On the Structure of Armstrong Relations for Functional Dependencies

CATRIEL BEERI

The Hebrew University, Jerusalem, Israel

MARTIN DOWD

Rutgers University, New Brunswick, New Jersey

RONALD FAGIN

IBM Research Laboratory, San Jose, California

AND

RICHARD STATMAN

Rutgers University, New Brunswick, New Jersey

Abstract. An Armstrong relation for a set of functional dependencies (FDs) is a relation that satisfies each FD implied by the set but no FD that is not implied by it. The structure and size (number of tuples) of Armstrong relations are investigated. Upper and lower bounds on the size of minimal-sized Armstrong relations are derived, and upper and lower bounds on the number of distinct entries that must appear in an Armstrong relation are given. It is shown that the time complexity of finding an Armstrong relation, given a set of functional dependencies, is precisely exponential in the number of attributes. Also shown is the falsity of a natural conjecture which says that almost all relations obeying a given set of FDs are Armstrong relations for that set of FDs. Finally, Armstrong relations are used to generalize a result, obtained by Demetrovics using quite complicated methods, about the possible sets of keys for a relation.

Categories and Subject Descriptors: F.4.1 [**Mathematical Logic and Formal Languages**]: Mathematical Logic—*model theory*; H.2.1 [**Database Management**]: Logic Design—*schema and subschema*

General Terms. Design, Languages, Management, Theory

Additional Key Words and Phrases: Armstrong relation, functional dependency, relational database, computational complexity, random relation, key, logical consequence

1. *Introduction*

Armstrong relations are objects of interest in relational database theory. Let Σ be a set of functional dependencies (FDs) [6], and let σ be a single FD. When we say that Σ *logically implies* σ or that σ is a *logical consequence* of Σ, we mean that whenever every FD in Σ holds for a relation R, then also σ holds for R. That is,

The research of the first author was supported in part by Grant 1849/79 of the USA-Israel Binational Science Foundation. The research of the fourth author was supported by the National Science Foundation under Grant MCS 79-23199.

Authors' addresses: C. Beeri, Department of Computer Science, Hebrew University, Givat Ram, 91904 Jerusalem, Israel; M. Dowd and R. Statman, Department of Mathematics, Rutgers University, New Brunswick, N.J. 08903; R. Fagin, IBM Research Laboratory, K55/281, 5600 Cottle Road, San Jose, Calif. 95193.

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.

© 1984 ACM 0004-5411/84/0100-0030 $00.75

EMP	DEPT	MGR
Hilbert	Math	Gauss
Pythagoras	Math	Gauss
Turing	Computer Science	von Neumann
Einstein	Physics	Gauss

FIGURE 1

there is no "counterexample relation" or "witness" R such that every FD in Σ holds for R, but such that σ fails in R. We write $\Sigma \models \sigma$ to mean that Σ logically implies σ (and $\Sigma \not\models \sigma$ to mean that Σ does not logically imply σ). For example, $\{A \rightarrow B, B \rightarrow C\} \models A \rightarrow C$. Let Σ be a set of FDs, and let Σ^* be the set of all FDs that are logical consequences of Σ. For each FD σ not in Σ^*, we know (by definition of $\models$) that there is a relation R_σ (a witness) such that R_σ obeys Σ but not σ. It follows from results of Armstrong [1] that there is a relation (a global witness) that can simultaneously serve the role of all of the R_σ's. That is, Armstrong showed that there is a relation that obeys Σ^* and *no other FDs*. Following Fagin [14], we call such a relation an *Armstrong relation* for Σ. Actually, Armstrong did not *explicitly* state or prove the existence of an Armstrong relation. Instead, he proved a result that implies both the completeness of a certain set of axioms about FDs (see [12]) and the existence of an Armstrong relation.

As an example [14], let Σ be the set {EMP → DEPT, DEPT → MGR}, containing two FDs. Then Σ^* contains the FDs in Σ, along with, for example, the FD EMP → MGR. It is easy to verify (by considering all possible FDs involving only EMP, DEPT, and MGR) that the relation (call it R) in Figure 1 is an Armstrong relation for Σ, that is, that it obeys every FD in Σ^* and no others. For example, the FD MGR → DEPT is not an FD in Σ^*, and indeed, R does *not* obey this FD, since Gauss is the manager of two distinct departments (Math and Physics).

We shall give a simple proof of the existence of Armstrong relations for sets of FDs. That proof is a slight modification of a proof by Beeri et al. [4], which contains a minor "bug" (in that dependencies whose left-hand side is the empty set are ignored).

The existence of Armstrong relations has been proved in the presence, not just of FDs, but of the much more general "embedded implicational dependencies" (for details, see [14]). A concept closely related to Armstrong relations in traditional mathematics is the free algebra with countably many generators [19], which obeys just a specified set of equations and their logical consequences, and no other equations. (However, although the free algebra just mentioned is unique to within iomorphism, Armstrong relations are not [14].) In ordinary first-order logic (where arbitrary first-order sentences, and not just, say, FDs are allowed), there can be no Armstrong relations. For example, let Σ be the empty set $\varnothing$. Assume that R were a relation that obeyed just Σ^* (that is, just the tautologies), and no other first-order sentences. Let σ be an arbitrary first-order sentence such that neither σ nor $\neg\sigma$ is a tautology. Clearly, R must obey one of σ or $\neg\sigma$; thus, R obeys a nontautology. This is a contradiction. Hence there is a witness for σ (a relation that shows that σ is not a tautology), and a witness for $\neg\sigma$ (a relation that shows that $\neg\sigma$ is not a tautology), but there is no global witness (a relation that simultaneously shows that σ is not a tautology and $\neg\sigma$ is not a tautology).

It is common to speak of a relation obeying an "accidental" dependency, that is, a dependency that is not a logical consequence of the collection of "specified" dependencies. Thus, each specified dependency is supposed to hold "for all time," that is, for every "snapshot" (instance) of the database, whereas an accidental dependency is one that happens to hold in some snapshot of the database, but may fail in other snapshots. An Armstrong relation is precisely one that obeys every specified dependency and no accidental dependency.

We note an interesting "practical" application for Armstrong relations. Silva and Melkanoff [22] have developed a database design aid, in which the database designer inputs a set of FDs and MVDs (multivalued dependencies) [11, 26]. The design aid then presents him with an Armstrong relation, that is, a "sample relation" that obeys just those dependencies that are logical consequences of those that he has inputted. (Armstrong relations exit in the presence of FDs and MVDs, and this is the case in which Silva and Melkanoff were interested.) Let us say, for example, that the designer gives as input the set {EMP → DEPT, DEPT → MGR} of FDs. The database design aid would then present the designer with an Armstrong relation, such as relation R in Figure 1, for this set of dependencies. The designer would then inspect the sample relation, and might observe, for example, "Here is a manager, namely Gauss, who manages two different departments. *Therefore*, the dependencies that I inputted must not have implied that no manager can manage two different departments. Since I want this to be a constraint for my database, I'd better input the FD MGR → DEPT."

In this example, the designer did not have to explicitly think about the dependency MGR → DEPT and whether or not it was a consequence of the dependencies that he input; rather, by seeing the Armstrong relation, and thinking about what it said, he simply *noticed* that the FD MGR → DEPT failed. Thus, Silva and Melkanoff's approach is a partial solution, in the spirit of Query-by-Example [27], to the problem of helping a designer think of what dependencies should be included. Unfortunately, one of the results in this paper shows a limitation of this approach: namely, a minimal-sized Armstrong relation for a set of FDs can be of exponential size (in the number of attributes).

