

Models for Incomplete and Probabilistic Information

Todd J. Green* and Val Tannen**

University of Pennsylvania
{tjgreen, val}@cis.upenn.edu

Abstract. We discuss, compare and relate some old and some new models for incomplete and probabilistic databases. We characterize the expressive power of c -tables over infinite domains and we introduce a new kind of result, algebraic completion, for studying less expressive models. By viewing probabilistic models as incompleteness models with additional probability information, we define completeness and closure under query languages of general probabilistic database models and we introduce a new such model, probabilistic c -tables, that is shown to be complete and closed under the relational algebra.

1 Introduction

The representation of incomplete information in databases has been an important research topic for a long time, see the references in [18], in Ch.19 of [2], in [31], in [35,25], as well as the recent [33,30,29]. Moreover, this work is closely related to recently active research topics such as inconsistent databases and repairs [4], answering queries using views [1], and data exchange [13]. The classic reference on incomplete databases remains [20] with the fundamental concept of c -table and its restrictions to simpler tables with variables. The most important result of [20] is the query answering algorithm that defines an algebra on c -tables that corresponds exactly to the usual relational algebra (\mathcal{RA}). A recent paper [29] has defined a hierarchy of incomplete database models based on finite sets of choices and optional inclusion. One of our contributions consists of **comparisons** between the models [29] and the tables with variables from [20].

Two criteria have been provided for comparisons among all these models: [20, 29] discuss *closure* under relational algebra operations, while [29] also emphasizes *completeness*, specifically the ability to represent all finite incomplete databases. We point out that the latter is not appropriate for tables with variables over an infinite domain, and we contribute another criterion, **\mathcal{RA} -completeness**, that fully characterizes the expressive power of c -tables.

We also introduce a new idea for the study of models that are not complete. Namely, we consider combining existing models with queries in various fragments of relational algebra. We then ask how big these fragments need to be to

* Partially supported by NSF IIS-0477972 and NSF SEIII IIS-0513778.

** Partially supported by NSF SEIII IIS-0513778.

obtain a combined model that is complete. We give a number of such **algebraic completion** results.

Early on, probabilistic models of databases were studied less intensively than incompleteness models, with some notable exceptions [7, 5, 28, 23, 10]. Essential progress was made independently in three papers [15, 22, 34] that were published at about the same time. [15, 34] assume a model in which tuples are taken independently in a relation with given probabilities. [22] assumes a model with a separate distribution for each attribute in each tuple. All three papers attacked the problem of calculating the probability of tuples occurring in query answers. They solved the problem by developing more general models in which rows contain additional information (“event expressions”, “paths”, “traces”), and they noted the similarity with the conditions in *c*-tables.

We go beyond the problem of individual tuples in query answers by defining **closure** under a query language for probabilistic models. Then we develop a new model, **probabilistic *c*-tables** that adds *to the *c*-tables themselves* probability distributions for the values taken by their variables. Here is an example of such a representation that captures the set of instances in which Alice is taking a course that is Math with probability 0.3; Physics (0.3); or Chemistry (0.4), while Bob takes the same course as Alice, provided that course is Physics or Chemistry and Theo takes Math with probability 0.85:

<i>Student</i>	<i>Course</i>	<i>Condition</i>
Alice	x	
Bob	x	$x = \text{phys} \vee x = \text{chem}$
Theo	math	$t = 1$

$$x = \begin{cases} \text{math} : 0.3 \\ \text{phys} : 0.3 \\ \text{chem} : 0.4 \end{cases} \quad t = \begin{cases} 0 : 0.15 \\ 1 : 0.85 \end{cases}$$

The concept of probabilistic *c*-table allows us to solve the closure problem by using the same algebra on *c*-tables defined in [20].

We also give a **completeness** result by showing that probabilistic boolean *c*-tables (all variables are two-valued and can appear only in the conditions, not in the tuples) can represent *any* probabilistic database.

An important conceptual contribution is that we show that, at least for the models we consider, the probabilistic database models can be seen, as **probabilistic counterparts** of incomplete database models. In an incompleteness model a tuple or an attribute value in a tuple may or may not be in the database. In its probabilistic counterpart, these are seen as elementary events with an assigned probability. For example, the models used in [15, 22, 34] are probabilistic counterparts of the two simplest incompleteness models discussed in [29]. As another example, the model used in [10] can be seen as the probabilistic counterpart of an incompleteness model one in which tuples sharing the same key have an exclusive-or relationship.

A consequence of this observation is that, in particular, query answering for probabilistic *c*-tables will allow us to solve the problem of calculating probabilities about query answers for any model that can be defined as a probabilistic counterpart of the incompleteness models considered in [20, 29].

This paper is purely theoretical. Nonetheless, it was motivated by the work the authors are doing with others on the Orchestra¹ and SHARQ² projects. These projects are concerned with certain aspects of collaborative information sharing. Incompleteness arises in Orchestra (a peer-to-peer data exchange system) in the process of update propagation between sites. Incompleteness is also exploited in query answering algorithms. Probabilistic models are used in SHARQ (a bio-informatics data sharing system) to model approximate mappings between schemas used by groups of researchers. The sources of uncertainty here include data from error-prone experiments and accepted scientific hypotheses that allow for the limited mismatch. We expect that the results of this paper will help us in choosing appropriate representation systems that will be used internally in the Orchestra and SHARQ systems.

2 Incomplete Information and Representation Systems

Our starting point is suggested by the work surveyed in [18], in Ch. 19 of [2], and in [31]. A database that provides incomplete information consists of a *set of possible instances*. At one end of this spectrum we have the conventional single instances, which provide “complete information.” At the other end we have the set of *all* allowable instances which provides “no information” at all, or “zero information.”

We adopt the formalism of relational databases over a fixed countably infinite domain \mathbb{D} . We use the unnamed form of the relational algebra. To simplify the notation we will work with relational schemas that consist of a single relation name of arity n . Everything we say can be easily reformulated for arbitrary relational schemas. We shall need a notation for the set of *all* (conventional) instances of this schema, i.e., all the finite n -ary relations over \mathbb{D} :

$$\mathcal{N} := \{I \mid I \subseteq \mathbb{D}^n, I \text{ finite}\}$$

Definition 1. An **incomplete(-information) database** (*i-database for short*), \mathcal{I} , is a set of conventional instances, i.e., a subset $\mathcal{I} \subseteq \mathcal{N}$.

The usual relational databases correspond to the cases when $\mathcal{I} = \{I\}$. The **no-information** or **zero-information database** consists of *all* the relations: \mathcal{N} .

Conventional relational instances are finite. However, because \mathbb{D} is infinite incomplete databases are in general infinite. Hence the interest in finite, syntactical, representations for incomplete information.

