

Conditioning Probabilistic Databases

Christoph Koch
Department of Computer Science
Cornell University
Ithaca, NY 14853, USA
koch@cs.cornell.edu

Dan Olteanu
Computing Laboratory
Oxford University
Oxford, OX1 3QD, UK
dan.olteanu@comlab.ox.ac.uk

ABSTRACT

Past research on probabilistic databases has studied the problem of answering queries on a static database. Application scenarios of probabilistic databases however often involve the conditioning of a database using additional information in the form of new evidence. The conditioning problem is thus to transform a probabilistic database of priors into a posterior probabilistic database which is materialized for subsequent query processing or further refinement. It turns out that the conditioning problem is closely related to the problem of computing exact tuple confidence values.

It is known that exact confidence computation is an NP-hard problem. This has led researchers to consider approximation techniques for confidence computation. However, neither conditioning nor exact confidence computation can be solved using such techniques. In this paper we present efficient techniques for both problems. We study several problem decomposition methods and heuristics that are based on the most successful search techniques from constraint satisfaction, such as the Davis-Putnam algorithm. We complement this with a thorough experimental evaluation of the algorithms proposed. Our experiments show that our exact algorithms scale well to realistic database sizes and can in some scenarios compete with the most efficient previous approximation algorithms.

1. INTRODUCTION

Queries on probabilistic databases have numerous applications at the interface of databases and information retrieval [13], data cleaning [4], sensor data, tracking moving objects, crime fighting [5], and computational science [9].

A core operation of queries on probabilistic databases is the computation of confidence values of tuples in the result of a query. In short, the confidence in a tuple t being in the result of a query on a probabilistic database is the combined probability weight of all possible worlds in which t is in the result of the query.

By extending the power of query languages for probabilis-

tic databases, new applications beyond the mere retrieval of tuples and their confidence become possible. An essential operation that allows for new applications is *conditioning*, the operation of removing possible worlds which do not satisfy a given condition from a probabilistic database. Subsequent query operations will apply to the reduced database, and a confidence computation will return *conditional probabilities* in the Bayesian sense with respect to the original database. Computing conditioned probabilistic databases has natural and important applications in virtually all areas in which probabilistic databases are useful. For example, in data cleaning, it is only natural to start with an uncertain database and then clean it – reduce uncertainty – by adding constraints or additional information. More generally, conditioning allows us to start with a database of prior probabilities, to add in some *evidence*, and take it to a posterior probabilistic database that takes the evidence into account.

Consider the example of a probabilistic database of social security numbers (SSN) and names of individuals extracted from paper forms using OCR software. If a symbol or word cannot be clearly identified, this software will offer a number of weighted alternatives. The database

R	SSN	NAME
	{ 1 (p=.2) 7 (p=.8) }	John
	{ 4 (p=.3) 7 (p=.7) }	Bill

represents four possible worlds (shown in Figure 1), modelling that John has either SSN 1 or 7, with probability .2 and .8 (the paper form may contain a hand-written symbol that can either be read as a European “1” or an American “7”), respectively, and Bill has either SSN 4 or 7, with probability .3 and .7, respectively. We assume independence between John’s and Bill’s alternatives, thus the world in which John has SSN 1 and Bill has SSN 7 has probability $.2 \cdot .7 = .14$.

If A_x denotes the event that Bill has SSN x , then $P(A_4) = .3$ and $P(A_7) = .7$. We can compute these probabilities in a probabilistic database by asking for the confidence values of the tuples in the result of the query

```
select SSN, conf(SSN) from R where NAME = 'Bill';
```

which will result in the table

Q	SSN	CONF
	4	.3
	7	.7

Now suppose we want to use the additional knowledge that social security numbers are unique. We can express this

R^1	SSN	NAME	R^2	SSN	NAME
	1	John		7	John
	4	Bill		4	Bill
P = .06			P = .24		
R^3	SSN	NAME	R^4	SSN	NAME
	1	John		7	John
	7	Bill		7	Bill
P = .14			P = .56		

Figure 1: The four worlds of the input database.

using a functional dependency $SSN \rightarrow NAME$. Asserting this constraint, or conditioning the probabilistic database using the constraint, means to eliminate all those worlds in which the functional dependency does not hold.

Let B be the event that the functional dependency holds. Conceptually, the database conditioned with B is obtained by removing world R^4 (in which John and Bill have the same SSN) and renormalizing the probabilities of the remaining worlds to have them again sum up to 1, in this case by dividing by $.06 + .24 + .14 = .44$. We will think of conditioning as an operation $\text{assert}[B]$ that reduces uncertainty by declaring worlds in which B does not hold impossible.

Computing tuple confidences for the above query on the original database will give us, for each possible SSN value x for Bill, the probabilities $P(A_x)$, while on the database conditioned with B it will give a table of social security numbers x and conditional probabilities $P(A_x | B)$. For example, the conditional probability of Bill having SSN 4 given that social security numbers are unique is

$$P(A_4 | B) = \frac{P(A_4 \wedge B)}{P(B)} = \frac{.3}{.44} \approx .68.$$

Using this definition, we could alternatively have computed the conditional probabilities by combining the results of two confidence computations,

```
select SSN, P1/P2
from (select SSN, conf(SSN) P1 from R, B
     where NAME = 'Bill'),
     (select conf() P2 from B);
```

where B is a Boolean query that is true if the functional dependency holds on R .

Unfortunately, both conditioning and confidence computation are NP-hard problems. Nevertheless, their study is justified by their obvious relevance and applications. While conditioning has not been previously studied in the context of probabilistic databases, previous work on confidence computation has aimed at cases that admit polynomial-time query evaluation and at approximating confidence values [9].

Previous work often assumes that confidence values are computed at the end of a query, closing the possible worlds semantics of the probabilistic database and returning a complete, nonprobabilistic relation of tuples with numerical confidence values that can be used for decision making. In such a context, techniques that return a reasonable approximation of confidence values may be acceptable.

In other scenarios we do not want to accept approximate confidence values because errors made while computing these estimates aggregate and grow, causing users to make wrong decisions based on the query results. This is particularly true in compositional query languages for probabilistic databases, where confidence values computed in a

subquery form part of an intermediate result that can be accessed and used for filtering the data in subsequent query operations [19].

Similar issues arise when confidence values can be inserted into the probabilistic database through updates and may be used in subsequent queries. For example, data cleaning is a scenario where we, on one hand, want to materialize the result of a data transformation in the database once and for all (rather than having to redo the cleaning steps every time a query is asked) and on the other hand do not want to store incorrect probabilities that may affect a very large number of subsequent queries. Here we need techniques for conditioning and exactly computing confidence values.

Exact confidence computation is particularly important in queries in which confidence values are used in comparison predicates. For an example, let us add a third person, Fred, to the database whose SSN is either 1 or 4, with equal probability. If we again condition using the functional dependency $SSN \rightarrow NAME$, we have only two possible worlds, one in which John, Bill, and Fred have social security numbers 1, 7, and 4, respectively, and one in which their SSN are 7, 4, and 1. If we now ask for the social security numbers that are in the database for certain,

```
select SSN from R where conf(SSN) = 1;
```

we should get three tuples in the result. Monte Carlo simulation based approximation algorithms will do very badly on such queries. Confidence approximation using a Karp-Luby-style algorithm [17, 9, 21] will independently underestimate each tuple's confidence with probability $\approx .5$. Thus the probability that at least one tuple is missing from the result of such a query is very high (see also [19]).

In this paper, we develop efficient algorithms for computing exact confidences and for conditioning probabilistic databases. The detailed contributions are as follows.

- In most previous models of probabilistic databases over finite world-sets, computing tuple confidence values essentially means the weighted counting of solutions to constraint structures closely related to disjunctive normal form formulas. Our notion of such structures are the world-set descriptor sets, or *ws-sets* for short. We formally introduce a probabilistic database model that is known to cleanly and directly generalize many previously considered probabilistic database models (cf. [3]) including, among others, various forms of tuple-independence models [9, 2], ULDBs [5], product decomposition [4], and c-table-based models [3]. We use this framework to study exact confidence computation and conditioning. The results obtained are thus of immediate relevance to all these models.
- We study properties of *ws-sets* that are essential to relational algebra query evaluation and to the design of algorithms for the two main problems of the paper.
- We exhibit the fundamental, close relationship between the two problems.
- We develop *ws-trees*, which capture notions of structural decomposition of *ws-sets* based on probabilistic independence and world-set disjointness. Once a *ws-tree* has been obtained for a given *ws-set*, both exact confidence computation and conditioning are feasible in linear time. The main problem is thus to efficiently find small *ws-tree* decompositions.