Let us call the collection of all relations (with a given set U of attributes) that obey a given set of FDs an *FD class* [14]. Let R be a fixed relation. In the spirit of Ginsburg and Zaiddan [17], we define *the FD class generated by* R to be the smallest FD class that contains R. It is easy to see that this class is simply those relations (with attributes the same as those of R) that obey Σ, where Σ is the set of all FDs obeyed by R. A natural question is whether every FD class has a generator. The answer [17] is yes: if the FD class $\mathscr{R}$ consists of all relations with attributes U that obey Σ, then let R be an Armstrong relation (with attributes U) for Σ; it is easy to see that R is generator for the class $\mathscr{R}$. Thus, a natural interpretation for Armstrong relations is as class generators.

Armstrong relations arise naturally in proofs in database theory. For example, Casanova et al. [5] make use of an "Armstrong database" to help show that for each k, there is no k-ary complete axiomatization for functional dependencies and "inclusion dependencies" together (a typical inclusion dependency [5] says, for example, that every manager is an employee). We give another application of Armstrong relations in Section 4.

For a history and summary of results on Armstrong relations, see [13].

In Section 2 we present basic definitions. In Section 3 we present a simple proof, which "patches" a proof by Beeri et al. [4] of the existence of Armstrong relations for sets of FDs. In Section 4 we give an application of Armstrong relations to generalize a result of Demetrovics about the possible sets of keys for a relation. In Section 5 we show that a natural conjecture is false. The conjecture says that almost all relations obeying a given set of functional dependencies are Armstrong relations for that set. In Section 6 we characterize Armstrong relations in terms of closed sets. We use this technique to obtain upper and lower bounds on the size of minimal Armstrong relations. We also give upper and lower bounds on the number of distinct entries that must appear in an Armstrong relation. In Section 7 we present some complexity results. We show that the time complexity of producing an Armstrong relation, given a set of functional dependencies, is precisely exponential in the number of attributes. By this we mean that there is an exponential-time algorithm, and furthermore, that there is an example in which the time simply to write down the Armstrong relation is exponential. Finally, we show that the problem of deciding whether there is a key of size at most given integer is NP-complete, whether the set of functional dependencies is presented explicitly, or implicitly via an Armstrong relation (Lucchesi and Osborn [20] already proved this result in the explicit case.)

2. *Basic Definitions*

We assume a finite set U of *attributes*. A *tuple* (over U) is a mapping with domain U, and a *relation* (over U) is a set of tuples (over U). If $X \subseteq U$, and if t is a tuple over U, then we denote the restriction of t to X by $t[X]$. If R is a relation over U, then $R[X] = \{t[X] : t \in R\}$. If A is an attribute of U, and if t is a tuple over U, then we may refer to $t[A]$ as an *entry*, in the A *column*.

A *functional dependency* (*over* U) [6], or an *FD*, is a statement, or sentence, $X \rightarrow Y$ where $X, Y \subseteq U$. A relation R over U *obeys* the FD $X \rightarrow Y$ if whenever t_1, t_2 are tuples of R with $t_1[X] = t_2[X]$, then $t_1[Y] = t_2[Y]$. We also say then that the FD *holds for* R. If the FD does not hold for R, then we say that the FD *fails* in R, or that R *violates* the FD.

3. *Constructing Armstrong Relations for FDs*

In this section we give a simple proof, which "patches" a proof by Beeri et al. [4], of the existence of Armstrong relations for sets of FDs. We begin with a simple lemma, which is not really necessary, but is convenient.

LEMMA 3.1 [11]. *Let* Σ *be a set of FDs, and let* σ *be a single FD such that* $\Sigma \not\models \sigma$. *Then there is a two-tuple relation that obeys* Σ *but not* σ.

PROOF. Since $\Sigma \not\models \sigma$, there is a relation R that obeys Σ but not σ. Let σ be the FD $X \rightarrow Y$. Since R does not obey $X \rightarrow Y$, there are tuples t_1 and t_2 of R such that $t_1[X] = t_2[X]$ but $t_1[Y] \neq t_2[Y]$. Let R' be a relation that contains precisely the tuples t_1 and t_2. It is easy to see that R' is the desired two-tuple relation. □

Although the proof of Lemma 3.1 is nonconstructive, we now sketch a simple constructive proof of the same result (this is of interest because we shall later use Lemma 3.1 to construct Armstrong relations). Let Σ be a fixed set of FDs, and let X be a set of attributes. We define X^* to be the set of all attributes A such that $\Sigma \models X \rightarrow A$. Let σ be an FD not in Σ^*. Now assume that $\Sigma \not\models \sigma$, and that σ is the FD $X \rightarrow Y$. Let R be a two-tuple relation, in which one tuple has all 0's as entries, and the other tuple has 0's in the X^* entries and 1 elsewhere. It is easy to verify [12] that R is a two-tuple relation that obeys Σ but not σ. A similar construction can be used [25] to prove the completeness of Armstrong's axioms for FDs.

We note that the claim in Lemma 3.1 is true about MVDs also. In fact, Sagiv et al. [21] show that if Σ is a set of dependencies (FDs or MVDs), if σ is a single dependency, and if R is a relation that obeys Σ but not σ, then R has a two-tuple subrelation that obeys Σ but not σ (a *subrelation* is a subset of the tuples). The proof is much harder than the proof of Lemma 3.1.

We now describe how to construct Armstrong relations for sets of FDs. We use the "disjoint union" technique of Beeri et al. [4]. The reader should note that there is a minor bug in their proof, caused by neglecting FDs with the empty set as their left-hand side. This bug is corrected below. See also Armstrong and Delobel [2].

THEOREM 3.2 [1, 4]. *For each set Σ of FDs, there is an Armstrong relation for Σ.*

PROOF. Let Σ be a set of FDs. We shall construct an Armstrong relation for Σ, that is, a relation that obeys the FDs in Σ^* and no other FDs. For each FD σ not in Σ^*, let R_σ be a relation that obeys Σ but not σ (by Lemma 3.1, we can even take R_σ to be a two-tuple relation). We can assume that the entries that appear in R_σ are distinct from those that appear in R_τ whenever σ and τ are distinct FDs, with the following exception. For each attribute A in $\varnothing^*$ (that is, for each attribute A such that $\Sigma \models \varnothing \rightarrow A$), every A entry in every tuple in every one of these relations is the same value. Let R be the union of all these relations R_σ, that is, $R = \bigcup \{R_\sigma : \sigma \notin \Sigma^*\}$. We now prove that R is an Armstrong relation for Σ.

We first show that R obeys every FD in Σ. Let $X \rightarrow Y$ be an FD in Σ, and let A be an attribute in Y. We need only show that R obeys the FD $X \rightarrow A$. There are two cases.

Case 1. $X \subseteq \varnothing^*$. Then $A \in \varnothing^*$, and so all of the tuples of R contain the same A entry. Hence, R obeys $X \rightarrow A$.

Case 2. $X \nsubseteq \varnothing^*$. Assume that t_1 and t_2 are tuples of R, and that $t_1[X] = t_2[X]$; we must show that $t_1[A] = t_2[A]$. Since $X \nsubseteq \varnothing^*$, there is an attribute B in $X - \varnothing^*$. Since $t_1[B] = t_2[B]$, and since $B \notin \varnothing^*$, we know by construction of R that t_1 and t_2 are tuples in the same relation R_σ, for some $\sigma \notin \Sigma^*$. Since R_σ obeys Σ, it follows that R_σ obeys the FD $X \rightarrow A$. Since also $t_1[X] = t_2[X]$, it follows that $t_1[A] = t_2[A]$. This was to be shown.