Definition 2. A **representation system** consists of a set (usually a syntactically defined “language”) whose elements we call tables, and a function *Mod* that associates to each table T an incomplete database $\text{Mod}(T)$.

The notation corresponds to the fact that T can be seen as a logical assertion such that the conventional instances in $\text{Mod}(T)$ are in fact the *models* of T (see also [27, 32]).

¹ <http://www.cis.upenn.edu/~zives/orchestra>

² <http://db.cis.upenn.edu/projects/SHARQ>

The classical reference [20] considers three representation systems: **Codd tables**, **v-tables**, and **c-tables**. *v*-tables are conventional instances in which variables can appear in addition to constants from \mathbb{D} . If T is a *v*-table then³

$$Mod(T) := \{\nu(T) \mid \nu : Var(T) \rightarrow \mathbb{D} \text{ is a valuation for the variables of } T\}$$

Codd tables are *v*-tables in which all the variables are distinct. They correspond roughly to the current use of nulls in SQL, while *v*-tables model “labeled” or “marked” nulls. *c*-tables are *v*-tables in which each tuple is associated with a condition — a boolean combination of equalities involving variables and constants. We typically use the letter φ for conditions. The tuple condition is tested for each valuation ν and the tuple is discarded from $\nu(T)$ if the condition is not satisfied.

Example 1. Here is an example of a *v*-table.

$$R := \begin{array}{|c|c|c|} \hline 1 & 2 & x \\ \hline 3 & x & y \\ \hline z & 4 & 5 \\ \hline \end{array} \quad Mod(R) = \left\{ \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 3 & 1 & 1 \\ \hline 1 & 4 & 5 \\ \hline \end{array}, \begin{array}{|c|c|c|} \hline 1 & 2 & 2 \\ \hline 3 & 2 & 1 \\ \hline 1 & 4 & 5 \\ \hline \end{array}, \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 3 & 1 & 2 \\ \hline 1 & 4 & 5 \\ \hline \end{array}, \dots, \begin{array}{|c|c|c|c|} \hline 1 & 2 & 77 \\ \hline 3 & 77 & 89 \\ \hline 97 & 4 & 5 \\ \hline \end{array}, \dots \right\}$$

Example 2. Here is an example of a *c*-table.

$$S := \begin{array}{|c|c|c|c|} \hline 1 & 2 & x & \\ \hline 3 & x & y & x = y \wedge z \neq 2 \\ \hline z & 4 & 5 & x \neq 1 \vee x \neq y \\ \hline \end{array} \quad Mod(S) = \left\{ \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 3 & 1 & 1 \\ \hline \end{array}, \begin{array}{|c|c|c|} \hline 1 & 2 & 2 \\ \hline 1 & 4 & 5 \\ \hline \end{array}, \dots, \begin{array}{|c|c|c|c|} \hline 1 & 2 & 77 \\ \hline 97 & 4 & 5 \\ \hline \end{array}, \dots \right\}$$

Several other representation systems have been proposed in a recent paper [29]. We illustrate here three of them and we discuss several others later. A **?-table** is a conventional instance in which tuples are optionally labeled with “?” meaning that the tuple may be missing. An **or-set-table** looks like a conventional instance but or-set values [21,26] are allowed. An or-set value $\langle 1, 2, 3 \rangle$ signifies that exactly one of 1, 2, or 3 is the “actual” (but unknown) value. Clearly, the two ideas can be combined yielding another representation systems that we might (awkwardly) call **or-set-?-tables**.⁴

Example 3. Here is an example of an or-set-?-table.

$$T := \begin{array}{|c|c|c|c|} \hline 1 & 2 & \langle 1, 2 \rangle & \\ \hline 3 & \langle 1, 2 \rangle & \langle 3, 4 \rangle & \\ \hline \langle 4, 5 \rangle & 4 & 5 & ? \\ \hline \end{array} \quad Mod(T) = \left\{ \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 3 & 1 & 3 \\ \hline 4 & 4 & 5 \\ \hline \end{array}, \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 3 & 1 & 3 \\ \hline \end{array}, \begin{array}{|c|c|c|} \hline 1 & 2 & 2 \\ \hline 3 & 1 & 3 \\ \hline 4 & 4 & 5 \\ \hline \end{array}, \dots, \begin{array}{|c|c|c|} \hline 1 & 2 & 2 \\ \hline 3 & 2 & 4 \\ \hline \end{array} \right\}$$

3 \mathcal{RA} -Completeness and Finite Completeness

“Completeness” of expressive power is the first obvious question to ask about representation systems. This brings up a fundamental difference between the

³ We follow [2, 29] and use the *closed-world assumption (CWA)*. [20] uses the *open-world assumption (OWA)*, but their results hold for CWA as well.

⁴ In [29] these three systems are denoted by $\mathcal{R}_?$, \mathcal{R}^A and $\mathcal{R}_?^A$.

representation systems of [20] and those of [29]. The presence of variables in a table T and the fact that \mathbb{D} is infinite means that $Mod(T)$ may be infinite. For the tables considered in [29], $Mod(T)$ is always finite.

[29] defines completeness as the ability of a representation system to represent “all” possible incomplete databases. For the kind of tables considered in [29] the question makes sense. But in the case of the tables with variables in [20] this is hopeless for trivial reasons. Indeed, in such systems there are only countably many tables while there are uncountably many incomplete databases (the subsets of \mathcal{N} , which is infinite). We will discuss separately below *finite completeness* for systems that only represent finite database. Meanwhile, we will develop a different yardstick for the expressive power of tables with variables that range over an infinite domain.

c -tables and their restrictions (v -tables and Codd tables) have an inherent limitation: the cardinality of the instances in $Mod(T)$ is at most the cardinality of T . For example, the zero-information database \mathcal{N} cannot be represented with c -tables. It also follows that among the incomplete databases that are representable by c -tables the “minimal”-information ones are those consisting for some m of all instances of cardinality up to m (which are in fact representable by Codd tables with m rows). Among these, we make special use of the ones of cardinality 1:

$$\mathcal{Z}_k := \{\{t\} \mid t \in \mathbb{D}^k\}.$$

Hence, \mathcal{Z}_k consists of *all* the one-tuple relations of arity k . Note that $\mathcal{Z}_k = Mod(Z_k)$ where Z_k is the Codd table consisting of a single row of k distinct variables.