W	Var	Dom	P	U_R	WSD	SSN	NAME
	j	1	.2		$\{j \mapsto 1\}$	1	John
	j	7	.8		$\{j \mapsto 7\}$	7	John
	b	4	.3		$\{b \mapsto 4\}$	4	Bill
	b	7	.7		$\{b \mapsto 7\}$	7	Bill

Figure 2: Probabilistic database with ws-descriptors made explicit and defined by world-table W .

- To this end, we develop a decomposition procedure motivated by the Davis-Putnam (DP) procedure for checking Propositional Satisfiability [12]. DP, while many decades old, is still the basis of the best exact solution techniques for the NP-complete Satisfiability problem. We introduce two decomposition rules, variable elimination (the main rule of DP) and a new independence decomposition rule, and develop heuristics for choosing among the rules.
- We develop a database conditioning algorithm based on ws-tree decompositions and prove its correctness.
- We study ws-set simplification and elimination techniques that can be either used as an alternative to the DP-based procedure or combined with it.
- We provide a thorough experimental evaluation of the algorithms presented in this paper. We also experimentally compare our exact techniques for confidence computation with approximation based on Monte Carlo simulation.

The structure of the paper follows the list of contributions.

2. PROBABILISTIC DATABASES

We define *sets of possible worlds* following U-relational databases [3]. Consider a finite set of independent random variables ranging over finite domains. Probability distributions over the possible worlds are defined by assigning a probability $P(\{x \mapsto i\})$ to each assignment of a variable x to a constant of its domain, $i \in \text{Dom}_x$, such that the probabilities of all assignments of a given variable sum up to one. We represent the set of variables, their domains, and probability distributions relationally by a *world-table* W consisting of all triples (x, i, p) of variables x , values i in the domain of x , and the associated probabilities $p = P(\{x \mapsto i\})$.

A *world-set descriptor* is a set of assignments $x \mapsto i$ with $i \in \text{Dom}_x$ that is functional, i.e. a partial function from variables to domain values. If such a world-set descriptor d is a *total* function, then it identifies a possible world. Otherwise, it denotes all those possible worlds $\omega(d)$ identified by total functions f that can be obtained by extension of d . (That is, for all x on which d is defined, $d(x) = f(x)$.) Because of the independence of the variables, the aggregate probability of these worlds is

$$P(d) = \prod_{\{x \mapsto i\} \subseteq d} P(\{x \mapsto i\}).$$

If $d = \emptyset$, then d denotes the set of all possible worlds.

We say that two ws-descriptors d_1 and d_2 are *consistent* iff their union (as sets of assignments) is functional.

A *ws-set* is a set of ws-descriptors S and represents the world-set computed as the union of the world-sets represented by the ws-descriptors in the set. We define the se-

mantics of ws-sets using the (herewith overloaded) function ω extended to ws-sets, $\omega(S) := \bigcup_{d \in S} (\omega(d))$.

A *U-relation* over schema Σ and world-table W is a set of tuples over Σ , where we associate to each tuple a ws-descriptor over W . A *probabilistic database* over schema $\{\Sigma_1, \dots, \Sigma_n\}$ and world-table W is a set of n U-relations, each over one schema Σ_i and W . A probabilistic database represents a set of databases, one database for each possible world defined by W . To obtain a possible world in the represented set, we first choose a total valuation f over W . We then process each probabilistic relation R_i tuple by tuple. If f extends the ws-descriptor d of a tuple t , then t is in the relation R_i of that database.

EXAMPLE 2.1. Consider again the probabilistic database of social security numbers and names given in Figure 1. Its representation in our formalism is given in Figure 2. The world-table W of Figure 2 defines two variables j and b modeling the social security numbers of John and Bill, with domains $\{1, 7\}$ and $\{4, 7\}$ respectively. The probability of the world defined by $f = \{j \mapsto 7, b \mapsto 7\}$ is $.8 \cdot .7 = .56$. The total valuation f extends the ws-descriptors of the second and fourth tuple of relation U_R , thus the relation R in world f is $\{(7, \text{John}), (7, \text{Bill})\}$. \square

REMARK 2.2. Leaving aside the probability distributions of the variables which are represented by the W table, U-relations are essentially restricted c-tables [16] in which the global condition is “true”, variables must not occur in the tuples, and each local condition must be a conjunction of conditions of the form $x = v$ where x is a variable and v is a constant. Nevertheless, it is known that U-relations are a complete representation system for probabilistic databases over nonempty finite sets of possible worlds.

U-relations can be used to represent attribute-level uncertainty using vertically decomposed relations. For details on this, we refer to [3]. All results in this paper work in the context of attribute-level uncertainty.

The efficient execution of the operations of positive relational algebra on such databases was described in that paper as well. Briefly, if U-relations U_R and U_S represent relations R and S , then selections $\sigma_\phi R$ and projections $\pi_{\bar{A}} R$ simply translate into $\sigma_\phi U_R$ and $\pi_{\bar{A}} U_R$, respectively. Joins $R \bowtie_\phi S$ translate into $U_R \bowtie_{\phi \wedge \psi} U_S$ where ψ is the condition that the ws-descriptors of the two tuples compared are consistent with each other (i.e., have a common extension into a total valuation). The set operations easily follow from the analogous operations on ws-sets that will be described below, in Section 3.2. \square

EXAMPLE 2.3. The functional dependency $\text{SSN} \rightarrow \text{NAME}$ on the probabilistic database of Figure 2 can be expressed as a boolean relational algebra query as the complement of $\pi_\emptyset(R \bowtie_\phi R)$ where $\phi := (1.\text{SSN} = 2.\text{SSN} \wedge 1.\text{NAME} \neq 2.\text{NAME})$. We turn this into the query

$$\pi_{\text{WSD}}(U_R \bowtie_{\phi \wedge 1.\text{WSD consistent with } 2.\text{WSD}} U_R).$$

over our representation, which results in the ws-set $\{j \mapsto 7, b \mapsto 7\}$. The complement of this with the world-set given by the W relation, $\{\{j \mapsto 1\}, \{j \mapsto 7\}, \{b \mapsto 4\}, \{b \mapsto 7\}\}$, is $\{\{j \mapsto 1\}, \{j \mapsto 7, b \mapsto 4\}\}$. (Note that this is just one among a set of equivalent solutions.) \square

3. PROPERTIES OF WS-DESCRIPTORS

In this section we investigate properties of ws-descriptors and show how they can be used to efficiently implement various set operations on world-sets without having to enumerate the worlds. This is important, for such sets can be extremely large in practice: [4, 3] report on experiments with 10^{10^6} worlds.

3.1 Mutex, Independence, and Containment

Two ws-descriptors d_1 and d_2 are (1) *mutually exclusive* (mutex for short) if they represent disjoint world-sets, i.e., $\omega(d_1) \cap \omega(d_2) = \emptyset$, and (2) *independent* if there is no valuation of the variables in one of the ws-descriptors that restricts the set of possible valuations of the variables in the other ws-descriptor (that is, d_1 and d_2 are defined only on disjoint sets of variables). A ws-descriptor d_1 is *contained* in d_2 if the world-set of d_1 is contained in the world-set of d_2 , i.e., $\omega(d_1) \subseteq \omega(d_2)$. *Equivalence* is mutual containment.

Although ws-descriptors represent very succinctly possibly very large world-sets, all aforementioned properties can be efficiently checked at the syntactical level: d_1 and d_2 , where all variables with singleton domains are eliminated, are (1) mutex if there is a variable with a different assignment in each of them, and (2) independent if they have no variables in common; d_1 is contained in d_2 if d_1 extends d_2 .

