseudo-code. (Again, $\mathbf{T}_{0}$ represents the cells $t$ with $L_{s}(t)=U_{s}(t)$.)

FOR every integer value $v_{l} \in\left[L_{s}\left(t_{0}\right), L_{s}\left(t_{0}\right)+1, \ldots, U_{s}\left(t_{0}\right)\right]$ DO

- Initialize new bound arrays $L^{v}(\mathbf{T})=L_{s}(\mathbf{T})$ and $U^{v}(\mathbf{T})=U_{s}(\mathbf{T})$.
- Set $V\left(t_{0}\right)=L^{v}\left(t_{0}\right)=U^{v}\left(t_{0}\right)=v$ and set $\mathbf{T}^{v}=\mathbf{T}_{0} \cup\left\{t_{0}\right\}$.
- Run the generalized shuttle algorithm to update $L^{v}\left(t_{0}\right), U^{v}\left(t_{0}\right)$ and $\mathbf{T}^{v}$.
- Using the procedure described in the previous section, find out whether $S\left[L^{v}\left(t_{0}\right), U^{v}\left(t_{0}\right)\right]$ is empty.
- If $S\left[L^{v}\left(t_{0}\right), U^{v}\left(t_{0}\right)\right]$ is not empty, set $L_{t_{0}}=v$ and stop the algorithm.


## END FOR

Determination of $U_{t_{0}}$ is similar: start with $v=U_{s}\left(t_{0}\right)$ and attempt to determine an array $V(\mathbf{T}) \in$ $S\left[L_{s}(\mathbf{T}), U_{s}(\mathbf{T})\right]$ with $V\left(t_{0}\right)=v$. If such an array exists, $U_{t_{0}}$ is equal to $v$; otherwise, we set $v=U_{s}\left(t_{0}\right)-1$ and continue. In this way, we obtain sharp bounds not only for the cells in $\mathcal{N}$ but also for any cell in a cross-classification obtained by collapsing $\mathbf{n}$ across categories.
Example 1. The left-hand panel of Table 1 contains a $2 \times 2 \times 2 \times 2$ contingency table. The generalized shuttle algorithm was used to compute the upper and lower bounds induced by fixing the 6 two-way marginals of this table. There is only one table consistent with this set of two-way marginals, so for all the cells in this table, the sharp integer upper and lower bounds are equal to the actual cell entries.

The right panel in Table 1 contains comparable bounds compute using linear programming (specifically, the simplex algorithm). For all the cells, one of the bounds is different from the corresponding sharp integer bound. In some cases, the simplex algorithm found fractional bounds. In other cases, although simplex converged to integer bounds, these bounds are not sharp. Especially, cell (1, 1,2, 1), marked with a box in Table 1, contains a count of 0 , while the actual upper bound is 1.67. Therefore the distance between the integer and the real bounds can be strictly greater than 1.

Adjusting the bounds can be done in parallel for all the cells in the table [4]. Other methods for reducing the computational effort required by the generalized shuttle algorithm in the case of large sparse multi-way tables, as well as an example of calculating bounds in parallel for a $2^{16}$, table are presented in [4].

|  |  | C |  |  | no |  | yes |  | C | no |  | yes |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| A | B | D | no | yes | no | yes | D | no | yes | no | yes |  |  |  |
| no | no |  | 1 | 0 | 0 | 1 |  | $[0,1]$ | $[0,0.67]$ | $[0,1.67]$ | $[0,1]$ |  |  |  |
|  | yes |  | 0 | 0 | 1 | 0 |  | $[0,0.67]$ | $[0,0.67]$ | $[0,1]$ | $[0,0.67]$ |  |  |  |
| yes | no | 0 | 0 | 1 | 0 |  | $[0,0.67]$ | $[0,0.67]$ | $[0,1]$ | $[0,0.67]$ |  |  |  |  |
|  | yes |  | 0 | 1 | 0 | 0 |  | $[0,0.67]$ | $[0,1]$ | $[0,0.67]$ | $[0,0.67]$ |  |  |  |

Table 1: A $2 \times 2 \times 2 \times 2$ table (left-hand panel) and the real bounds computed using linear programming (right-hand panel).