Definition 3. *An incomplete database \mathcal{I} is \mathcal{RA} -definable if there exists a relational algebra query q such that $\mathcal{I} = q(\mathcal{Z}_k)$, where k is the arity of the input relation name in q .*

Theorem 1. *If \mathcal{I} is an incomplete database representable by a c -table T , i.e., $\mathcal{I} = Mod(T)$, then \mathcal{I} is \mathcal{RA} -definable.*

Proof. Let T be a c -table, and let $\{x_1, \dots, x_k\}$ denote the variables in T . We want to show that there exists a query q in \mathcal{RA} such that $q(Mod(\mathcal{Z}_k)) = Mod(T)$. Let n be the arity of T . For every tuple $t = (a_1, \dots, a_n)$ in T with condition φ_t , let $\{x_{i_1}, \dots, x_{i_j}\}$ be the variables in φ_t which do not appear in t . For $1 \leq i \leq n$, define C_i to be the singleton $\{c\}$, if $a_i = c$ for some constant c , or $\pi_j(Z_k)$, if $a_i = x_j$ for some variable x_j . For $1 \leq j \leq k$, define C_{n+j} to be the expression $\pi_{i_j}(Z_k)$, where x_j is the j th variable in φ_t which does not appear in t . Define q to be the query

$$q := \bigcup_{t \in T} \pi_{1, \dots, n}(\sigma_{\psi_t}(C_1 \times \dots \times C_{n+k})),$$

where ψ_t is obtained from φ_t by replacing each occurrence of a variable x_i with the index j of the term C_j in which x_i appears. To see that $q(Mod(\mathcal{Z}_k)) = Mod(T)$, since Z_k is a c -table, we can use Theorem 4 and check that, in fact, $\bar{q}(Z_k) = T$ where \bar{q} is the translation of q into the c -tables algebra (see the proof of Theorem 4). Note that we only need the *SPJU* fragment of \mathcal{RA} . □

Example 4. The c -table from Example 2 is definable as $Mod(S) = q(\mathcal{Z}_3)$ where q is the following query with input relation name V of arity 3: $q(V) := \pi_{123}(\{1\} \times \{2\} \times V) \cup \pi_{123}(\sigma_{2=3,4 \neq 2}(\{3\} \times V)) \cup \pi_{512}(\sigma_{3 \neq 1, 3 \neq 4}(\{4\} \times \{5\} \times V))$.

Remark 1. It turns out that the i -databases representable by c -tables are also definable via \mathcal{RA} starting from the absolute zero-information instance, \mathcal{N} . Indeed, it can be shown (Proposition 4) that for each k there exists an \mathcal{RA} query q such that $\mathcal{Z}_k = q(\mathcal{N})$. From there we can apply Theorem 1. The class of incomplete databases $\{\mathcal{I} \mid \exists q \in \mathcal{RA} \text{ s.t. } \mathcal{I} = q(\mathcal{N})\}$ is strictly larger than that representable by c -tables, but it is still countable hence strictly smaller than that of all incomplete databases. Its connections with FO-definability in finite model theory might be interesting to investigate.

Hence, c -tables are in some sense “no more powerful” than the relational algebra. But are they “as powerful”? This justifies the following:

Definition 4. A representation system is **\mathcal{RA} -complete** if it can represent any \mathcal{RA} -definable i -database.

Since Z_k is itself a c -table the following is an immediate corollary of the fundamental result of [20] (see Theorem 4 below). It also states that the converse of Theorem 1 holds.

Theorem 2. c -tables are \mathcal{RA} -complete.

This result is similar in nature to Corollary 3.1 in [18]. However, the exact technical connection, if any, is unclear, since Corollary 3.1 in [18] relies on the certain answers semantics for queries.

We now turn to the kind of completeness considered in [29].

Definition 5. A representation system is **finitely complete** if it can represent any finite i -database.

The finite incompleteness of $?$ -tables, or-set-tables, or-set- $?$ -tables and other systems is discussed in [29] where a finitely complete representation system \mathcal{R}_{prop}^A is also given (we repeat the definition in the Appendix). Is finite completeness a reasonable question for c -tables, v -tables, and Codd tables? In general, for such tables $Mod(T)$ is infinite (all that is needed is a tuple with at least one variable and with an infinitely satisfiable condition). To facilitate comparison with the systems in [29] we define *finite-domain* versions of tables with variables.

Definition 6. A **finite-domain c -table** (v -table, Codd table) consists of a c -table (v -table, Codd table) T together with a finite $dom(x) \subset \mathbb{D}$ for each variable x that occurs in T .

Note that finite-domain Codd tables are equivalent to or-set tables. Indeed, to obtain an or-set table from a Codd table, one can see $dom(x)$ as an or-set and substitute it for x in the table. Conversely, to obtain a Codd table from an or-set

table, one can substitute a fresh variable x for each or-set and define $\text{dom}(x)$ as the contents of the or-set.

In light of this connection, finite-domain v -tables can be thought of as a kind of “correlated” or-set tables. Finite-domain v -tables are strictly more expressive than finite Codd tables. Indeed, every finite Codd table is also a finite v -table. But, the set of instances represented by e.g. the finite v -table $\{(1, x), (x, 1)\}$ where $\text{dom}(x) = \{1, 2\}$ cannot be represented by any finite Codd table. Finite-domain v -tables are themselves finitely incomplete. For example, the i -database $\{(1, 2), (2, 1)\}$ cannot be represented by any finite v -table.

It is easy to see that finite-domain c -tables are finitely complete and hence equivalent to [29]’s $\mathcal{R}_{\text{prop}}^A$ in terms of expressive power. In fact, this is true even for the fragment of finite-domain c -tables which we will call *boolean c -tables*, where the variables take only boolean values and are only allowed to appear in conditions (never as attribute values).

Theorem 3. *Boolean c -tables are finitely complete (hence finite-domain c -tables are also finitely complete).*

Proof. Let $\mathcal{I} = \{I_1, I_2, \dots, I_m\}$ be a finite i -database. Construct a boolean c -table T such that $\text{Mod}(T) = \mathcal{I}$ as follows. Let $\ell := \lceil \lg m \rceil$. For $1 \leq i < m$, put all the tuples from I_i into T with condition φ_i , defined

$$\varphi_i := \bigwedge_j \neg x_j \wedge \bigwedge_k x_k,$$

where the first conjunction is over all $1 \leq j \leq \ell$ such that j th digit in the ℓ -digit binary representation of $i - 1$ is 0, and the second conjunction is over all $1 \leq k \leq \ell$ such that the k th digit in the ℓ -digit binary representation of $i - 1$ is 1. Finally, put all the tuples from I_m into T with condition $\varphi_m \vee \dots \vee \varphi_{2^\ell}$. \square

Although boolean c -tables are complete there are clear advantages to using variables in tuples also, chief among them being *compactness* of representations

Example 5. Consider the finite c -table $\{(x_1, x_2, \dots, x_m : \text{true})\}$ where $\text{dom}(x_1) = \text{dom}(x_2) = \dots = \text{dom}(x_m) = \{1, 2, \dots, n\}$. The equivalent boolean c -table has n^m tuples.