Thus, R obeys each FD in Σ (and hence in Σ^*). If σ is an FD not in Σ^*, then R violates σ, since its subrelation R_σ violates σ. Hence, R is an Armstrong relation for Σ. □

4. *An Application*

We now give an application of Armstrong relations to database theory. A *key* of a relation is a set K of attributes such that $K \rightarrow U$ holds in the relation but such that for every proper subset K' of K, the FD $K' \rightarrow U$ does not hold in it. A key gives a minimal unique identifier for each tuple in a relation.

If J is the set of keys of a relation, then clearly every pair of keys in J is incomparable under set inclusion. We now show that the converse holds.

THEOREM 4.1. *Let J be a nonempty collection of incomparable subsets of a finite set U. Then there is a relation with attributes U for which the set of keys is precisely J.*

PROOF. Let Σ be the set $\{K \to U : K \in J\}$ of FDs, and let R be an Armstrong relation for Σ. We claim that the set of keys of R is precisely J. To show this, it is sufficient to show that if $K' \subseteq U$, then

$$R \text{ obeys the FD } K' \to U \text{ if and only if } K' \text{ is a superset of a member of } J. \quad (4.1)$$

Since R is an Armstrong relation for Σ, we know that the statement (4.1) is equivalent to the statement,

$$\Sigma \models K' \to U \text{ if and only if } K' \text{ is a superset of a member of } J. \quad (4.2)$$

We now show (4.2). Assume first that K' is a superset of a member K of J. Then the FD $K \to U$ is in Σ, and so $\Sigma \models K' \to U$. Conversely, assume that $\Sigma \models K' \to U$, but that K' is not a superset of a member of J; we shall derive a contradiction. Let R be a two-tuple relation such that one tuple has all 0's as entries, and the other tuple has 0's in the K' entries and 1 elsewhere. Since K' is not a superset of a member of J, it follows easily that R obeys Σ. However, R does not obey the FD $K' \to U$. This contradicts our assumption that $\Sigma \models K' \to U$. □

It is well-known [23] that the biggest set J of incomparable subsets of an n-element set U has $S(n)$ members, where $S(n)$ is the binomial coefficient $\binom{n}{\lfloor n/2 \rfloor}$, and where $\lfloor x \rfloor$ is the greatest integer not exceeding x. This set J is the set of all subsets of U that contain precisely $\lfloor n/2 \rfloor$ members. Demetrovics [7] proved, by a complicated construction, that there is a relation, with n attributes, that has $S(n)$ keys. Of course, this result is an immediate consequence of our Theorem 4.1 and the above remarks. Demetrovics' result is of interest because it shows that the maximum possible number of keys in a relation with n attributes can be obtained.

5. *Random Relations versus Armstrong Relations*

Let us hold fixed a set U of attributes. Let $\mathscr{A}_k$ be the set of all relations with attributes U such that every entry of the relation is a member of $\{1, \ldots, k\}$. Thus, $\mathscr{A}_k$ contains 2^{k^u} members, where u is the number of attributes (that is, the size of U). If $\mathscr{P}$ is a property of relations, then we say that "almost all relations have property $\mathscr{P}$" (or "a random relation has property $\mathscr{P}$") if the fraction of members of $\mathscr{A}_k$ with property $\mathscr{P}$ converges to 1 as $k \to \infty$. Fagin [10] showed that if $\mathscr{P}$ is a first-order property of relations, then either almost all relations have property $\mathscr{P}$ or almost all relations violate property $\mathscr{P}$. From his characterization, it follows easily that for each nontrivial FD σ, almost all relations (over the appropriate attributes) violate σ. Since there are only a finite number of FDs over a given set of attributes, it follows that almost all relations simultaneously violate every nontrivial FD. Thus, almost all relations are Armstrong relations (with respect to FDs) for the empty set of FDs.

If $\mathscr{P}$ and $\mathscr{Q}$ are properties of relations, then we say that "almost all relations with property $\mathscr{Q}$ have property $\mathscr{P}$" if the number of members of $\mathscr{A}_k$ with both properties $\mathscr{Q}$ and $\mathscr{P}$ divided by the number with property $\mathscr{Q}$ converges to 1 as $k \to \infty$. A natural conjecture is that almost all relations that obey a given set Σ of FDs is an Armstrong relation for Σ (with respect to FDs). As we noted earlier, the conjecture is true when Σ is empty. We now show that the conjecture is false in general. In fact, we shall show that if the attributes U are $\{A, B, C, D\}$, then almost all relations obeying the FD $A \to BCD$ also obey the FD $BCD \to A$, and so are certainly not Armstrong relations for $\{A \to BCD\}$.

Let a_k be the number of 4-ary relations over $\{1, \ldots, k\}$ obeying $A \rightarrow BCD$, and let b_k be the number of 4-ary relations over $\{1, \ldots, k\}$ obeying $\{A \rightarrow BCD, BCD \rightarrow A\}$. We shall show that $b_k/a_k \rightarrow 1$ as $k \rightarrow \infty$.

Let a_{kr} be the number of 4-ary r-tuple relations over $\{1, \ldots, k\}$ obeying $A \rightarrow BCD$, and let b_{kr} be the number of 4-ary r-tuple relations over $\{1, \ldots, k\}$ obeying $\{A \rightarrow BCD, BCD \rightarrow A\}$. Clearly a_{kr} and b_{kr} are positive precisely when $0 \leq r \leq k$, since each tuple of 4-ary relation obeying $A \rightarrow BCD$ must have a distinct A value, and there are only k possible values.

We now show that $a_{kr} = \binom{k}{r}k^{3r}$. This is because if R obeys $A \rightarrow BCD$, then each tuple of R must have a distinct A entry; there are $\binom{k}{r}$ possible choices for the set of r distinct A entries, and there are k^{3r} choices for how to fill out the remaining $3r$ entries in the B, C, and D columns.

We now show that $b_{kr} = \binom{k}{r}(k^3)(k^3 - 1) \cdots (k^3 - r + 1)$. This is because if R obeys $\{A \rightarrow BCD, BCD \rightarrow A\}$, then each tuple of R must have a distinct A entry and each tuple of R must have a distinct BCD entry; there are $\binom{k}{r}$ possible choices for the set of r distinct A entries, and there are $(k^3)(k^3 - 1) \cdots (k^3 - r + 1)$ choices for how to fill out the remaining $3r$ entries (which can be thought of as r distinct BCD tuples) in the B, C and D columns, given that no two BCD tuples can be the same. So

$$\begin{aligned}\frac{b_{kr}}{a_{kr}} &= \frac{k^3(k^3 - 1) \cdots (k^3 - r + 1)}{k^{3r}} \\ &= \frac{k^3}{k^3}\,\frac{k^3 - 1}{k^3} \cdots \frac{k^3 - r + 1}{k^3} \\ &= 1\left(1 - \frac{1}{k^3}\right)\left(1 - \frac{2}{k^3}\right) \cdots \left(1 - \frac{r-1}{k^3}\right) \\ &\geq \left(1 - \frac{1}{k^2}\right)^k, \qquad \text{since} \quad r \leq k.\end{aligned}$$

Now

$$\begin{aligned}\frac{b_k}{a_k} &\geq \min_{0 \leq r \leq k} \frac{b_{kr}}{a_{kr}} \\ &\geq \left(1 - \frac{1}{k^2}\right)^k,\end{aligned}$$

as we just showed.