EXAMPLE 3.1. Consider the world-table of Figure 2 and the ws-descriptors $d_1 = \{j \mapsto 1\}$, $d_2 = \{j \mapsto 7\}$, $d_3 = \{j \mapsto 1, b \mapsto 4\}$, and $d_4 = \{b \mapsto 4\}$. Then, the pairs (d_1, d_2) and (d_2, d_3) are mutex, d_3 is contained in d_1 , and the pairs (d_1, d_4) and (d_2, d_4) are independent. \square

We also consider the mutex, independence, and equivalence properties for ws-sets. Two ws-sets S_1 and S_2 are mutex (independent) iff d_1 and d_2 are mutex (independent) for any $d_1 \in S_1$ and $d_2 \in S_2$. Two ws-sets are equivalent if they represent the same world-set.

EXAMPLE 3.2. We continue Example 3.1. The ws-set $\{d_1\}$ is mutex with $\{d_2\}$. $\{d_1, d_2\}$ is independent from $\{d_4\}$. At a first glance, it looks like $\{d_1, d_2\}$ and $\{d_3, d_4\}$ are neither mutex nor independent, because d_1 and d_3 overlap. However, we note that $d_3 \subseteq d_4$ and then $\omega(\{d_3, d_4\}) = \omega(\{d_4\})$ and $\{d_4\}$ is independent from $\{d_1, d_2\}$. \square

3.2 Set Operations on ws-sets

Various relevant computation tasks, ranging from decision procedures like tuple possibility [1] to confidence computation of answer tuples, and conditioning of probabilistic databases, require symbolic manipulations of ws-sets. For example, checking whether two tuples of a probabilistic relation can co-occur in some worlds can be done by intersecting their ws-descriptors; both tuples co-occur in the worlds defined by the intersection of the corresponding world-sets. We next define set operations on ws-sets.

- *Intersection.* $\text{Intersect}(S_1, S_2) := \{d_1 \cap d_2 \mid d_1 \in S_1, d_2 \in S_2, d_1 \text{ is consistent with } d_2\}$.
- *Union.* $\text{Union}(S_1, S_2) := S_1 \cup S_2$.
- *Difference.* The definition is inductive, starting with singleton ws-sets. If ws-descriptors d_1 and d_2 are in-

consistent, $\text{Diff}(\{d_1\}, \{d_2\}) := \{d_1\}$. Otherwise,

$$\begin{aligned} \text{Diff}(\{d_1\}, \{d_2\}) := \\ \{d_1 \cup \{x_1 \mapsto w_1, \dots, x_{i-1} \mapsto w_{i-1}, x_i \mapsto w'_i\} \mid \\ d_2 - d_1 = \{x_1 \mapsto w_1, \dots, x_k \mapsto w_k\}, \\ 1 \leq i \leq k, w'_i \in \text{dom}_{x_i}, w_i \neq w'_i\}. \end{aligned}$$

$$\text{Diff}(\{d_1\}, S \cup \{d_2\}) := \text{Diff}(\text{Diff}(\{d_1\}, S), \{d_2\}).$$

$$\text{Diff}(\{d_1, \dots, d_n\}, S) := \bigcup_{1 \leq i \leq n} \text{Diff}(\{d_i\}, S).$$

EXAMPLE 3.3. Consider $d_1 = \{j \mapsto 1\}$, $d_2 = \{j \mapsto 7\}$, and $d_3 = \{j \mapsto 1, b \mapsto 4\}$. Then, $\text{Intersect}(\{d_1\}, \{d_2\}) = \text{Intersect}(\{d_2\}, \{d_3\}) = \emptyset$ because d_2 is inconsistent with d_1 and d_3 . $\text{Intersect}(\{d_1\}, \{d_3\}) = \{d_3\}$, because d_3 is contained in d_1 . $\text{Diff}(\{d_2\}, \{d_1\}) = \text{Diff}(\{d_2\}, \{d_3\}) = \{d_2\}$ because d_2 is mutex with d_1 and d_3 . $\text{Diff}(\{d_1\}, \{d_3\}) = \{\{j \mapsto 1, b \mapsto 7\}\}$. $\text{Diff}(\{d_3\}, \{d_1\}) = \{d_3\}$, because d_3 and d_1 are inconsistent. \square

PROPOSITION 3.4. *The above definitions of set operations on ws-sets are correct:*

1. $\omega(\text{Union}(S_1, S_2)) = \omega(S_1) \cup \omega(S_2)$.
2. $\omega(\text{Intersect}(S_1, S_2)) = \omega(S_1) \cap \omega(S_2)$.
3. $\omega(\text{Diff}(S_1, S_2)) = \omega(S_1) - \omega(S_2)$.

The ws-descriptors in $\text{Diff}(S_1, S_2)$ are pairwise mutex.

4. WORLD-SET TREES

The ws-sets have important properties, like succinctness, closure under set operations, and natural relational encoding, and [3] employed them to achieve the purely relational processing of positive relational algebra on U-relational databases. When it comes to the manipulation of probabilities of query answers or of worlds violating given constraints, however, ws-sets are in most cases inadequate. This is because ws-descriptors in a ws-set may represent non-disjoint world-sets, and for most manipulations of probabilities a substantial computational effort is needed to identify common world-subsets across possibly many ws-descriptors.

We next introduce a new compact representation of world-sets, called *world-set tree* representation, or ws-tree for short, that makes the structure in the ws-sets explicit. This representation formalism allows for efficient exact probability computation and conditioning and has strong connections to *knowledge compilation*, as it is used in system modelling and verification [11]. There, too, various kinds of decision diagrams, like binary decision diagrams (BDDs) [7], are employed for the efficient manipulation of propositional formulas.

DEFINITION 4.1. Given a world-table W , a ws-tree over W is a tree with inner nodes \otimes and \oplus , leaves holding the ws-descriptor \emptyset , and edges annotated with weighted variable assignments consistent with W . The following constraints hold for a ws-tree:

- A variable defined in W occurs at most once on each root-to-leaf path.

W	V	D	P
x	1	.1	
x	2	.4	
x	3	.5	
y	1	.2	
y	2	.8	
z	1	.4	
z	2	.6	
u	1	.7	
u	2	.3	
v	1	.5	
v	2	.5	




Figure 3: World-set table W , a ws-tree R over W , and an equivalent ws-set S .

- Each of its \oplus -nodes is associated with a variable v such that each outgoing edge is annotated with a different assignment of v .
- The sets of variables occurring in the subtrees rooted at the children of any \otimes -node are disjoint. \square

We define the semantics of ws-trees in strict analogy to that of ws-sets based on the observation that the set of edge annotations on each root-to-leaf path in a ws-tree represents a ws-descriptor. The world-set represented by a ws-tree is precisely represented by the ws-set consisting of the annotation sets of all root-to-leaf paths. The inner nodes have a special semantics: the children of a \otimes -node use disjoint variable sets and are thus independent, and the children of a \oplus -node follow branches with different assignments of the same variable and are thus mutually exclusive.

EXAMPLE 4.2. Figure 3 shows a ws-tree and the ws-set consisting of all its root-to-leaf paths. \square

4.1 Constructing world-set trees

The key idea underlying our translation of ws-sets into ws-trees is a divide-and-conquer approach that exploits the relationships between ws-descriptors, like independence and variable sharing.

Figure 4 gives our translation algorithm. We proceed recursively by partitioning the ws-sets into independent disjoint partitions (when possible) or into (possibly overlapping) partitions that are consistent with different assignments of a variable. In the case of independent partitioning, we create \otimes -nodes whose children are the translations of the independent partitions. In the second case, we simplify the problem by eliminating a variable: we choose a variable x and create an \oplus -node whose outgoing edges are annotated with different assignments $x \mapsto i$ of x and whose children are translations of the subsets of the ws-set consisting of ws-descriptors consistent with $x \mapsto i$ ¹. If at any recursion step the input ws-set contains the nullary ws-descriptor, which by definition represents the whole world-set, then we stop from recursion and create a ws-tree leaf \emptyset . This can happen after several variable elimination steps that reduced some of the input ws-descriptors to \emptyset .