If we additionally restrict boolean c -tables to allow conditions to contain only **true** or a single variable which appears in no other condition, then we obtain a representation system which is equivalent to $?$ -tables.

Since finite c -tables and $\mathcal{R}_{\text{prop}}^A$ are each finitely complete there is an obvious naïve algorithm to translate back and forth between them: list all the instances the one represents, then use the construction from the proof of finite completeness for the other. Finding a more practical “syntactic” algorithm is an interesting open question.

4 Closure Under Relational Operations

Definition 7. *A representation system is **closed** under a query language if for any query q and any table T there is a table T' that represents $q(\text{Mod}(T))$.*

(For notational simplicity we consider only queries with one input relation name, but everything generalizes smoothly to multiple relation names.)

This definition is from [29]. In [2], a *strong* representation system is defined in the same way, with the significant addition that T' should be *computable* from T and q . It is not hard to show, using general recursion-theoretic principles, that there exist representation systems (even ones that only represent finite i -databases) which are closed as above but not strong in the sense of [2]. However, the concrete systems studied so far are either not closed or if they are closed then the proof provides also the algorithm required by the definition of strong systems. Hence, we see no need to insist upon the distinction.

Theorem 4 ([20]). *c-tables, finite-domain c-tables, and boolean c-tables are closed under the relational algebra.*

Proof. (Sketch.) We repeat here the essentials of the proof, including most of the definition of the c -table algebra. For each operation u of the relational algebra [20] defines an operation \bar{u} on c -tables as follows. For projection, we have

$$\bar{\pi}_\ell(T) := \{(t' : \varphi_{t'}) \mid t \in T \text{ s.t. } \pi_\ell(t) = t', \varphi_{t'} = \bigvee \varphi_t\}$$

where ℓ is a list of indexes and the disjunction is over all t in T such that $\pi_\ell(t) = t'$. For selection, we have

$$\bar{\sigma}_c(T) := \{(t : \varphi_t \wedge c(t)) \mid (t, \varphi_t) \in T\}$$

where $c(t)$ denotes the result of evaluating the selection predicate c on the values in t (for a boolean c -table, this will always be **true** or **false**, while for c -tables and finite-domain c -tables, this will be in general a boolean formula on constants and variables). For cross product and union, we have

$$\begin{aligned} T_1 \bar{\times} T_2 &:= \{(t_1 \times t_2 : \varphi_{t_1} \wedge \varphi_{t_2}) \mid t_1 \in T_1, t_2 \in T_2\} \\ T_1 \bar{\cup} T_2 &:= T_1 \cup T_2 \end{aligned}$$

Difference and intersection are handled similarly. By replacing u 's by \bar{u} we translate any relational algebra expression q into a c -table algebra expression \bar{q} and it can be shown that

Lemma 1. *For all valuations ν , $\nu(\bar{q}(T)) = q(\nu(T))$.*

From this, $Mod(\bar{q}(T)) = q(Mod(T))$ follows immediately. □

5 Algebraic Completion

None of the incomplete representation systems we have seen so far is closed under the full relational algebra. Nor are two more representation systems considered in [29], $\mathcal{R}_{\text{sets}}$ and $\mathcal{R}_{\oplus \equiv}$ (we repeat their definitions in the Appendix).

Proposition 1 ([20, 29]). *Codd tables and v-tables are not closed under e.g. selection. Or-set tables and finite v-tables are also not closed under e.g. selection. ?-tables, \mathcal{R}_{sets} , and $\mathcal{R}_{\oplus\equiv}$ are not closed under e.g. join.*

We have seen that “closing” minimal-information one-row Codd tables (see before Definition 4) $\{Z_1, Z_2, \dots\}$, by relational algebra queries yields equivalence with the *c*-tables. In this spirit, we will investigate “how much” of the relational algebra would be needed to complete the other representation systems considered. We call this kind of result *algebraic completion*.

Definition 8. *If (\mathcal{T}, Mod) is a representation system and \mathcal{L} is a query language, then the representation system obtained by closing \mathcal{T} under \mathcal{L} is the set of tables $\{(T, q) \mid T \in \mathcal{T}, q \in \mathcal{L}\}$ with the function $Mod : \mathcal{T} \times \mathcal{L} \rightarrow \mathcal{N}$ defined by $Mod(T, q) := q(Mod(T))$.*

We are now ready to state our results regarding algebraic completion.

Theorem 5 (\mathcal{RA} -Completion).

1. *The representation system obtained by closing Codd tables under SPJU queries is \mathcal{RA} -complete.*
2. *The representation system obtained by closing v-tables under SP queries is \mathcal{RA} -complete.*

Proof. (Sketch.) For each case we show that given a arbitrary *c*-table T one can construct a table S and a query q of the required type such that $\bar{q}(S) = T$. Case 1 is a trivial corollary of Theorem 1. The details for Case 2 are in the Appendix. □

Note that in general there may be a “gap” between the language for which closure fails for a representation system and the language required for completion. For example, Codd tables are not closed under selection, but at the same time closing Codd tables under selection does not yield an \mathcal{RA} -complete representation system. (To see this, consider the incomplete database represented by the *v*-table $\{(x, 1), (x, 2)\}$. Intuitively, selection alone is not powerful enough to yield this incomplete database from a Codd table, as, selection operates on one tuple at a time and cannot correlate two un-correlated tuples.) On the other hand, it is possible that some of the results we present here may be able to be “tightened” to hold for smaller query languages, or else proved to be “tight” already. This is an issue we hope to address in future work.

We give now a set of analogous completion results for the finite case.

Theorem 6 (Finite-Completion).

1. *The representation system obtained by closing or-set-tables under PJ queries is finitely complete.*
2. *The representation system obtained by closing finite v-tables under PJ or S^+P queries is finitely complete.*
3. *The representation system obtained by closing \mathcal{R}_{sets} under PJ or PU queries is finitely complete.*

4. The representation system obtained by closing $\mathcal{R}_{\oplus\equiv}$ under S^+PJ queries is finitely complete.

Proof. (Sketch.) In each case, given an arbitrary finite incomplete database, we construct a table and query of the required type which yields the incomplete database. The details are in the Appendix. \square

Note that there is a gap between the \mathcal{RA} -completion result for Codd tables, which requires $SPJU$ queries, and the finite-completion result for finite Codd tables, which requires only PJ queries. A partial explanation is that proof of the latter result relies essentially on the finiteness of the i -database.

More generally, if a representation system can represent arbitrarily-large i -databases, then closing it under \mathcal{RA} yields a finitely complete representation system, as the following theorem makes precise (see Appendix for proof).