By l'Hospital's rule, $(1 - 1/k^2)^k \rightarrow 1$ as $k \rightarrow \infty$. So, since $1 \geq (b_k/a_k) \geq (1 - 1/k^2)^k$, it follows that $b_k/a_k \rightarrow 1$, which was to be shown.

It is interesting to consider the question of the probability (as a function of Σ) that a relation obeying Σ is an Armstrong relation for Σ, that is, the limiting ratio (if it exists) of the number of Armstrong relations for Σ over $\{1, \ldots, k\}$ divided by the number of relations over $\{1, \ldots, k\}$ obeying Σ.

6. *Size of Minimal Armstrong Relations*

A *minimal* Armstrong relation is an Armstrong relation R for a set Σ of FDs such that every Armstrong relation for Σ has at least at many tuples as R. In this section we consider the size of (i.e., number of tuples in) minimal Armstrong relations, in terms of the number of intersection generators (defined soon), and also in terms of the number of attributes. We also consider the number of distinct entries that must appear in an Armstrong relation. We need some preliminary results, which are interesting in their own right.

Let Σ be a set of FDs, over the set U of attributes. A subset $V \subseteq U$ is *closed* if for every *FD* $X \rightarrow Y$ in Σ for which $X \subseteq V$, also $Y \subseteq V$. It is easy to see [1] that the intersection of closed sets is closed. Note that the minimal closed set containing X is X^*, where, as before, X^* is the set of all attributes A such that $\Sigma \models X \rightarrow A$.

Let M be a family of subsets of a finite set, closed under intersection. Then M contains a unique minimal subfamily M' such that the members of M' generate M by intersection [2]. Thus, M' is the smallest set such that $M = \{S_1 \cap \cdots \cap S_k: k \geq 0 \text{ and } S_1, \ldots, S_k \in M'\}$. The members of M' are the *intersection generators* of M. In fact, it is not hard to see that a member V of M is in M' if and only if V is properly contained in the intersection of the members of M that properly contain V. For a given set of Σ of FDs, denote by CL(Σ) the family of closed sets defined by Σ. As we noted, CL(Σ) is closed under intersection. Denote by GEN(Σ) the intersection generators of CL(Σ). Note that U is in CL(Σ) but not in GEN(Σ), since it is the intersection of the empty collection of sets.

Let t_1 and t_2 be tuples, and let X be a set of attributes. We say that t_1 and t_2 *agree exactly on* X if $t_1[X] = t_2[X]$, and if $t_1[A] \neq t_2[A]$ for each attribute A not in X. If R is a relation, then we define agr(R) to be $\{X$: there is a pair of distinct tuples in R that agree exactly on $X\}$. The next theorem is extremely useful as a characterization of Armstrong relations.

THEOREM 6.1. *Let Σ be a set of FDs, and let R be a relation. Then R is an Armstrong relation for Σ if and only if GEN(Σ) $\subseteq$ agr(R) $\subseteq$ CL(Σ).*

PROOF. $\Leftarrow$: Assume that GEN(Σ) $\subseteq$ agr(R) $\subseteq$ CL(Σ). Let $X \rightarrow A$ be in Σ, and assume that two tuples of R agree on X. Since by assumption agr(R) $\subseteq$ CL(Σ), we know that the tuples agree on a closed set. Hence they agree on X^*, and so they agree also on A. It follows that R satisfies Σ. Now let $X \rightarrow A$ be an FD not in Σ^*. Then $A \notin X^*$. By the discussion above, there exists a set $X' \in$ GEN(Σ) such that $X \subseteq X'$ and $A \in U - X'$. But GEN(Σ) $\subseteq$ agr(R), and so R contains two tuples that agree exactly on X'. Thus, R does not satisfy $X \rightarrow A$. We have shown that R is an Armstrong relation for Σ.

$\Rightarrow$: Let R be an Armstrong relation for Σ. We show first that agr(R) $\subseteq$ CL(Σ), that is, that every two tuples of R agree on a closed set (with respect to Σ). Indeed, since R satisfies Σ, if two tuples agree on a set then they agree on its closure. Hence agr(R) $\subseteq$ CL(Σ).

We now show that GEN(Σ) $\subseteq$ agr(R). Assume $X \in$ GEN(Σ); we shall show that $X \in$ agr(R). Let Y be the intersection of all closed sets that properly contain X. Then Y is a closed set. Since $X \in$ GEN(Σ), we know that $Y - X$ is not empty. Take A in $Y - X$. Since X is a closed set, the FD $X \rightarrow A$ is not in Σ^*, so R contains two tuples t_1 and t_2 that agree (at least) on X and disagree on A. We claim that t_1 and t_2 agree exactly on X. For, assume not. Then t_1 and t_2 agree on a closed set Z that properly contains X. By the definition of Y, they agree on Y, and hence on A. This is a contradiction. So, t_1 and t_2 agree exactly on X. Thus, $X \in$ agr(R). This was to be shown. □

We note that Ginsburg and Hull [16] independently proved that agr(R) $\subseteq$ CL(Σ) if and only if R obeys Σ; this result follows immediately from our proof of Theorem 6.1.

Let us define a *GEN set* to be a collection Q of subsets of U such that (a) $U \notin Q$, and (b) no member of Q is equal to the intersection of other members of Q. From what we have said, it is clear that GEN(Σ) is a GEN set for each set Σ of FDs. Conversely, Armstrong showed [1] that for every GEN set Q, there is a set Σ of FDs such that Q = GEN(Σ). Let $Q\dagger$ be the closure of Q under intersection. Thus, $Q\dagger = \{S_1 \cap \cdots \cap S_k : k \geq 0 \text{ and } S_1, \ldots, S_k \in Q\}$. We may say that a relation R is an *Armstrong relation for Q* if $Q \subseteq \text{agr}(R) \subseteq Q\dagger$. By Theorem 6.1, if Σ is a set of FDs and if Q = GEN(Σ), then R is an Armstrong relation for Σ (as defined earlier) if and only if R is an Armstrong relation for Q (as we just defined it). We shall usually denote a set Q fulfilling (a) and (b) above (that is, a GEN set) by GEN, since, as noted above, Q equals GEN(Σ) for some Σ. Similarly, we may write CL for $Q\dagger$.

Let us associate with each Armstrong relation R for GEN a graph $G(R)$. The nodes of the graph are the tuples of R. Two tuples are connected by an edge if they agree exactly on a set in GEN.

The next theorem provides a useful necessary condition for minimal Armstrong relations.

THEOREM 6.2. *If R is a minimal Armstrong relation for GEN, then $G(R)$ is connected.*

PROOF. Let R be a minimal Armstrong relation for GEN. Assume that $G(R)$ is not connected. Let $G(R)$ contain the connected components $H_1, \ldots, H_k$. Then $k > 1$, since $G(R)$ is not connected. Obtain relation R' from R by changing the values in the tuples of H_1 to be distinct from the values in the tuples of $G(R) - H_1$ except in the columns $\varnothing^*$. Denote by H_1' the connected component of $G(R')$ that corresponds to H_1 (note that H_1 and H_1' have the same number of tuples). It is easy to see that $\text{agr}(H_1') = \text{agr}(H_1)$, that $\text{agr}(G(R') - H_1') = \text{agr}(G(R) - H_1)$, and that tuples of H_1' agree with tuples of $G(R') - H_1'$ in $\varnothing^*$, a member of CL. It follows easily from these facts and from the fact that GEN $\subseteq \text{agr}(R) \subseteq$ CL (since R is an Armstrong relation for GEN) that GEN $\subseteq \text{agr}(R') \subseteq$ CL. Thus, the new relation R' is still an Armstrong relation for GEN. Now obtain R'' from R' by choosing arbitrarily tuples t_1 in H_1' and t_2 in H_2', and identifying t_1 and t_2. By "identifying t_1 and t_2," we mean to replace $t_1[A]$ by $t_2[A]$ for each attribute A everywhere $t_1[A]$ appears in H_1'. The new relation R'' has one less tuple than R'. Denote by H_1'' the result of transforming the tuples of H_1' (again H_1' and H_1'' have the same number of tuples). For tuples $r_1 \in H_1''$ and $r_2 \in G(R'') - H_1''$, tuple r_1 agrees with t_1 on a set X_1 in CL, and r_2 agrees with t_2 on a set X_2 in CL. Hence, in R'', tuples r_1 and r_2 agree exactly on $X_1 \cap X_2$, which is in CL. It is easy to show that what we have said implies that GEN $\subseteq \text{agr}(R'') \subseteq$ CL. Thus, R'' is an Armstrong relation for GEN, which contradicts minimality of R. □