¹Our translation abstracts out implementation details. For instance, for those assignments of x that do not occur in S we have $T \cup S_{x \mapsto i} = T$ and can translate T only once.

```

ComputeTree (WS-Set S) returns WS-Tree
if ( $S = \emptyset$ ) then return  $\perp$ 
else if ( $\emptyset \in S$ ) //S contains a universal ws-desc
then return  $\emptyset$ 
else choose one of the following:
(independent partitioning)
if there are non-empty and independent ws-sets
 $S_1, \dots, S_{|I|}$  such that  $S = S_1 \cup \dots \cup S_{|I|}$ 
then return  $\bigotimes_{i \in I} (\text{ComputeTree}(S_i))$ 
(variable elimination)
choose a variable  $x$  in  $S$ ;
 $T := \{d \mid d \in S, \exists i \in \text{dom}_x : \{x \mapsto i\} \subseteq d\}$ ;
 $\forall i \in \text{dom}_x : S_{x \mapsto i} := \{\{y_1 \mapsto j_1, \dots, y_m \mapsto j_m\} \mid$ 
 $\{x \mapsto i, y_1 \mapsto j_1, \dots, y_m \mapsto j_m\} \in S\}$ ;
return  $\bigoplus_{i \in \text{dom}_x} (x \mapsto i : \text{ComputeTree}(S_{x \mapsto i} \cup T))$ 

```

Figure 4: Translating ws-sets into ws-trees.

EXAMPLE 4.3. We show how to translate the ws-set S into the ws-tree R (Figure 3). We first partition S into two (minimally) independent ws-sets S_1 and S_2 : S_1 consists of the first three ws-descriptors of S , and S_2 consists of the remaining two. For S_1 , we can eliminate any of the variables x , y , or z . Consider we choose x and create two branches for $x \mapsto 1$ and $x \mapsto 2$ respectively (there is no ws-descriptor consistent with $x \mapsto 3$). For the first branch, we stop with the ws-set $\{\emptyset\}$, whereas for the second branch we continue with the ws-set $\{\{y \mapsto 1\}, \{z \mapsto 1\}\}$. The latter ws-set can be partitioned into independent subsets in the *context* of the assignment $x \mapsto 2$. We proceed similarly for S_2 and choose to eliminate variable u . We create an \oplus -node with outgoing edges for assignments $u \mapsto 1$ and $u \mapsto 2$ respectively. We are left in the former case with the ws-set $\{\{v \mapsto 1\}\}$ and in the latter case with $\{\emptyset\}$.

Different variable choices can lead to different ws-trees. This is the so-called *variable ordering problem* that applies to the construction of binary decision diagrams. Later in




Figure 5: A ws-tree equivalent to R of Figure 3.

this section we discuss heuristics for variable orderings. \square

THEOREM 4.4. *Given a ws-set S , $\text{ComputeTree}(S)$ and S represent the same world-set.*

Our translation can yield ws-trees of exponential size (similar to BDDs). This rather high worst-case complexity needs to be paid for efficient exact probability computation and conditioning. It is known that counting models of propositional formulas and exact probability computation are $\#\text{P}$ -hard problems [9]. This complexity result does not preclude, however, BDDs from being very successful in practice. We expect the same for ws-trees. The key observation for a good behaviour in practice is that we should partition ws-sets into independent subsets whenever possible and we should carefully choose a good ordering for variable eliminations. Both methods greatly influence the size of the ws-trees and the translation time, as shown in the next example.

EXAMPLE 4.5. Consider again the ws-set S of Figure 3 and a different ordering for variable eliminations that leads to the ws-tree of Figure 5. We shortly discuss the construction of this ws-tree. Assume we choose to eliminate the variable y and obtain the ws-sets

$$S_{y \mapsto 2} = \{\{x \mapsto 1\}, \{x \mapsto 2, z \mapsto 1\}, \{u \mapsto 1, v \mapsto 1\}, \{u \mapsto 2\}\}$$

$$S_{y \mapsto 1} = S_{y \mapsto 2} \cup \{\{x \mapsto 2\}\}$$

In contrast to the computation of the ws-tree R of Figure 3, our variable choice creates intermediary ws-sets that overlap at large, which ultimately leads to a large increase in the size of the ws-tree. This bad choice need not necessarily lead to *redundant* computation, which we could easily detect. In fact, the only major savings in case we detect and eliminate redundancy here are the subtrees α and β , which still leave a graph larger than R . \square

4.2 Heuristics

We next study heuristics for variable elimination and independent partitioning that are compared experimentally in Section 7. We devise a simple cost estimate, which we use to decide at each step whether to partition or which variable to eliminate. We assume that, in worst case, the cost of translating a ws-set S is $2^{|S|}$ (following the exponential formula of the inclusion-exclusion principle).

In case of independent partitioning, the partitions S_1, \dots, S_n are disjoint and can be computed in polynomial time (by

Estimate (WS-Set S , variable x in S) returns Real

```

missing_assignment := false;
foreach  $i \in \text{dom}_x$  do
  compute  $S_{x \mapsto i}$  and  $T$  as shown in Figure 4
  if  $|S_{x \mapsto i}| > 0$  then  $s_i = |S_{x \mapsto i} \cup T|$ 
  else  $s_i = 0$ ; missing_assignment = true; endif
if (missing_assignment) then  $e = |T|$  else  $e = 0$ 
foreach  $j \in \text{dom}_x$  such that  $s_j > 0$  do
   $e = e + \log_k(1 + k^{s_j - e})$ 
return  $e$ 

```

Figure 6: Log cost estimate for a variable choice.

computing the connected components of the graph of variables co-occurring within ws-descriptors). We thus reduce the computation cost from $2^{|S|}$ to $2^{|S_1|} + \dots + 2^{|S_n|}$. This method is, however, not always applicable and we need to apply variable elimination.

The main advantage of variable elimination is that S is divided into subsets $T \cup S_{x \mapsto i}$ without the dependencies enforced by variable x and thus subject to independent partitioning in the context of $x \mapsto i$. Consider s_i the size of the ws-set $T \cup S_{x \mapsto i}$. Then, the cost of choosing x is $\sum_{i \in \text{dom}_x} 2^{s_i}$. Of course, for those assignments of x that do not occur in S we have $T \cup S_{x \mapsto i} = T$ and can translate T only once. The computation cost using variable elimination can match that of independent partitioning only in the case that the assignments of the chosen variable partition the input ws-set S and thus T is empty.

Our first heuristic, called *minlog*, chooses a variable that minimizes $\log(\sum_{i=1}^{\text{dom}_x} 2^{s_i})$. Figure 6 shows how to compute incrementally the cost estimate by avoiding summation of potentially large numbers. The variable `missing_assignment` is used to detect whether there is at least one assignment of x not occurring in S for which T will be translated; in this case, T is only translated once (and not for every missing assignment).

The second heuristic, called *minmax*, approximates the cost estimate and chooses a variable that minimizes the maximal ws-set $T \cup S_{x \mapsto i}$. Both heuristics need time linear in the sizes of all variable domains plus of the ws-set. In addition to *minmax*, *minlog* needs to perform log and exp operations.

REMARK 4.6. To better understand our heuristics, we give one scenario where *minmax* behaves suboptimal. Consider S of size n and two variables. Variable x occurs with the same assignment in $n-1$ ws-descriptors and thus its *minmax* estimate is n , and variable y occurs twice with different assignments, and thus its *minmax* estimate is $n-1$. Using *minmax*, we choose y , although the *minlog* would choose differently: $e(y) = \log(2 \cdot 2^{n-1} + 2^{n-2}) > \log(2 \cdot 2^{n-1}) = e(x)$. \square

4.3 Probability computation

We next give an algorithm for computing the exact probability of a ws-set by employing the translation of ws-sets into ws-trees discussed in Section 4. Figure 7 defines the function P to this effect. This function is defined using pattern matching on the node types of ws-trees. The probability of an \otimes -node is the joint probability of its independent children $S_1, \dots, S_{|I|}$. The probability of an \oplus -node is the joint prob-

$$\begin{aligned}
P(\bigotimes_{i \in I} S_i) &= 1 - \prod_{i \in I} (1 - P(S_i)) \\
P(\bigoplus_{i \in I} (x \mapsto i : S_i)) &= \sum_{i \in I} P(\{x \mapsto i\}) \cdot P(S_i) \\
P(\emptyset) &= 1 & P(\perp) &= 0
\end{aligned}$$

Figure 7: Probability computation for ws-trees.

ability of its mutually exclusive children, where the probability of each child S_i is weighted by the probability of the variable assignment $x \mapsto i$ annotating the incoming edge of S_i . Finally, the probability of a leaf represented by the nullary ws-descriptor is 1 and of \perp is 0.