Theorem 7 (General Finite-Completion). *Let \mathcal{T} be a representation system such that for all $n \geq 1$ there exists a table T in \mathcal{T} such that $|\text{Mod}(T)| \geq n$. Then the representation system obtained by closing \mathcal{T} under \mathcal{RA} is finitely-complete.*

Corollary 1. *The representation system obtained by closing $?$ -tables under \mathcal{RA} queries is finitely complete.*

6 Probabilistic Databases and Representation Systems

Finiteness assumption. For the entire discussion of probabilistic database models we will assume that *the domain of values \mathbb{D} is finite*. Infinite domains of values are certainly interesting in practice; for some examples see [22, 33, 29]. Moreover, in the case of incomplete databases we have seen that they allow for interesting distinctions.⁵ However, finite probability spaces are much simpler than infinite ones and we will take advantage of this simplicity. We leave for future investigations the issues related to probabilistic databases over infinite domains.

We wish to model probabilistic information using a probability space whose possible outcomes are all the conventional instances. Recall that for simplicity we assume a schema consisting of just one relation of arity n . The finiteness of \mathbb{D} implies that there are only finitely many instances, $I \subseteq \mathbb{D}^n$.

By **finite probability space** we mean a probability space (see e.g. [11]) $(\Omega, \mathcal{F}, \mathbb{P}[\cdot])$ in which the set of outcomes Ω is *finite* and the σ -field of events \mathcal{F} consists of *all* subsets of Ω . We shall use the equivalent formulation of pairs (Ω, p) where Ω is the finite set of outcomes and where the *outcome probability assignment* $p : \Omega \rightarrow [0, 1]$ satisfies $\sum_{\omega \in \Omega} p(\omega) = 1$. Indeed, we take $\mathbb{P}[A] = \sum_{\omega \in A} p(\omega)$.

⁵ Note however that the results remain true if \mathbb{D} is finite; we just require an infinite supply of *variables*.

Definition 9. A **probabilistic(-information) database** (sometimes called in this paper a **p-database**) is a finite probability space whose outcomes are all the conventional instances, i.e., a pair (\mathcal{N}, p) where $\sum_{I \in \mathcal{N}} p(I) = 1$.

Demanding the direct specification of such probabilistic databases is unrealistic because there are 2^N possible instances, where $N := |\mathbb{D}|^n$, and we would need that many (minus one) probability values. Thus, as in the case of incomplete databases we define **probabilistic representation systems** consisting of “probabilistic tables” (prob. tables for short) and a function *Mod* that associates to each prob. table *T* a probabilistic database *Mod*(*T*). Similarly, we define **completeness** (finite completeness is the only kind we have in our setting).

To define closure under a query language we face the following problem. Given a probabilistic database (\mathcal{N}, p) and a query *q* (with just one input relation name), how do we define the probability assignment for the instances in $q(\mathcal{N})$? It turns out that this is a common construction in probability theory: image spaces.

Definition 10. Let (Ω, p) be a finite probability space and let $f : \Omega \rightarrow \Omega'$ where Ω' is some finite set. The **image** of (Ω, p) under *f* is the finite probability space (Ω', p') where $p'(\omega') := \sum_{f(\omega)=\omega'} p(\omega)$.

Again we consider as query languages the relational algebra and its sublanguages defined by subsets of operations.

Definition 11. A probabilistic representation system is **closed** under a query language if for any query *q* and any prob. table *T* there exists a prob. table *T'* that represents $q(\text{Mod}(T))$, the image space of *Mod*(*T*) under *q*.

7 Probabilistic ?-Tables and Probabilistic Or-Set Tables

Probabilistic ?-tables (*p*-?-tables for short) are commonly used for probabilistic models of databases [34, 15, 16, 9] (they are called “independent tuple representation in [30]). Such tables are the probabilistic counterpart of ?-tables where each “?” is replaced by a probability value. Example 6 below shows such a table. The tuples not explicitly shown are assumed tagged with probability 0. Therefore, we define a *p*-?-table as a mapping that associates to each $t \in \mathbb{D}^n$ a probability value p_t . In order to represent a probabilistic database, papers using this model typically include a statement like “every tuple *t* is in the outcome instance with probability p_t , independently from the other tuples” and then a statement like

$$\mathbb{P}[I] = \left(\prod_{t \in I} p_t \right) \left(\prod_{t \notin I} (1 - p_t) \right).$$

In fact, to give a rigorous semantics, one needs to define the events $E_t \subseteq \mathcal{N}$, $E_t := \{I \mid t \in I\}$ and then to prove the following.

Proposition 2. *There exists a unique probabilistic database such that the events E_t are jointly independent and $\mathbb{P}[E_t] = p_t$.*

⁶ It is easy to check that the $p'(\omega')$'s do actually add up to 1.

This defines p -?-tables as a probabilistic representation system. We shall however provide an equivalent but more perspicuous definition. We shall need here another common construction from probability theory: product spaces.

Definition 12. Let $(\Omega_1, p_1), \dots, (\Omega_n, p_n)$ be finite probability spaces. Their **product** is the space $(\Omega_1 \times \dots \times \Omega_n, p)$ where⁷ $p(\omega_1, \dots, \omega_n) := p_1(\omega_1) \dots p_n(\omega_n)$.

This definition corresponds to the intuition that the n systems or phenomena that are modeled by the spaces $(\Omega_1, p_1), \dots, (\Omega_n, p_n)$ behave without “interfering” with each other. The following formal statements summarize this intuition.

Proposition 3. Consider the product of the spaces $(\Omega_1, p_1), \dots, (\Omega_n, p_n)$. Let $A_1 \subseteq \Omega_1, \dots, A_n \subseteq \Omega_n$.

1. We have $\mathbb{P}[A_1 \times \dots \times A_n] = \mathbb{P}[A_1] \dots \mathbb{P}[A_n]$.
2. The events $A_1 \times \Omega_2 \times \dots \times \Omega_n, \Omega_1 \times A_2 \times \dots \times \Omega_n, \dots, \Omega_1 \times \Omega_2 \times \dots \times A_n$ are jointly independent in the product space.

Turning back to p -?-tables, for each tuple $t \in \mathbb{D}^n$ consider the finite probability space $B_t := (\{\text{true}, \text{false}\}, p)$ where $p(\text{true}) := p_t$ and $p(\text{false}) = 1 - p_t$. Now consider the product space

$$P := \prod_{t \in \mathbb{D}^n} B_t$$

We can think of its set of outcomes (abusing notation, we will call this set P also) as the set of functions from \mathbb{D}^n to $\{\text{true}, \text{false}\}$, in other words, predicates on \mathbb{D}^n . There is an obvious function $f : P \rightarrow \mathcal{N}$ that associates to each predicate the set of tuples it maps to **true**.