The next theorem deals with the case where $\varnothing \in$ GEN. The corollary that follows this theorem will prove useful to us later. Let us associate with each Armstrong relation R for GEN a graph $H(R)$. The nodes of the graph are the tuples of R. Two tuples are connected by an edge if they agree exactly on a set in GEN $- \{\varnothing\}$. Recall that the graph $G(R)$ was defined similarly, except that in $G(R)$, two tuples are connected by an edge if they agree exactly on a set in GEN.

THEOREM 6.3. *Let R be a minimal Armstrong relation for GEN, where $\varnothing \in$ GEN. Then $H(R)$ contains exactly two connected components. Furthermore, in each column the values in the two components are distinct (that is, it is false that there* *are a column A and tuples t_1 and t_2 in distinct connected components such that $t_1[A] = t_2[A]$).*

PROOF. It is convenient to label the edges of $H(R)$ as follows. Let e be an arbitrary edge of $H(R)$ between nodes t_1 and t_2. Thus t_1 and t_2 are tuples of R. Assume that t_1 and t_2 agree exactly on X. We then label the edge e with the label X.

We shall prove the following fact, which we shall utilize several times.

FACT (∗). *Each pair of tuples from the same connected component of $H(R)$ agree on a nonempty closed set.*

We now prove Fact (∗). Let K be a connected component of $H(R)$. Each pair of tuples of K agree at least on the intersection of the labels on a path connecting them. Since the labels are from GEN − $\{\varnothing\}$, and since $\varnothing \in$ GEN, it follows that this intersection properly contains $\varnothing$. Thus, each pair of tuples of $H(R)$ agree on a nonempty set. The set on which the pair of tuples agree is also closed, since agr(R) ⊆ CL. This proves Fact (∗).

Now R must contain two tuples that agree precisely on $\varnothing$ (since GEN ⊆ agr(R)). By Fact (∗), these tuples are not in the same connected component of $H(R)$. Hence, $H(R)$ contains more than one connected component.

If $H(R)$ were to contain more than two connected components, then by a similar argument to that used in the proof of Theorem 6.2, it would be possible to merge two of the components, to yield an Armstrong relation with fewer tuples. (For those readers who are worried as to why *all* of the components cannot be merged into a single component, the answer is that GEN would not then be a subset of agr(R), since $\varnothing$ would not be in agr(R).)

We conclude the proof by showing that in each column, the values in distinct components are distinct. If not, assume that two tuples t_1 and t_2 from distinct components agree precisely on a nonempty set X_0. Of course, X_0 is closed, since agr(R) ⊆ CL. Since GEN ⊆ agr(R), there are tuples s_1 and s_2 that agree precisely on $\varnothing$, that is, $s_1[A] \neq s_2[A]$ for each attribute A. By Fact (∗) above, we know that s_1 and s_2 cannot be in the same connected component. So s_1 and s_2 are in distinct connected components. Since there are only two connected components, we can assume (by relabeling if necessary) that s_1 is in the same connected component as t_1, and that s_2 is in the same connected component as t_2. By Fact (∗), we know that s_1 and t_1 agree on a nonempty closed set X_1, and similarly, we know that s_2 and t_2 agree on a nonempty closed set X_2. So s_1 and s_2 agree at least on $X_0 \cap X_1 \cap X_2$, which is a nonempty closed set (it is nonempty because $\varnothing \in$ GEN). This contradicts our assumption that s_1 and s_2 agree precisely on $\varnothing$. □

COROLLARY 6.4. *If $\varnothing \in$ GEN, then a minimal Armstrong relation for GEN contains precisely one more tuple than a minimal Armstrong relation for GEN − $\{\varnothing\}$.*

PROOF. Let s be the number of rows in a minimal Armstrong relation for GEN, and let t be the number of rows in a minimal Armstrong relation for GEN − $\{\varnothing\}$. We must show that $s = t + 1$. Let R be a minimal Armstrong relation for GEN − $\{\varnothing\}$, and let R' be the result of adding to R one more tuple of all new distinct values. It follows easily from our characterization of Armstrong relations in Theorem 6.1 that R' is an Armstrong relation for GEN. Thus, $s \leq t + 1$. We conclude the proof by showing that $s \geq t + 1$.

Let S be a minimal Armstrong relation for GEN. By Theorem 6.3, we know that $H(S)$ contains exactly two connected components K_1 and K_2, and that in each column, the entries in each tuple in K_1 are distinct from those in K_2. Choose arbitrarily tuples t_1 in K_1 and t_2 in K_2 and identify them, as in the proof of Theorem 6.2. Call the resulting relation S. By an argument very similar to that used in the proof of Theorem 6.2, it follows easily that S is an Armstrong relation for GEN $-$ $\{\emptyset\}$. Hence, $t \leq s - 1$, and so $s \geq t + 1$, which was to be shown. □

The next theorem characterizes precisely how small or how big a minimal Armstrong relation can be, in terms of the size of GEN. In this theorem, $\lceil x \rceil$ is the result of rounding x up to an integer.

THEOREM 6.5. *Every minimal Armstrong relation for GEN contains at least* $\lceil (1 + (1 + 8r)^{1/2})/2 \rceil$ *tuples and at most* $r + 1$ *tuples, where* r *is the number of elements in GEN. Both bounds are attainable for each positive integer* r.

PROOF. Let R be a minimal Armstrong relation for Σ. By Theorem 6.2, we know that $G(R)$ is connected. But a connected graph with r edges has at most $r + 1$ nodes. Thus, $G(R)$ has at most $r + 1$ nodes, and so R has at most $r + 1$ tuples. As for the lower bound, assume that $G(R)$ has m nodes. Then the number of edges of $G(R)$ is at most $m(m - 1)/2$. Hence, $m(m - 1)/2 \geq r$, that is, $m \geq \lceil (1 + (1 + 8r)^{1/2})/2 \rceil$. Thus, R has at least $\lceil (1 + (1 + 8r)^{1/2})/2 \rceil$ tuples. Finally, the fact that the lower and upper bounds are attainable follows from two examples, which follow. □

Example 6.6. Assume that there are n attributes $A_1, \ldots, A_n$. Let GEN consist of the singleton sets $\{A_1\}, \{A_2\}, \ldots, \{A_n\}$. In particular, distinct members of GEN are disjoint. In this case, $r = n$. Let G be a graph with a minimal number of nodes that has r edges. (Thus G has m nodes, where $m = \lceil (1 + (1 + 8r)^{1/2})/2 \rceil$.) Label each edge of G with a member of GEN, such that distinct edges are labeled with distinct sets. It can be easily seen that we can construct a relation with tuples corresponding to the nodes of G such that two tuples agree exactly on A_i if their nodes are connected by an edge labeled with $\{A_i\}$, and have nothing in common otherwise. This relation is an Armstrong relation for GEN with a minimal number of tuples equal to the lower bound of Theorem 6.5. □

Example 6.7. Consider a sequence of strictly monotonically increasing sets $X_1 \subsetneq X_2 \subsetneq \cdots \subsetneq X_r \subsetneq U$, and let GEN $= \{X_1, \ldots, X_r\}$. GEN is closed under intersection, so GEN = CL. Let R be a minimal Armstrong relation for GEN. We prove, using induction on r, that R must contain $r + 1$ tuples.