EXAMPLE 4.7. The probability of the ws-tree R of Figure 3 can be computed as follows (we label the inner nodes with l for left child and r for right child):

$$\begin{aligned}
P(R) &= 1 - (1 - P(l)) \cdot (1 - P(r)) \\
P(l) &= P(\{x \mapsto 1\}) \cdot P(\emptyset) + P(\{x \mapsto 2\}) \cdot P(lr) \\
P(lr) &= 1 - (1 - P(\{y \mapsto 1\}) \cdot P(\emptyset)) \cdot (1 - P(\{z \mapsto 1\}) \cdot P(\emptyset)) \\
P(r) &= P(\{u \mapsto 1\}) \cdot P(\{v \mapsto 1\}) \cdot P(\emptyset) + P(\{u \mapsto 2\}) \cdot P(\emptyset)
\end{aligned}$$

We can now replace the probabilities for variable assignments and ws-descriptor \emptyset and obtain

$$\begin{aligned}
P(r) &= 0.7 \cdot 0.5 \cdot 1 + 0.3 = 0.65 \\
P(lr) &= 1 - (1 - 0.2 \cdot 1) \cdot (1 - 0.4 \cdot 1) = 0.52 \\
P(l) &= 0.1 \cdot 1 + 0.4 \cdot 0.52 = 0.308 \\
P(R) &= 1 - (1 - 0.308) \cdot (1 - 0.65) = 0.7578 \quad \square
\end{aligned}$$

The probability of a ws-tree R can be computed in one bottom-up traversal of R and does not require the precomputation of R . The translation and probability computation functions can be easily composed to obtain the function $\text{ComputeTree} \circ P$ by inlining P in ComputeTree . As a result, the construction of the nodes \oplus , \otimes , and \emptyset is replaced by the corresponding probability computation given by P .

5. CONDITIONING

In this section we study the problem of conditioning a probabilistic database, i.e., the problem of removing all possible worlds that do not satisfy a given condition (say, by a Boolean relational calculus query) and normalizing the database such that, if there is at least one world left, the probability weights of all worlds sum up to one.

We will think of conditioning as a query or update operation assert_ϕ , where ϕ is the condition, i.e., a Boolean query. Processing relational algebra queries on probabilistic databases was discussed in Section 2. We will now assume the result of the Boolean query given as a ws-set defining the worlds on which ϕ is true.

EXAMPLE 5.1. Consider again the data cleaning example from the Introduction, formalized by the U-relational database of Figure 2. Relation W represents the set of possible worlds and U represents the tuples in these worlds.

As discussed in Example 2.3, the set of ws-descriptors $S = \{\{j \mapsto 1\}, \{j \mapsto 7, b \mapsto 4\}\}$ represents the three worlds on which the functional dependency $\text{SSN} \rightarrow \text{NAME}$ holds.

cond: conditioning algorithm
In: ws-tree R representing the new nonempty world-set.
ws-set U from the U-relations
Out: (confidence value, ws-set U')
if $R = \emptyset$ **then return** $(1, U)$;
if $R = \bigotimes_{i \in I} (R_i)$ **then**
 foreach $i \in I$ **do** $(c_i, U_i) := \text{cond}(R_i, U)$;
 return $(1 - \prod_{i \in I} (1 - c_i), \bigcup_{i \in I} U_i)$;
if $R = \bigoplus_{i \in \text{dom}_x} (x \mapsto i : R_i)$ **then**
 foreach $i \in \text{dom}_x$ **do**
 $U_i :=$ the subset of U consistent with $x \mapsto i$;
 $(c_i, U'_i) := \text{cond}(R_i, U_i)$;
 $c := \sum_{i \in I} P(\{x \mapsto i\}) \cdot c_i$;
 let x' be a new variable;
 foreach $i \in \text{dom}_x$ **such that** $c_i \neq 0$ **do**
 add $\langle x', i, \frac{P(\{x \mapsto i\}) \cdot c_i}{c} \rangle$ to the W relation;
 replace each occurrence of x in U'_i by x' ;
 return $(c, \bigcup_{i \in \text{dom}_x} U'_i)$;

Figure 8: The conditioning algorithm.

The world $\{j \mapsto 7, b \mapsto 7\}$ is excluded and thus the confidence of S does not add up to one but to $.2 + .8 \cdot .3 = .44$. What we now want to do is transform this database into one that represents the three worlds identified by S and preserves their tuples as well as their relative weights, but with a sum of world weights of one. This can of course be easily achieved by multiplying the weight of each of the three remaining worlds by $1/.44$. However, we want to do this in a smart way that in general does not require to consider each possible world individually, but instead preserves a succinct representation of the data and runs efficiently.

Such a technique exists and is presented in this section. It is based on running our confidence computation algorithm for ws-trees and, while returning from the recursion, renormalizing the world-set by introducing new variables whose assignments are normalized using the confidence values obtained. For this example, the conditioned database will be

W	Var	Dom	P
	b	4	.3
	b	7	.7
	j'	1	.2/.44
	j'	7	.8 \cdot .3/.44

U	WSD	SSN	NAME
	$\{j' \mapsto 1\}$	1	John
	$\{j' \mapsto 7\}$	7	John
	$\{j' \mapsto 1, b \mapsto 4\}$	4	Bill
	$\{j' \mapsto 1, b \mapsto 7\}$	7	Bill
	$\{j' \mapsto 7\}$	4	Bill

Note that the W relation actually models four possible worlds, but two of them, $\{j' \mapsto 7, b \mapsto 4\}$ and $\{j' \mapsto 7, b \mapsto 7\}$ are equal (contain the same tuples). Example 5.2 will show in detail how conditioning works. \square

Figure 8 gives our efficient algorithm for conditioning a U-relational database. The input is a U-relational database

U	WSD	A
	$\{y \mapsto 2, u \mapsto 1\}$	a_1
	$\{u \mapsto 1, v \mapsto 2\}$	a_2

ΔW	Var	Dom	P
	x'	1	.1/.308
	x'	2	.208/.308
	y'	1	1
	z'	1	1
	u'	1	.35/.65
	u'	2	.3/.65
	v'	1	1




Figure 9: U-relation U , additions ΔW to the W -relation, and a renormalized ws-tree.

and a ws-tree R that describes the subset of the possible worlds of the database that we want to condition it to. The output is a modified U-relational database and, as a by-product, since we recursively need to compute confidences for the renormalization, the confidence of R in the input database. The confidence of R in the output database will of course be 1. The renormalization works as follows. The probability of each branch of an inner node n of R is re-weighted such that the probability of n becomes 1. We reflect this re-weighting by introducing new variable whose assignments reflect the new weights of the branches of n .

This algorithm is essentially the confidence computation algorithm of Figure 7. We just add some lines of code along the line of recursively computing confidence that renormalize the weights of alternative assignments of variables for which some assignments become impossible. Additionally, we pass around a set of ws-descriptors (associated with tuples from the input U-relational database) and extend each ws-descriptor in that set by $x \mapsto i$ whenever we eliminate variable x , for each of its alternatives i .

EXAMPLE 5.2. Consider the U-relational database consisting of the W -relation of Figure 3 and the U-relation U of Figure 9. Let us run the algorithm to condition the database on the ws-tree R of Figure 3 (R need not be precomputed for conditioning).