All this gives us a p -database, namely the image of P under f . It remains to show that it satisfies the properties in Proposition 2. Indeed, since f is a bijection, this probabilistic database is in fact *isomorphic* to P . In P the events that are in bijection with the E_t 's are the Cartesian product in which there is exactly one component $\{\text{true}\}$ and the rest are $\{\text{true}, \text{false}\}$. The desired properties then follow from Proposition 3.

We define now another simple probabilistic representation system called **probabilistic or-set-tables** (p -or-set-tables for short). These are the probabilistic counterpart of or-set-tables where the attribute values are, instead of or-sets, finite probability spaces whose outcomes are the values in the or-set. p -or-set-tables correspond to a simplified version of the ProbView model presented in [22], in which plain probability values are used instead of confidence intervals.

Example 6. A p -or-set-table S , and a p -?-table T .

$$S := \begin{array}{|c|c|} \hline 1 & \langle 2 : 0.3, 3 : 0.7 \rangle \\ \hline 4 & 5 \\ \hline \langle 6 : 0.5, 7 : 0.5 \rangle & \langle 8 : 0.1, 9 : 0.9 \rangle \\ \hline \end{array} \qquad T := \begin{array}{|c|c|c|} \hline 1 & 2 & 0.4 \\ \hline 3 & 4 & 0.3 \\ \hline 5 & 6 & 1.0 \\ \hline \end{array}$$

⁷ Again, it is easy to check that the outcome probability assignments add up to 1.

A p -or-set-table determines an instance by choosing an outcome in each of the spaces that appear as attribute values, *independently*. Recall that or-set tables are equivalent to finite-domain Codd tables. Similarly, a p -or-set-table corresponds to a Codd table T plus for each variable x in T a finite probability space $\text{dom}(x)$ whose outcomes are in \mathbb{D} . This yields a p -database, again by image space construction, as shown more generally for c -tables next in section 8.

Query answering. The papers [15,34,22] have considered, independently, the problem of calculating the probability of tuples appearing in query answers. This does *not* mean that in general $q(\text{Mod}(T))$ can be represented by another tuple table when T is some p -?-table and $q \in \mathcal{RA}$ (neither does this hold for p -or-set-tables). This follows from Proposition 1. Indeed, if the probabilistic counterpart of an incompleteness representation system \mathcal{T} is closed, then so is \mathcal{T} . Hence the lifting of the results in Proposition 1 and other similar results.

Each of the papers [15, 34, 22] recognizes the problem of query answering and solves it by developing a more general model in which rows contain additional information *similar in spirit* to the conditions that appear in c -tables (in fact [15]’s model is essentially what we call probabilistic boolean c -tables, see next section). We will show that we can actually use a probabilistic counterpart to c -tables themselves together with the algebra on c -tables given in [20] to achieve the same effect.

8 Probabilistic c -Tables

Definition 13. A **probabilistic c -table** (*pc-tables for short*) consists of a c -table T together with a finite probability space $\text{dom}(x)$ (whose outcomes are values in \mathbb{D}) for each variable x that occurs in T .

To get a probabilistic representation system consider the product space

$$V := \prod_{x \in \text{Var}(T)} \text{dom}(x)$$

The outcomes of this space are in fact the *valuations* for the c -table T ! Hence we can define the function $g : V \rightarrow \mathcal{N}, g(\nu) := \nu(T)$ and then define $\text{Mod}(T)$ as the image of V under g .

Similarly, we can talk about boolean pc -tables, pv -tables and probabilistic Codd tables (the latter related to [22], see previous section). Moreover, the p -?-tables correspond to restricted boolean pc -tables, just like ?-tables.

Theorem 8. *Boolean pc -tables are complete (hence pc -tables are also complete).*

Proof. Let I_1, \dots, I_k denote the instances with non-zero probability in an arbitrary probabilistic database, and let p_1, \dots, p_k denote their probabilities. Construct a probabilistic boolean c -table T as follows. For $1 \leq i \leq k - 1$, put the tuples from I_i in T with condition $\neg x_1 \wedge \dots \wedge \neg x_{i-1} \wedge x_i$. Put the tuples from I_k in T with condition $\neg x_1 \wedge \dots \wedge \neg x_{k-1}$. For $1 \leq i \leq k - 1$, set $\mathbb{P}[x_i = \text{true}] := p_i / (1 - \sum_{j=1}^{i-1} p_j)$. It is straightforward to check that this yields a table such that $\mathbb{P}[I_i] = p_i$. □

The previous theorem was independently observed in [30].

Theorem 9. *pc-tables (and boolean pc-tables) are closed under the relational algebra.*

Proof. (Sketch.) For any *pc*-table T and any \mathcal{RA} query q we show that the probability space $q(\text{Mod}(T))$ (the image of $\text{Mod}(T)$ under q) is in fact the same as the space $\text{Mod}(\bar{q}(T))$. The proof of Theorem 4 already shows that the outcomes of the two spaces are the same. The fact that the probabilities assigned to each outcome are the same follows from Lemma 1. \square

The proof of this theorem gives in fact an algorithm for constructing the answer as a p -database itself, represented by a *pc*-table. In particular this will work for the models of [15, 22, 34] or for models we might invent by adding probabilistic information to v -tables or to the representation systems considered in [29]. The interesting result of [9] about the applicability of an “extensional” algorithm to calculating answer tuple probabilities can be seen also as characterizing the conjunctive queries q which for any p -? -table T are such that the c -table $\bar{q}(T)$ is in fact equivalent to some p -? -table.

9 Some Ideas for Further Work

The new results on algebraic completion may not be as tight as they can be. Ideally, we would like to be able show that for each representation system we consider, the fragment of \mathcal{RA} we use is minimal in the sense that closing the representation system under a more restricted fragment does not obtain a complete representation system.

We did not consider c -tables with global conditions [17] nor did we describe the exact connection to logical databases [27, 32]. Even more importantly, we did not consider complexity issues as in [3]. All of the above are important topics for further work, especially the complexity issues and the related issues of *succinctness/compactness* of the table representations.

As we see, in *pc*-tables the probability distribution is on the values taken by the variables that occur in the table. The variables are assumed independent here. This is a lot more flexible (as the example shows) than independent tuples, but still debatable. Consequently, as part of the proposed work, trying to make *pc*-tables even more flexible, we plan to investigate models in which the assumption that the variables take values independently is relaxed by using conditional probability distributions [14].