For $r = 1$, at least two tuples are needed so that the nontrivial FDs with left-hand side X_1 are false in the resulting relation. Assume that all for $r \leq k$, at least $r + 1$ tuples are needed, and let $r = k + 1$. Now, by the induction hypothesis, a minimal Armstrong relation for $\{X_2 - X_1, \ldots, X_r - X_1\}$ contains r tuples. Hence, by Corollary 6.4, a minimal Armstrong relation for $\{\emptyset, X_2 - X_1, \ldots, X_r - X_1\}$ contains $r + 1$ tuples. Let R be a minimal Armstrong relation for GEN. It is not hard to see that $R[U - X_1]$ is then a minimal Armstrong relation for $\{\emptyset, X_2 - X_1, \ldots, X_r - X_1\}$ (and they have the same number of tuples). But we just showed that a minimal Armstrong relation for $\{\emptyset, X_2 - X_1, \ldots, X_r - X_1\}$ contains $r + 1$ tuples. So, a minimal Armstrong relation for GEN contains $r + 1$ tuples. This completes the induction. We remark that another example of a GEN set with r members, whose minimal Armstrong relation contains $r + 1$ tuples, is given by letting GEN be the collection of all $(r - 1)$-element subsets of $\{1, \ldots, r\}$. □

In Theorem 6.5, we considered the size of a minimal Armstrong relation as a function of the size of GEN. We now consider the size of a minimal Armstrong relation as a function of the number of attributes. As before, let us denote the binomial coefficient $\binom{n}{\lfloor n/2 \rfloor}$ by $S(n)$. By Stirling's formula, it follows that $S(n)$ is asymptotic to $(2/\pi)^{1/2}2^n n^{-1/2}$.

LEMMA 6.8. *There is a constant c such that each GEN set over n attributes has less than* $S(n)(1 + (c/n^{1/2}))$ *members.*

PROOF. Kleitman (see [9]) showed that there is c such that $S(n)(1 + c/n^{1/2}))$ is an upper bound on the size of a set of subsets of $\{1, \ldots, n\}$ with the property that no one is the intersection of two others. Now a GEN set has the even stronger property that no member is the intersection of any collection of the others. The result follows immediately. □

We now give upper and lower bounds on the size (number of tuples in) the biggest minimal Armstrong relation with n attributes.

THEOREM 6.9. *There is a constant c such that every minimal Armstrong relation contains less than* $S(n)(1 + (c/n^{1/2}))$ *tuples, where n is the number of attributes. There is a set* Σ *of FDs such that each Armstrong relation for* Σ *contains more than* $S(n)/n^2$ *tuples. Thus, let* $p(n)$ *be the size of the biggest minimal Armstrong relation with n attributes. That is,* $p(n)$ *is the maximum (over all sets* Σ *of FDs, where* Σ *involves n attributes) of the minimum number of tuples (over all Armstrong relations for* Σ*). Then* $S(n)/n^2 < p(n) < S(n)(1 + (c/n^{1/2}))$.

PROOF. We first consider the upper bound $S(n)(1 + c/n^{1/2}))$. By Theorem 6.5, the number of tuples is at most $r + 1$, where r is the size of GEN(Σ). By Lemma 6.8, we know that $r < S(n)(1 + c_1/n^{1/2})$ for some constant c_1. So, the number of tuples is less than $1 + S(n)(1 + c_1/n^{1/2})$. Now, $1 + S(n)(1 + c_1/n^{1/2}) \leq S(n)(1 + (c_1 + 1)/n^{1/2})$, since $S(n) \geq n^{1/2}$. Hence, we can take $c = c_1 + 1$.

We now prove the lower bound $S(n)/n^2$. Our proof is a slight modification of a proof of a related result by Ronyai, which appears in [8]. Let us denote $G(n)$ the number of GEN sets over n attributes. Recall that a GEN set over n attributes $A_1, \ldots, A_n$ is a set of subsets of $\{A_1, \ldots, A_n\}$, such that no member equals the intersection of a collection of other members (and such that $U = \{A_1, \ldots, A_n\}$ is not a member). The set V, consisting of all subsets of $\lfloor n/2 \rfloor$ attributes, is a GEN set. Furthermore, every subset of a GEN set is a GEN set, and so every subset of V is a GEN set. Since V contains $S(n)$ members, it has $2^{S(n)}$ distinct subsets. Hence,

$$G(n) \geq 2^{S(n)}. \tag{6.1}$$

Suppose an Armstrong relation has t tuples. Then, without loss of generality, it can be assumed (by renaming if necessary) that its entries are taken from the set $\{1 \ldots, t\}$. Let us call a relation *special* if

(1) it is a minimal Armstrong relation for some GEN set, and
(2) if it contains t tuples, then its entries are taken from the set $\{1, \ldots, t\}$.

Let us denote the number of special relations over n attributes by $H(n)$. Then

$$H(n) \geq G(n). \tag{6.2}$$

This follows from the fact that two Armstrong relations for two distinct GEN sets are nonisomorphic, hence distinct.

Let $p(n)$ be as in the statement of the theorem. Each special relation has at most $np(n)$ entries, since there are n columns and at most $p(n)$ tuples. Further, each entry lies in $\{1, \ldots, p(n)\}$, because the largest entry in a special relation with t tuples is at most t. Since there are at most $np(n)$ entries, and since the entries lie in $\{1, \ldots, p(n)\}$, the total number of special relations is at most $p(n)^{np(n)}$. That is,

$$p(n)^{np(n)} \geq H(n). \tag{6.3}$$

By (6.1)–(6.3), we obtain

$$p(n)^{np(n)} \geq 2^{S(n)}. \tag{6.4}$$

We now show that (6.4) implies that $p(n) > S(n)/n^2$. Assume not. Then for some n, we have $p(n) \leq S(n)/n^2$. Thus, $S(n)/n^2 > 1$. Since the function mapping x into x^{nx} is monotone increasing in x (for $x \geq 1$), the fact that $S(n)/n^2 \geq p(n)$ implies that

$$\left(\frac{S(n)}{n^2}\right)^{nS(n)/n^2} \geq p(n)^{np(n)}. \tag{6.5}$$

By (6.4) and (6.5), it follows that

$$\left(\frac{S(n)}{n^2}\right)^{nS(n)/n^2} \geq 2^{S(n)}. \tag{6.6}$$

An easy simplification of (6.6) gives

$$S(n) \geq n^2 2^n. \tag{6.7}$$

But $S(n) < 2^n$ (in fact, it is well known that $2^n = \sum_{i=0}^{n} \binom{n}{i}$, where one of the terms in the sum is $S(n)$). Since $S(n) < 2^n$, it is all the more true that $S(n) < n^2 2^n$. This contradicts (6.7). □

We remark that by a simple modification of our proof, we can show that for each constant k, we have $p(n) > S(n)/(n - k)^2$ for n sufficiently large.