We recursively call function `cond` at each node in the ws-tree R starting at the root. To simplify the explanation, let us assume a numbering of the nodes and of the ws-sets we pass around: If R_w is a (sub)tree then $R_{w,i}$ is its i -th child. The ws-set passed in the recursion with R_w is U_w and the ws-set returned is U'_w . The ws-sets passed on at the nodes of R are:

$$\begin{aligned}
U_1 &= U_2 = U \\
U_{1,1} &= x \mapsto 1 : U = \{\{x \mapsto 1, y \mapsto 2, u \mapsto 1\}, \\
&\quad \{x \mapsto 1, u \mapsto 1, v \mapsto 2\}\} \\
U_{1,2} &= x \mapsto 2 : U = \{\{x \mapsto 2, y \mapsto 2, u \mapsto 1\}, \\
&\quad \{x \mapsto 2, u \mapsto 1, v \mapsto 2\}\} \\
U_{1,2,1,1} &= y \mapsto 1 : U_{1,2} = \{\{y \mapsto 1, x \mapsto 2, u \mapsto 1, v \mapsto 2\}\} \\
U_{1,2,2,1} &= z \mapsto 1 : U_{1,2} = \{\{z \mapsto 1, x \mapsto 2, y \mapsto 2, u \mapsto 1\}, \\
&\quad \{z \mapsto 1, x \mapsto 2, u \mapsto 1, v \mapsto 2\}\} \\
U_{2,1} &= u \mapsto 1 : U_2 = U_2 \\
U_{2,2} &= u \mapsto 2 : U_2 = \emptyset \\
U_{2,1,1} &= v \mapsto 1 : U_{2,1} = \{\{v \mapsto 1, y \mapsto 2, u \mapsto 1\}\}
\end{aligned}$$

When we reach the leaves of R , we start returning from recursion and do the following. We first compute the probabilities of the nodes of R – in this case, they are already computed in Example 4.7. Next, for each \oplus -node representing the elimination of a variable, say α , we create a new variable α' with the assignments of α present at that node. In contrast to α , the assignments of α' are re-weighted by the probability of that \oplus -node so that the sum of their weights is 1. The new variables and their weighted assignments are given in Figure 9 along the original ws-tree R and in the ΔW relation to be added to the world table W .

When we return from recursion, we also compute the new ws-sets U'_i from U_i . These ws-sets are equal in case of leaves and \otimes -nodes, but, in case of \oplus -nodes, the variable eliminated at that node is replaced by the new one we created. In case of \oplus and \otimes nodes, we also return the union of all U'_i of their children. We finally return from the first call with the following ws-set U' :

$$\begin{aligned}
&\{\{x' \mapsto 1, y \mapsto 2, u \mapsto 1\}, \\
&\quad \{x' \mapsto 1, u \mapsto 1, v \mapsto 2\}, \\
&\quad \{x' \mapsto 2, y' \mapsto 1, u \mapsto 1, v \mapsto 2\}, \\
&\quad \{x' \mapsto 2, z' \mapsto 1, y \mapsto 2, u \mapsto 1\}, \\
&\quad \{x' \mapsto 2, z' \mapsto 1, u \mapsto 1, v \mapsto 2\}, \\
&\quad \{u' \mapsto 1, v' \mapsto 1, y \mapsto 2\}\}.
\end{aligned}$$

□

Let us view a probabilistic database semantically, as a set of pairs (I, p) of instances I with probability weights p .

THEOREM 5.3. Given a representation of probabilistic database $\mathbf{W} = \{(I_1, p_1), \dots, (I_n, p_n)\}$ and a ws-tree R identifying a nonempty subset of the worlds of \mathbf{W} , the algorithm of Figure 8 computes a representation of probabilistic database

$$\{(I_j, \frac{p_j}{c}) \mid (I_j, p_j) \in \mathbf{W}, I_j \in \omega(R)\}$$

such that the probabilities p_j add up to 1.

Thus, of course, c is the confidence of R .

Three simple optimizations of this algorithm that simplify the world table W and the output ws-descriptors are worth mentioning.

1. Variables that do not appear anywhere in the U-relations can be dropped from W .

2. Variables with a single domain value (obviously of weight 1) can be dropped everywhere from the database.
3. Variables x' and x'' obtained from the same variable x (by creation of a new variable in the case of variable elimination on x in two distinct branches of the recursion) can be merged into the same variable if the alternatives and their weights in the W relation are the same. In that case we can replace x'' by x' everywhere in the database.

EXAMPLE 5.4. In the previous example, we can remove the variables y' , z' , and v' from the W -relation and all variable assignments involving these variables from the U -relation because of (1). Furthermore, we can remove the variables x and z because of (1). The resulting database is

U'	WSD	A
	$\{x' \mapsto 1, y \mapsto 2, u \mapsto 1\}$	a_1
	$\{x' \mapsto 1, u \mapsto 1, v \mapsto 2\}$	a_2
	$\{x' \mapsto 2, u \mapsto 1, v \mapsto 2\}$	a_2
	$\{x' \mapsto 2, y \mapsto 2, u \mapsto 1\}$	a_1
	$\{x' \mapsto 2, u \mapsto 1, v \mapsto 2\}$	a_2
	$\{u' \mapsto 1, y \mapsto 2\}$	a_1

W'	Var	Dom	P
	x'	1	.1/.308
	x'	2	.208/.308
	y	1	.2
	y	2	.8
	u	1	.7
	u	2	.3
	u'	1	.35/.65
	u'	2	.3/.65
	v	1	.5
	v	2	.5

Finally, we state an important property of conditioning (expressed by the assert operation) useful for query optimization.

THEOREM 5.5. *Assert-operations commute with other asserts and the operations of positive relational algebra.*

6. WS-DESCRIPTOR ELIMINATION

We next present an alternative to exact probability computation using ws-trees based on the difference operation on ws-sets, called here ws-descriptor elimination. The idea is to incrementally eliminate ws-descriptors from the input ws-set. Given a ws-set S and a ws-descriptor d_1 in S , we compute two ws-sets: the original ws-set S without d_1 , and the ws-set representing the difference of $\{d_1\}$ and the first ws-set. The probability of S is then the sum of the probabilities of the two computed ws-sets, because the two ws-sets are mutex, as stated below by function P_w :

$$P_w(\emptyset) = 0 \quad P_w(\{\emptyset\}) = 1$$

$$P_w(S) = P_w(\{d_2, \dots, d_n\}) + \sum_{d \in (\{d_1\} - \{d_2, \dots, d_n\})} P(d)$$

The function P computes here the probability of a ws-descriptor.

EXAMPLE 6.1. Consider the ws-set $\{d_1, d_2, d_3\}$ of Example 3.1. The ws-descriptor d_2 is mutex with both d_1 and d_3

and we can eliminate it: $P_w(\{d_1, d_2, d_3\}) = P_w(\{d_1, d_3\}) + P(d_2)$. We now choose any to eliminate d_3 and obtain $P_w(\{d_1, d_3\}) = P_w(\{d_3\} - \{d_1\}) + P(d_1) = P(d_1)$, as explained in Example 3.3. Thus $P_w(\{d_1, d_2, d_3\}) = P(d_2) + P(d_1) = 1$. \square

This method exploits the fact that the difference operation preserves the mutex property and is world-set monotone.

LEMMA 6.2. *The following equations hold for any ws-sets S_1, S_2 , and S_3 :*

$$\begin{aligned} \omega(S_1 - S_2) &\subseteq \omega(S_1) \\ \omega(S_1) \cup \omega(S_2) &= \omega(S_1 - S_2) \cup \omega(S_2) \\ \emptyset &= \omega(S_1 - S_2) \cap \omega(S_2) \\ \omega(S_1) \cap \omega(S_2) = \emptyset &\Rightarrow \omega(S_1 - S_3) \cap \omega(S_2 - S_3) = \emptyset \end{aligned}$$

The correctness of probability computation by ws-descriptor elimination follows immediately from Lemma 6.2.

THEOREM 6.3. *Given a ws-set S , the function P_w computes the probability of S .*

As a corollary, we have that

COROLLARY 6.4 (THEOREM 6.3). *Any ws-set $\bigcup_{i=1}^n \{d_i\}$ has the equivalent mutex ws-set $\bigcup_{i=1}^{n-1} (\{d_i\} - \bigcup_{j=i+1}^n \{d_j\}) \cup \{d_n\}$.*

Like the translation of ws-sets into ws-trees, this method can take exponential time in the size of the input ws-set. Moreover, the equivalent mutex ws-set given above can be exponential. On the positive side, computing the exact probability of such mutex ws-sets can be done in linear time. Additionally, the probability of $\{d\} - S_d$ can be computed on the fly without requiring to first generate all ws-descriptors in the difference ws-set. This follows from the fact that the difference operation on ws-descriptors only generates mutex and distinct ws-descriptors. After generating a ws-descriptor from the difference ws-set we can thus add its probability to a running sum and discard it before generating the next ws-descriptor. The next section reports on experiments with an implementation of this method.