Space limitations prevent us from giving details, but there is a good reason why the c -table algebra was in essence rediscovered in [15, 22, 34] and to some extent in [28]. The condition that decorates a tuple t in $\bar{q}(T)$ can be seen as the *lineage* [8], a.k.a. the *why-provenance* [6], of the tuple t . We plan to discuss elsewhere the connection between algorithms for computing why-provenance and the c -table algebra.

It would be interesting to connect this work to the extensive literature on *disjunctive databases*, see e.g., [24], and to the work on probabilistic object-oriented databases [12].

Probabilistic modeling is by no means the only way to model uncertainty in information systems. In particular it would be interesting to investigate *possibilistic* models [19] for databases, perhaps following again, as we did here, the parallel with incompleteness.

References

1. S. Abiteboul and O. M. Duschka. Complexity of Answering Queries Using Materialized Views. In *PODS*, pages 254–263, 1998.
2. S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison–Wesley, Reading, MA, 1995.
3. S. Abiteboul, P. Kanellakis, and G. Grahne. On the representation and querying of sets of possible worlds. *Theor. Comput. Sci.*, 78(1):159–187, 1991.
4. M. Arenas, L. E. Bertossi, and J. Chomicki. Answer sets for consistent query answering in inconsistent databases. *TPLP*, 3(4-5):393–424, 2003.
5. D. Barbara, H. Garcia-Molina, and D. Porter. A probabilistic relational data model. In *EDBT*, pages 60–74, New York, NY, USA, 1990.
6. P. Buneman, S. Khanna, and W. C. Tan. Why and Where: A Characterization of Data Provenance. In *ICDT*, pages 316–330, 2001.
7. R. Cavallo and M. Pittarelli. The Theory of Probabilistic Databases. In *VLDB*, pages 71–81, 1987.
8. Y. Cui, J. Widom, and J. L. Wiener. Tracing the lineage of view data in a warehousing environment. *ACM Trans. Database Syst.*, 25(2):179–227, 2000.
9. N. Dalvi and D. Suciu. Efficient Query Evaluation on Probabilistic Databases. In *VLDB*, pages 864–875, 2004.
10. D. Dey and S. Sarkar. A Probabilistic Relational Model and Algebra. *ACM TODS*, 21(3):339–369, 1996.
11. R. Durrett. *Probability: Theory and Examples*. Duxbury Press, 3rd edition, 2004.
12. T. Eiter, J. J. Lu, T. Lukasiewicz, and V. S. Subrahmanian. Probabilistic object bases. *ACM Trans. Database Syst.*, 26(3):264–312, 2001.
13. R. Fagin, P. G. Kolaitis, R. J. Miller, and L. Popa. Data exchange: Semantics and query answering. In *ICDT*, pages 207–224, London, UK, 2003. Springer-Verlag.
14. N. Friedman, L. Getoor, D. Koller, and A. Pfeffer. Learning probabilistic relational models with structural uncertainty. In *Proc. ICML*, 2001.
15. N. Fuhr and T. Rölleke. A Probabilistic Relational Algebra for the Integration of Information Retrieval and Database Systems. *ACM TODS*, 14(1):32–66, 1997.
16. E. Grädel, Y. Gurevich, and C. Hirsch. The Complexity of Query Reliability. In *PODS*, pages 227–234, 1998.
17. G. Grahne. Horn Tables - An Efficient Tool for Handling Incomplete Information in Databases. In *PODS*, pages 75–82. ACM Press, 1989.
18. G. Grahne. *The Problem of Incomplete Information in Relational Databases*, volume 554 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 1991.
19. J. Y. Halpern. *Reasoning About Uncertainty*. MIT Press, Cambridge, MA, 2003.
20. T. Imieliński and W. Lipski, Jr. Incomplete Information in Relational Databases. *J. ACM*, 31(4):761–791, 1984.

21. T. Imieliński, S. A. Naqvi, and K. V. Vadaparty. Incomplete objects — a data model for design and planning applications. In *SIGMOD*, pages 288–297, 1991.
22. L. V. S. Lakshmanan, N. Leone, R. Ross, and V. S. Subrahmanian. ProbView: a Flexible Probabilistic Database System. *ACM TODS*, 22(3):419–469, 1997.
23. L. V. S. Lakshmanan and F. Sadri. Probabilistic deductive databases. In *ILPS*, pages 254–268, Cambridge, MA, USA, 1994. MIT Press.
24. N. Leone, F. Scarcello, and V. S. Subrahmanian. Optimal Models of Disjunctive Logic Programs: Semantics, Complexity, and Computation. *IEEE Trans. Knowl. Data Eng.*, 16(4):487–503, 2004.
25. L. Libkin. *Aspects of Partial Information in Databases*. PhD thesis, University of Pennsylvania, 1994.
26. L. Libkin and L. Wong. Semantic representations and query languages for or-sets. *J. Computer and System Sci.*, 52(1):125–142, 1996.
27. R. Reiter. A sound and sometimes complete query evaluation algorithm for relational databases with null values. *J. ACM*, 33(2):349–370, 1986.
28. F. Sadri. Modeling Uncertainty in Databases. In *ICDE*, pages 122–131. IEEE Computer Society, 1991.
29. A. D. Sarma, O. Benjelloun, A. Halevy, and J. Widom. Working Models for Uncertain Data. In *To appear in ICDE*, April 2006.
30. D. Suciu and N. Dalvi. Foundations of probabilistic answers to queries (tutorial). In *SIGMOD*, pages 963–963, New York, NY, USA, 2005. ACM Press.
31. R. van der Meyden. Logical Approaches to Incomplete Information: A Survey. In J. Chomicki and G. Saake, editors, *Logics for Databases and Information Systems*. Kluwer Academic Publishers, Boston, 1998.
32. M. Y. Vardi. Querying Logical Databases. *JCSS*, 33(2):142–160, 1986.
33. J. Widom. Trio: A System for Integrated Management of Data, Accuracy, and Lineage. In *CIDR*, Jan. 2005.
34. E. Zimányi. Query evaluation in probabilistic databases. *Theoretical Computer Science*, 171(1–2):179–219, 1997.
35. E. Zimányi and A. Pirotte. Imperfect information in relational databases. In *Uncertainty Management in Information Systems*, pages 35–88. Kluwer, 1996.