## 5 Conclusions

In this paper we have outlined how to deal with the challenging issues associated with storing and manipulating large, sparse tables in the context of statistical disclosure limitation. Dissemination systems and software implementing these techniques are described in $[6,9]$. To the extent that they exploit sparsity, these techniques exhibit strong scalability properties. For instance, the fundamental data structures and methods for calculating marginal sub-tables seem able to handle almost any table that is likely to arise in practice. By contrast, the generalized shuttle algorithm of Section 4, in its current form, entails computations for every cell in the underlying table, and consequently does not scale well. Deriving a scalable version is just one of many research challenges that remain.

## Acknowledgments

Support for the research was provided by National Science Foundation grant EIA-9876619 to NISS. The authors thank Stephen E. Fienberg and Stephen F. Roehrig for suggestions, comments and useful discussions. The example in Table 1 was suggested by Bernd Sturmfels.

## References

[1] Y. M. M. Bishop, S. E. Fienberg, and P. W. Holland. Discrete Multivariate Analyses: Theory and Practice. MIT Press, Cambridge, MA, 1975.
[2] L. Buzzigoli and A. Giusti. An algorithm to calculate the lower and upper bounds of the elements of an array given its marginals. In Statistical Data Protection (SDP'98) Proceedings, pages 131-147, Luxembourg, 1999. Eurostat.
[3] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. Introduction to Algorithms. MIT Press/McGraw-Hill, 1990.
[4] A. Dobra. Statistical Tools for Disclosure Limitation in Multi-way Contingency Tables. PhD thesis, Department of Statistics, Carnegie Mellon University, 2002.
[5] A. Dobra and S. E. Fienberg. Bounds for cell entries in contingency tables given marginal totals and decomposable graphs. Proc. Nat. Acad. Sciences, 97:11885-11892, 2000.
[6] A. Dobra, A. F. Karr, S. E. Fienberg, and A. P. Sanil. Software systems for tabular data releases. Int. J. Uncertainty, Fuzziness and Knowledge Based Systems, 2002.
[7] S. E. Fienberg. Fréchet and bonferroni bounds for multi-way tables of counts with applications to disclosure limitation. In Statistical Data Protection (SDP'98) Proceedings, pages 115-129, Luxembourg, 1999. Eurostat.
[8] V. Harinarayan, A. Rajaraman, and J. D. Ullman. Implementing data cubes efficiently. In Proceedings of theACM SIGMOD International Conference on Management of Data, volume 25, 2 of ACM SIGMOD Record, pages 205-216, New York, June 4-6 1996. ACM Press.
[9] A. F. Karr, A. Dobra, and A. P. Sanil. Table servers: Protecting confidentiality in tabular data releases. Comm. ACM, 2003. To appear.
[10] D. E. Knuth. Fundamental Algorithms, volume 1 of The Art of Computer Programming. AddisonWesley, Reading, Massachusetts, third edition, 1997.
[11] S. L. Lauritzen. Graphical Models. Clarendon Press, Oxford, 1996.
[12] D. Madigan and J. York. Bayesian graphical models for discrete data. Internat. Statis. Rev., 63:215232, 1995.
[13] Andrew W. Moore and Mary S. Lee. Cached sufficient statistics for efficient machine learning with large datasets. J. Artificial Intell. Res., 8:67-91, 1998.
[14] Bjarne Stroustrup. The C++ Programming Language (3rd Edition). Addison-Wesley, Reading, MA, 1997.
[15] Sun Microsystems. Java programming language. http://java.sun.com, 2002.
[16] L. C. R. J. Willenborg and T. de Waal. Elements of Statistical Disclosure Control. Springer-Verlag, New York, 2001.


[^0]:    *Now at Duke University, Durham, NC 27708-0251.