7. EXPERIMENTS

The experiments were conducted on an Athlon-X2 (4600+) x86-64bit/1.8GB/ Linux 2.6.20/gcc 4.1.2 machine.

We considered two synthetic data sets.

TPC-H data and queries. The first data set consists of tuple-independent probabilistic databases obtained from relational databases produced by TPC-H 2.7.0, where each tuple is associated with a Boolean random variable and the probability distribution is chosen at random. We evaluated the two Boolean queries of Figure 10 on each probabilistic database and then computed the probability of the ws-set consisting of the ws-descriptors of all the answer tuples. Among the two queries, only the second is *safe* and thus admits PTIME evaluation on tuple-independent probabilistic databases [9]. As we rewrite constraints into Boolean queries, we consider this querying scenario equally relevant to conditioning.

Query	Size of ws-desc.	TPC-H Scale	#Input Vars	Size of ws-set	User Time(s)
Q_1 : select true from customer c, orders o, lineitem l where c.mktsegment = 'BUILDING' and c.custkey = o.custkey and o.orderkey = l.orderkey and o.orderdate > '1995-03-15'	3	0.01	77215	9836	5.10
		0.05	382314	43498	99.76
		0.10	765572	63886	356.56
Q_2 : select true from lineitem where shipdate between '1994-01-01' and '1996-01-01' and discount between '0.05' and '0.08' and quantity < 24	1	0.01	60175	3029	0.20
		0.05	299814	15545	8.24
		0.10	600572	30948	33.68

Figure 10: TPC-H scenario: Queries, data characteristics, and performance of INDVE(minlog).

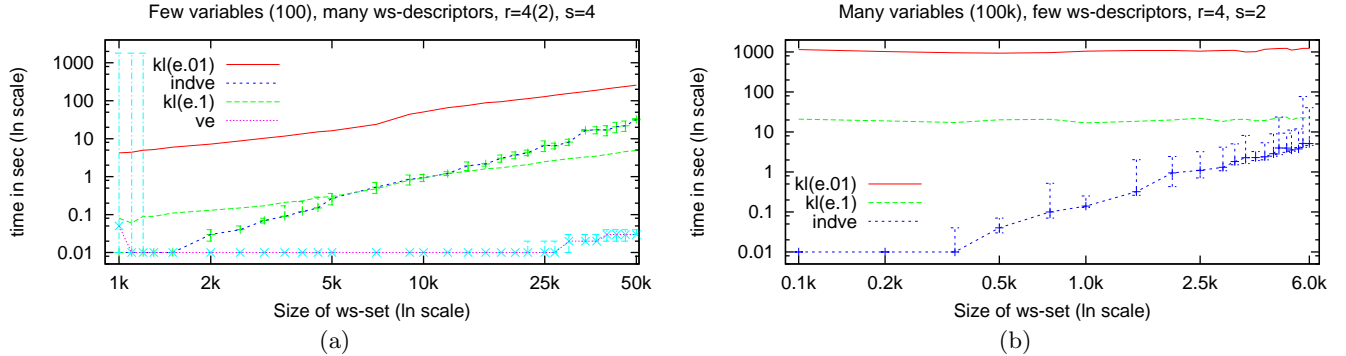


Figure 11: The two cases when the numbers of variables and of ws-descriptors differ by orders of magnitude.

#P-hard cases. The second data set consists of ws-sets similar to those associated with the answers of nonhierarchical conjunctive queries without self-joins on tuple-independent probabilistic databases, i.e. join queries such as $Q_s = R_1 \bowtie \dots \bowtie R_s$ for schemas $R_i(A_i, A_{i+1})$ in which all relations are joined together, but there is no single column common to all of them. Such queries are known to be #P-hard [9].

The data generation is simple: we partition the set of variables into s equally-sized sets V_1, \dots, V_s and then sample ws-sets $\{x_1 \mapsto a_1, \dots, x_s \mapsto a_s\}$ where x_i is from V_i and a_i is a random alternative for x_i , for $1 \leq i \leq s$. It is easy to verify that each such ws-set is actually the result of query Q_s on some tuple-independent probabilistic database. (For $s = 3$ this fact is used in the #P-hardness proof of [9].)

We use the following parameters in our experiments: number n of variables ranging from 50 to 100K, number r of possible alternatives per variable (2 or 4), length s of ws-descriptors, which equals the number of joined relations (2 or 4), and number w of ws-descriptors ranging from 5 to 60K. For each variable, the alternatives have uniform probabilities $1/r$: our exact algorithms are not sensitive to changing probability values as long as the numbers of alternatives of the variables remain constant.

Note that the focus on Boolean queries means no loss of generality for confidence computation; rather, the projection of a query result to a nullary relation causes all the ws-sets to be unioned and large.

Algorithms. We experimentally compared three versions of our exact algorithm: one that employs independent partitioning and variable elimination (INDVE), one that employs variable elimination only (VE), and one with ws-descriptor elimination (WE). We considered INDVE with the two heuristics minlog and minmax. These implementations compute confidence values and the modified world table (ΔW in Example 5.2), but do not materialize the modified, conditioned U-relations (U' in Example 5.2). We have verified that the computation of these additional data structures adds only

a small overhead over confidence computation in practice. We therefore do not distinguish in the sequel between confidence computation and conditioning. Note that our implementation is based on the straightforward composition of the ComputeTree and conditioning algorithms and does not need to materialize the ws-trees.

Although we also implemented a brute-force algorithm for probability computation, its timing is extremely bad and not reported. At a glance, this algorithm iterates over all worlds and sums up the probabilities of those that are represented by some ws-descriptors in the input ws-set. We also tried a slight improvement of the brute-force algorithm by first partitioning the input ws-set into independent subsets [22]. This version, too, performed bad and is not reported, as the partitioning can only be applied once at the beginning on the whole ws-set, yet most of our input ws-sets only exhibit independence in the context of variable eliminations.

We experimentally compared INDVE against a Monte Carlo simulation algorithm for confidence computation [21, 9] which is based on the Karp-Luby (KL) fully polynomial randomized approximation scheme (FPRAS) for DNF counting [17]. Given a DNF formula with m clauses, the base algorithm computes an (ϵ, δ) -approximation \hat{c} of the number of solutions c of the DNF formula such that

$$\Pr[|c - \hat{c}| \leq \epsilon \cdot c] \geq 1 - \delta$$

for any given $0 < \epsilon < 1$, $0 < \delta < 1$. It does so within $\lceil 4 \cdot m \cdot \log(2/\delta) / \epsilon^2 \rceil$ iterations of an efficiently computable estimator. This algorithm can be easily turned into an (ϵ, δ) -FPRAS for tuple confidence computation (see [19]). In our experiments, we use the optimal Monte-Carlo estimation algorithm of [8]. This is a technique to determine a small sufficient number of Monte-Carlo iterations (within a constant factor from optimal) based on first collecting statistics on the input by running the Monte Carlo simulation a small number of times. We use the version of the Karp-Luby unbiased estimator described in the book [24], which converges faster than the basic algorithm of [17], adapted to the problem of

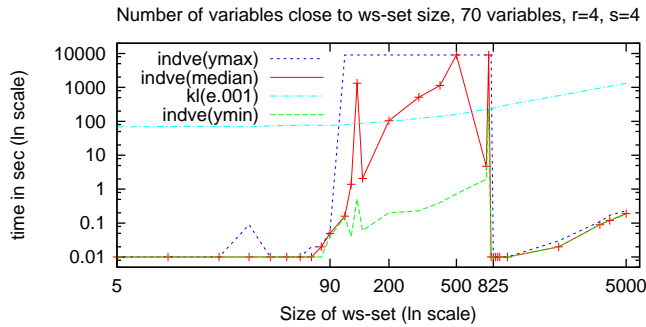


Figure 12: Performance of INDVE and KL when numbers of variables and ws-descriptors are close.

computing confidence values. This algorithm is similar to the self-adjusting coverage algorithm of [18].