Appendix

Proposition 4. *There exists a relational query q such that $q(\mathcal{N}) = \mathcal{Z}_n$.*

Proof. Define sub-query q' to be the relational query

$$q'(V) := V - \pi_\ell(\sigma_{\ell \neq r}(V \times V)),$$

where ℓ is short for $1, \dots, n$ and $\ell \neq r$ is short for $1 \neq n + 1 \vee \dots \vee n \neq 2n$. Note that q' yields V if V consists of a single tuple and \emptyset otherwise. Now define q to be the relational query

$$q(V) := q'(V) \cup (\{t\} - \pi_\ell(\{t\} \times q'(V))),$$

where t is a tuple chosen arbitrarily from \mathbb{D}^n . It is clear that $q(\mathcal{N}) = \mathcal{Z}_n$. □

Definition 14. A table in the representation system \mathcal{R}_{sets} is a multiset of sets of tuples, or blocks, each such block optionally labeled with a ‘?’. If T is an \mathcal{R}_{sets} table, then $Mod(T)$ is the set of instances obtained by choosing one tuple from each block not labeled with a ‘?’, and at most one tuple from each block labeled with a ‘?’.

Definition 15. A table in the representation system $\mathcal{R}_{\oplus\equiv}$ is a multiset of tuples $\{t_1, \dots, t_m\}$ and a conjunction of logical assertions of the form $i \oplus j$ (meaning t_i or t_j must be present in an instance, but not both) or $i \equiv j$ (meaning t_i is present in an instance iff t_j is present in the instance). If T is an $\mathcal{R}_{\oplus\equiv}$ table then $Mod(T)$ consists of all subsets of the tuples satisfying the conjunction of assertions.

Definition 16. A table in the representation system \mathcal{R}_{prop}^A is a multiset of or-set tuples $\{t_1, \dots, t_m\}$ and a boolean formula on the variables $\{t_1, \dots, t_m\}$. If T is an \mathcal{R}_{prop}^A table then $Mod(T)$ consists of all subsets of the tuples satisfying the boolean assertion, where the variable t_i has value **true** iff the tuple t_i is present in the subset.

Theorem 5 (RA-Completion).

1. The representation system obtained by closing Codd tables under SPJU queries is RA-complete.
2. The representation system obtained by closing v-tables under SP queries is RA-complete.

Proof. In each case we show that given an arbitrary c-table T , one can construct a table S and a query q such that $\bar{q}(S) = T$.

1. Trivial corollary of Theorem 1.
2. Let k be the arity of T . Let $\{t_1, \dots, t_m\}$ be an enumeration of the tuples of T , and let $\{x_1, \dots, x_n\}$ be an enumeration of the variables which appear in T . Construct a v-table S with arity $k + n + 1$ as follows. For every tuple t_i in T , put exactly one tuple t'_i in S , where t'_i agrees with t_i on the first k columns, the $k + 1$ st column contains the constant i , and the last m columns contain the variables x_1, \dots, x_m . Now let q be the SP query defined

$$q := \pi_{1, \dots, k}(\sigma_{\bigvee_{i=1}^m k+1=i' \wedge \psi_i}(S))$$

where ψ_i is obtained from the condition φ_{t_i} of tuple t_i by replacing variable names with their corresponding indexes in S . □

Theorem 6 (Finite Completion).

1. The representation system obtained by closing or-set-tables under PJ queries is finitely complete.
2. The representation system obtained by closing finite v-tables under PJ or S^+P queries is finitely complete.

3. The representation system obtained by closing \mathcal{R}_{sets} under PJ or PU queries is finitely complete.
4. The representation system obtained by closing $\mathcal{R}_{\oplus\equiv}$ under S^+PJ queries is finitely complete.

Proof. Fix an arbitrary finite incomplete database $\mathcal{I} = \{I_1, \dots, I_n\}$ of arity k . It suffices to show in each case that one can construct a table T in the given representation system and a query q in the given language such that $q(\text{Mod}(T)) = \mathcal{I}$.

1. We construct a pair of or-set-tables S and T as follows. (They can be combined together into a single table, but we keep them separate to simplify the presentation.) For each instance I_i in \mathcal{I} , we put all the tuples of I_i in S , appending an extra column containing value i . Let T be the or-set-table of arity 1 containing a single tuple whose single value is the or-set $\langle 1, 2, \dots, n \rangle$. Now let q be the S^+PJ query defined:

$$q := \pi_{1, \dots, k} \sigma_{k+1=k+2}(S \times T).$$

2. Completion for PJ follows from Case 1 and the fact that finite v -tables are strictly more expressive than or-set tables. For S^+P , take the finite v -table representing the cross product of S and T in the construction from Case 1, and let q be the obvious S^+P query.
3. Completion for PJ follows from Case 1 and the fact (shown in [29]) that or-set-tables are strictly less expressive than \mathcal{R}_{sets} . Thus we just need show the construction for PU . We construct an \mathcal{R}_{sets} table T as follows. Let m be the cardinality of the largest instance in \mathcal{I} . Then T will have arity km and will consist of a single block of tuples. For every instance I_i in \mathcal{I} , we put one tuple in T which has every tuple from I_i arranged in a row. (If the cardinality of I_i is less than m , we pad the remainder with arbitrary tuples from I_i .) Now let q be the PU query defined as follows:

$$q := \bigcup_{i=0}^{m-1} \pi_{ki, \dots, ki+k-1}(T)$$

4. We construct a pair of $\mathcal{R}_{\oplus\equiv}$ -tables S and T as follows. (S can be encoded as a special tuple in T , but we keep it separate to simplify the presentation.) Let $m = \lceil \lg n \rceil$. T is constructed as in Case 2. S is a binary table containing, for each i , $1 \leq i \leq m$, a pair of tuples $(0, i)$ and $(1, i)$ with an exclusive-or constraint between them. Let sub-query q' be defined

$$q' := \prod_{i=1}^m \pi_1(\sigma_{2=i}(S))$$

The S^+PJ query q is defined as in Case 2, but using this definition of q' . \square

Theorem 7 (General Finite Completion). *Let \mathcal{T} be a representation system such that for all $n \geq 1$ there exists a table T in \mathcal{T} such that $|\text{Mod}(T)| \geq n$. Then the representation system obtained by closing \mathcal{T} under \mathcal{RA} is finitely-complete.*

Proof. Let \mathcal{T} be a representation system such that for all $n \geq 1$ there is a table T in \mathcal{T} such that $|\text{Mod}(T)| \geq n$. Let $\mathcal{I} = \{I_1, \dots, I_k\}$ be an arbitrary non-empty finite set of instances of arity m . Let T be a table in \mathcal{T} such that $\text{Mod}(T) = \{J_1, \dots, J_\ell\}$, with $\ell \geq k$. Define \mathcal{RA} query q to be

$$q(V) := \bigcup_{1 \leq i \leq k-1} I_i \times q_i(V) \cup \bigcup_{k \leq i \leq \ell} I_i \times q_i(V),$$

where I_i is the query which constructs instance I_i and $q_i(V)$ is the boolean query which returns true iff V is identical to I_i (which can be done in \mathcal{RA}). Then $q(\text{Mod}(T)) = \mathcal{I}$. □