We now consider how large the domain size must be in an Armstrong relation, that is, we consider the number of distinct entries in each column.

THEOREM 6.10. *There is a constant c such that every minimal Armstrong relation contains less than $S(n)(1 + c/n^{1/2})$ distinct entries in each column, where n is the number of attributes. There is a set Σ of FDs such that each Armstrong relation for Σ contains more than $S(n)/(2n^2)$ entries in some column.*

PROOF. The upper bound follows from Theorem 6.9, since the number of distinct entries in each column is bounded by the number of tuples. We now consider the lower bound.

Let $m = n - 1$. By Theorem 6.9, where we let m play the role of n, we know that there is a set Σ' of FDs (over $n - 1$ attributes $A_1, \ldots, A_{n-1}$) such that each Armstrong relation for Σ' contains more than $S(n - 1)/(n - 1)^2$ tuples. Let Σ contain Σ', along with exactly one more FD $A_n \rightarrow A_1 \ldots A_{n-1}$. Thus, the new FD says that the new attribute A_n is a key. Each Armstrong relation R for Σ contains more than $S(n - 1)/(n - 1)^2$ tuples, since the projection of R onto the first $(n - 1)$ attributes is an Armstrong relation for Σ' with as many tuples as R. Since A_n is a key, every tuple has a distinct A_n entry. Thus, the A_n column contains more than $S(n - 1)/(n - 1)^2$ entries. Simple algebra shows that $S(n - 1)/(n - 1)^2 \geq S(n)/(2n^2)$. The result follows. □

The lower bound of Theorem 6.10 provides a negative answer to Gold's question [18] as to whether, for each set Σ of FDs, there is an Armstrong relation for Σ that is 0–1 valued (that is, whose only entries are 0's and 1's). It also strengthens a result, obtained independently by Ginsburg and Zaiddan [17], that there is a set Σ of FDs such that each Armstrong relation for Σ contains more than three entries in some column.

7. *Complexity Results*

In this section we show that the complexity of finding an Armstrong relation, given a set of FDs, is precisely exponential in the number of attributes. We also show that the problem of deciding if there is a key of size at most k is NP-complete [15], whether the set of FDs is presented explicitly, or implicitly via an Armstrong relation.

We begin with the issue of the complexity of finding an Armstrong relation, given a set of FDs. We shall show:

(i) There is an algorithm for obtaining an Armstrong relation, given the set Σ of FDs, where the running time of the algorithm is exponential in the number of attributes (by *exponential*, we mean time c^n for some constant c; in fact, our proof shows that c can be taken to be $2 + \epsilon$ for arbitrary $\epsilon > 0$); and

(ii) There is a set Σ of FDs in which the number of tuples in each minimal Armstrong relation for Σ is exponential—thus, an exponential amount of time is required in this case simply to write down the relation. We shall, in fact, show that for each $\epsilon > 0$ there is N such that if $n > N$, then there is a set Σ of FDs over n attributes such that every Armstrong relation for Σ has at least $(2 - \epsilon)^n$ tuples.

Because of (i) and (ii) above, we say that the complexity of finding an Armstrong relation is *precisely exponential* in the number of attributes.

We can prove a result that is stronger than (ii). Specifically, we can (and shall) exhibit a set Σ of functional dependencies such that the number of tuples in a minimal Armstrong relation for Σ is exponential, not only in the number of attributes (as demanded by (ii)), but also in the number of functional dependencies. We begin by proving (i) and (ii).

THEOREM 7.1. *The complexity of finding an Armstrong relation, given a set of FDs, is precisely exponential in the number of attributes.*

PROOF. We first present an exponential-time algorithm for finding an Armstrong relation, given a set Σ of FDs. Our construction is very similar to that of Gold [18]. It is also reminiscent of the "partially disconnected augmentation" technique of Ginsburg and Hull [16]. Let n be the number of attributes. The algorithm first cycles through each of the 2^n subsets of attributes, and checks which are closed (with respect to Σ). Let S be the collection $\mathrm{CL}(\Sigma)$ of closed sets. (We could get away with using $\mathrm{GEN}(\Sigma)$ instead of $\mathrm{CL}(\Sigma)$ as S in the construction that follows, but we do not wish to spend the time to prune out the nongenerators). Assume that the distinct members of S are $S_1, \ldots, S_r$. Let t_i $(1 \leq i \leq r)$ be a tuple, where $t[A] = 0$ if A is an attribute in S_i, and where $t[A] = i$ for each of the other attributes. The desired relation contains a tuple of all 0's, along with each of the tuples t_i $(1 \leq i \leq r)$. By Theorem 6.1, it follows easily that as long as $\mathrm{GEN}(\Sigma) \subseteq S \subseteq \mathrm{CL}(\Sigma)$, this contruction produces an Armstrong relation for Σ. It is clear that this algorithm has an exponential running time (exponential in the number of attributes), since the size of Σ is at most exponential in the number of attributes, and checking whether a set X is closed can be done in time polynomial in the size of Σ and the set X [3].

To complete the proof of the theorem, we must show (ii) above: that there is a set Σ of FDs for which the number of tuples in each minimal Armstrong relation for Σ is exponential in the number of attributes. By Theorem 6.9, there is a set Σ of FDs such that a minimal Armstrong relation for Σ has more than $S(n)/n^2$ tuples. By Stirling's formula, $S(n)$ is asymptotic to $(2/\pi)^{1/2}2^n n^{-1/2}$. Thus, the size of a minimal Armstrong relation is asymptotically greater than $(2/\pi)^{1/2}2^n n^{-5/2}$. But this value is asymptotically greater than $(2 - \epsilon)^n$ for each $\epsilon > 0$. This shows (ii) above. □

It is easy to obtain an explicit example of a GEN set GEN such that the size of a minimal Armstrong relation is exponential in the number of attributes. We simply take GEN to be the set of all $\lfloor n/2 \rfloor$-sized subsets of the n attributes. Since this GEN set has $S(n)$ members, it follows from the lower bound of Theorem 6.5 that each Armstrong relation for GEN has at least $\lceil (1 + (1 + 8S(n))^{1/2})/2 \rceil$ tuples. This number asymptotically dominates k^n for each $k < 2^{1/2}$.

In the example we just presented, the number of FDs was itself exponential in the number of attributes. We now exhibit another example, in which the size of each Armstrong relation is exponential, but for which the number of FDs is small (in fact, less than the number of attributes). Let there be $2m + 1$ attributes $A_1, A_2, \ldots, A_{2m}, B$. The set Σ of FDs is

$$\begin{gathered} A_1A_2 \rightarrow B, \\ A_3A_4 \rightarrow B, \\ \vdots \\ A_{2m-1}A_{2m} \rightarrow B. \end{gathered}$$

Let T be an arbitrary set containing exactly one attribute from the left-hand side of each of the FDs above, and not containing B. Thus, T has exactly m attributes, and there are exactly 2^m such sets T. We now show that $T \in \mathrm{GEN}(\Sigma)$.

If T is not in GEN(Σ), then it is the intersection $\bigcap T_i$ of a family of closed sets, each of which properly contains T. Since $B \notin T$, clearly $B \notin T_j$ for some j. We shall show that for this j, necessarily $T = T_j$, a contradiction. Clearly $T \subseteq T_j$. If T_j were to contain some attribute A_i not in T, then T_j would contain the entire left-hand of one of the FDs above. But then, T_j would contain B, which it does not. Hence, $T = T_j$, as desired.

We have shown that GEN(Σ) contains each such set T, and so contains (at least) 2^m members. So by Theorem 6.9, each Armstrong relation for Σ contains at least $\lceil (1 + (1 + 8r)^{1/2})/2 \rceil$ tuples, where $r = 2^m$. Now $m = (n - 1)/2$, where n is the number of attributes. From a simple computation, we conclude that a minimal Armstrong relation for Σ contains at least k^n tuples (where $k = 2^{1/4}$). This was to be shown.

We close by showing that the problem of deciding whether there is a key of size at most k is NP-complete. There are (at least) two possible methods of presenting an input to this problem:

(a) by presenting a set Σ of FDs, and
(b) by presenting an Armstrong relation for Σ.

THEOREM 7.2. *For either type of presentation, the problem of deciding whether there is a key of size of most k is NP-complete.*

PROOF. The problem is clearly in NP. We shall show that the vertex cover problem is reducible to this problem. Let G be a graph with V as its set of vertices and E as its set of edges. The vertex cover problem asks if there is a set X of at most k vertices in V such that every edge contains a vertex in X. It is well known [15] that the vertex cover problem is NP-complete.

Let R be a relation with attributes V. The entries of the relation R will be members of E, along with a new value z. There will be $|E| + 1$ tuples, indexed by $E \cup \{z\}$. Every entry in the z tuple is z. If e is in E, and if vertex a is in edge e in G, then the a entry in the e tuple is e; otherwise the a entry in the e tuple is z. It is straightforward to verify that if X is a set of attributes and if a is a single attribute not in X, then the FD $X \rightarrow a$ holds in R if and only if $\{b:(a, b) \in E\} \subseteq X$. Hence, X is a key (for relation R) if and only if X is a vertex cover (for graph G). This proves the result for presentations of type (b). Now, let Σ be the set $\{X \rightarrow a: X = \{b:(a, b) \in E\}\}$ of FDs. As we noted, these are precisely the FDs that hold in R. This proves the result for presentations of type (a). □

Lucchesi and Osborn [20] prove the theorem for presentations of type (a). The problem of determining whether the FD $X \rightarrow Y$ holds, given a presentation of type (b), can be solved in logspace. On the other hand, the problem of determining whether the FD $X \rightarrow Y$ holds, given a presentation of type (a), is logspace-complete for P; there is an easy reduction of whether a propositional Horn formula follows from a set of Horn formulas (see [24]).

REFERENCES

1. ARMSTRONG, W.W. Dependency structures of database relationships. In *Proc. IFIP 74*, North Holland, Amsterdam, 1974, pp. 580–583.
2. ARMSTRONG, W.W., AND DELOBEL, C. Decomposition and functional dependencies in relations. *ACM Trans Database Syst 5*, 4 (Dec. 1980), 404–430.
3. BEERI, C., AND BERNSTEIN, P A Computational problems related to the design of normal form relational schemas. *ACM Trans. Database Syst 4*, 1 (Mar. 1979), 30–59.
4. BEERI, C., FAGIN, R., AND HOWARD, J.H. A complete axiomatization for functional and multivalued dependencies in database relations. In *Proc ACM-SIGMOD Int. Conf. on Management of Data* (Toronto, Aug. 1977), ACM, New York, pp. 47–61.
5. CASANOVA, M.A., FAGIN, R., AND PAPADIMITRIOU, C. Inclusion dependencies and their interaction with functional dependencies. *SIAM J. Comput* (to appear); extended abstract appeared in *Proc. 1st ACM SIGACT-SIGMOD Principles of Database Systems* (Los Angeles, Calif., 1982), ACM, New York, pp. 171–176.
6. CODD, E.F. Further normalization of the database relational model. In *Data Base Systems*, Courant Computer Science Symposia 6, R. Rustin, Ed., Prentice Hall, Englewood Cliffs, N.J., pp. 65–98.
7. DEMETROVICS, J. On the number of candidate keys. *Inf Proc Lett 7*, 6 (Oct. 1978), 266–269.
8. DEMETROVICS, J., AND GYEPESI, G. On the functional dependency and some generalizations of it. *Acta Cyberne 5*, 3 (1981), 295–305.
9. ERDOS, P., AND KLEITMAN, D. Extremal problems among subsets of a given set. In *Paul Erdos: The Art of Counting—Selected Writings*, J. Spencer, Ed., MIT Press, 1973, pp. 482–506.
10 FAGIN, R. Probabilities on finite models. *J Symb Logic 41*, 1 (Mar. 1976), 50–58.
11. FAGIN, R. Multivalued dependencies and a new normal form for relational databases. *ACM Trans. Database Syst. 2*, 3 (Sept. 1977), 262–278.
12. FAGIN, R. Functional dependencies in a relational database and propositional logic. *IBM J. Res Devel 21*, 6 (Nov. 1977), 534–544.
13. FAGIN, R. Armstrong databases. In *Proc 7th IBM Symp. on Mathematical Foundations of Computer Science* (Kanagawa, Japan, May 1982); also appeared as Res. Rep. RJ3440, IBM Research Lab., San Jose, Calif., Apr. 1982.

14. FAGIN, R. Horn clauses and database dependencies. *J. ACM 29*, 4 (Oct. 1982), 952–985; extended abstract appeared in *Proc. 12th Ann ACM Symp. on Theory of Computing* (Los Angeles, Calif., 1980), ACM, New York, pp. 123–134.
15. GAREY, M.R., AND JOHNSON, D.S. *Computers and Intractibility: A Guide to the Theory of NP-Completeness*, Freeman, San Francisco, 1979.
16. GINSBURG, S., AND HULL, R. Characterizations for functional dependency and Boyce–Codd normal form families. Tech. Rep., Univ. of Southern California, Los Angeles, Calif., Feb. 1982.
17. GINSBURG, S., AND ZAIDDAN, S.M. Properties of functional-dependency families. *J. ACM 29*, 3 (July 1982), 678–698.
18. GOLD, E.M. Axiomatization of attribute dependencies. Publ. #173, Départemenf d'Informatique, Université de Montréal, May 1974.
19 GRATZER, G. *Universal Algebra*, 2nd ed. Springer-Verlag, New York, 1979.
20. LUCCHESI, C.L., AND OSBORN, S.L. Candidate keys for relations, *J. Comput Syst. Sci. 17*, 2 (Oct. 1978), 270–279.
21. SAGIV, Y., DELOBEL, C., PARKER, D.S., AND FAGIN, R. An equivalence between relational database dependencies and a fragment of propositional logic. *J. ACM 28*, 3 (July 1981), 435–453.
22. SILVA, A.M., AND MELKANOFF, M.A. A method for helping discover the dependencies of a relation. In *Advances in Database Theory, Vol.* 1, H. Gallaire, J. Minker, and J-M. Nicolas, Eds., Plenum Publishing, New York, 1981.
23. SPERNER, E., Eine Satz uber Untermensen einer endlichen Menge. *Math. Zeit. 27* (1928), 544–548.
24. STATMAN, R. Worst case exponential lower bounds for input resolution with paramodulation. *SIAM J Comput.* (to appear).
25. ULLMAN, J.D. *Principles of Database Systems*, 2nd ed. Computer Science Press, Rockville, Md., 1982.
26. ZANIOLO, C. Analysis and design of relational schemata for database systems. Ph.D. Dissertation, Tech. Rep. UCLA-ENG-7669, Univ. of California, Los Angeles, Calif., July 1976.
27. ZLOOF, M.M. Query-by-example. In *Proc. 1975 National Computer Conf. 44*, AFIPS Press, Arlington, Va., pp. 431–438.

RECEIVED SEPTEMBER 1980; REVISED JULY 1982; ACCEPTED JULY 1982