1. Queries on TPC-H data. Figure 10 shows that INDVE(minlog) performs within hundreds of seconds in case of queries with equi-joins (Q_1) and selection-projection (Q_2) on tuple-independent probabilistic TPC-H databases with over 700K variables and 60K ws-descriptors. In the answers of query Q_2 , ws-descriptors are pairwise independent, and INDVE can effectively employ independence partitioning, making confidence computation more efficient than for Q_1 .

The remaining experiments use the second data generator.

2. The numbers of variables and of ws-descriptors differ by orders of magnitude. If there are much more ws-descriptors than variables, many ws-descriptors share variables (or variable assignments) and a good choice for variable elimination can effectively partition the ws-set. On the other hand, independence partitioning is unlikely to be very effective, and the time for checking it is wasted. Figure 11(a) shows that in such cases VE and INDVE (with minlog heuristic) are very stable and not influenced by fluctuations in data correlations. In particular, VE performs better than INDVE and within a second for 100 variables with domain size 4 (and nearly the same for 2), ws-descriptors of length 4, and ws-set size above 1.2k. We witnessed a sharp hard-easy transition at 1.2k, which suggests that the computation becomes harder when the number of ws-descriptors falls under one order of magnitude greater than the number of variables. Experiment 3 studies easy-hard-easy transitions in more detail. The plot data were produced from 25 runs and record the median value and ymin/ymax for the error bars.

In case of many variables and few ws-descriptors, the independence partitioning clearly pays off. This case naturally occurs for query evaluation on probabilistic databases, where a small set of tuples (and thus of ws-descriptors) is selected from a large database. As shown in Figure 11(b), INDVE(minlog) performs within seconds for the case of 100K variables and 100 to 6K ws-descriptors of size $s = 2$, and with variable domain size $r = 4$. Two further findings are not shown in the figure: (1) VE performs much worse than INDVE, as it cannot exploit the independence of tuples and thus creates partitions that overlap at large; (2) the case of $s = 4$ has a few (2 in 25) outliers exceeding 600 seconds.

3. The numbers of variables and of ws-descriptors are close. It is known from literature on knowledge compi-

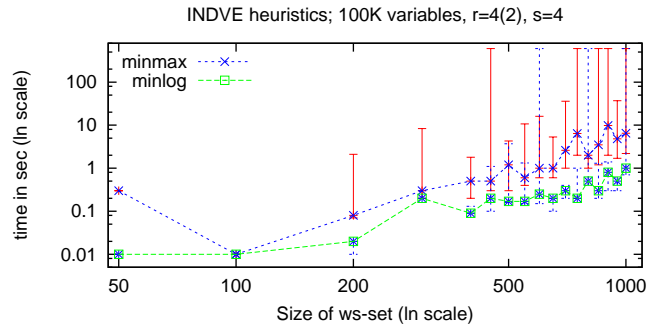


Figure 13: Heuristics: minmax versus minlog.

lation and model counting [6] that the computation becomes harder in this case. Figure 12 shows the easy-hard-easy pattern of INDVE(minlog) by plotting the minimal, maximal, and median computation time of 20 runs (max allowed time of 9000s). We experimentally observed the expected sharp transitions: When the numbers of ws-descriptors and of variables become close, the computation becomes hard and remains so until the number of ws-descriptors becomes one order of magnitude larger than the number of variables. The behavior of WE (not shown in the figure) follows very closely the easy-hard transition of INDVE, but in our experiment WE does not return anymore to the easy case within the range of ws-set sizes reported on in the figure.

4. Exact versus approximate computation. We experimentally verified our conjecture that the Karp-Luby approximation algorithm (KL) converges rather slowly. In case the numbers of variables and of ws-descriptors differ by orders of magnitude, INDVE(minlog) and VE(minlog) are definitely competitive when compared to KL with parameters $\epsilon = 0.1$ resp. $\epsilon = 0.01$, and $\delta = 0.01$, see Figure 11.

In Figure 11(b), KL uses about the same number of iterations for all the ws-set sizes, a sufficient number to warrant the running time. The reason for the near-constant line for KL is that for $s = 2$ and 100k variables, ws-descriptors are predominantly pairwise independent, and the confidence is close to $1 - (3/4)^w$, where w is the number of ws-descriptors. But this quickly gets close to 1, and the optimal algorithm can decide on a small number of iterations that does not increase with w . In case the numbers of variables and ws-descriptors are close (Figure 12), KL with $\epsilon = 0.001$ only performs better than INDVE(minlog) in the hard cases.

5. Heuristics for variable elimination. Figure 13 shows that, although the minmax heuristic is cheaper to compute than the minlog heuristic, using minlog we find in general better choices of variables and INDVE remains less sensitive to data correlations. The plot data are produced from 10 runs and show the median value and ymin/ymax for the error bars. Although VE exceeds the allocated time of 600 seconds for different data points, it does this less than five times (the median value is closer to ymin).

8. RELATED WORK

To the best of our knowledge, this paper is the first to study the conditioning problem for probabilistic databases. In this section, we survey related work in the areas of probabilistic databases and knowledge compilation procedures.

U-relations capture most other representation formalisms for uncertain data that were recently proposed in the literature, including those of MystiQ [9], Trio [5], and MayBMS [3].

For each of these formalisms, natural applications in data cleaning and other areas have been described [5, 4, 9].

Graphical models are a class of rich formalisms for representing probabilistic information which perform well in scenarios in which conditional probabilities and a known graph of dependencies and independences between events are available. There are, for instance, Bayesian network learning algorithms that produce just such data. Unfortunately, if probabilistic data is obtained by queries on tuple-independent or similar databases, the corresponding graphical models tend to be relatively flat [23] but have high tree-width, which causes techniques widely used for confidence computation on graphical models to be highly inefficient. Graphical models are more succinct than U-relations, yet their succinctness does not benefit the currently known query evaluation techniques. This justifies the development of conditioning techniques specifically for the c-table-like representations (such as U-relations) developed by the database community.

It has been long known that computing tuple confidence values on DNF-like representations of sets of possible worlds is a generalization of the DNF model counting problem and is #P-complete [10]. Monte Carlo approximation techniques for confidence computation have been known since the original work by Karp, Luby, and Madras [18]. Within the database field, this approach has first been followed in work on query reliability [14] and in the MystiQ project [9]. Section 7 reports on an experimental comparison of approximation and our exact algorithms.

Our variable elimination technique is based on Davis-Putnam procedure for satisfiability checking [12]. This procedure was already used for model counting [6]. Our approach combines it with independent partitioning for efficiently solving two more difficult problems: exact confidence computation and conditioning. [6] uses the minmax heuristic (which we benchmark against) and discusses experiments for CNF formulas with up to 50 variables and 200 clauses only. Our experiments also discuss new settings that are more natural in a database context: for instance, when the size of a query answer (and thus the number of ws-descriptors) is small in comparison to the size of the input database (and thus of variables). Follow-up work [15] reports on techniques for compiling ws-sets generated by conjunctive queries with inequalities into decision diagrams with polynomial-time guarantees.

Finally, there is a strong connection between ws-trees and ordered binary decision diagrams (OBDDs). Both make the structure of the propositional formulas explicit and allow for efficient manipulation. They differ, however, in important aspects: binary versus multistate variables, same variable ordering on all paths in case of OBDDs, and the new ws-tree \otimes -node type, which makes independence explicit. It is possible to reduce the gap between the two formalisms, but this affects the representation size. For instance, different variable orderings on different paths allows for exponentially more succinct BDDs [20]. Multistate variables can be easily translated into binary variables at a price of a logarithmic increase in the number of variables [25].

9. REFERENCES

- [1] S. Abiteboul, P. Kanellakis, and G. Grahne. “On the Representation and Querying of Sets of Possible Worlds”. *Theor. Comput. Sci.*, **78**(1):158–187, 1991.
- [2] P. Andritsos, A. Fuxman, and R. J. Miller. “Clean Answers over Dirty Databases: A Probabilistic Approach”. In *Proc. ICDE*, 2006.
- [3] L. Antova, T. Jansen, C. Koch, and D. Olteanu. “Fast and Simple Relational Processing of Uncertain Data”. In *Proc. ICDE*, 2008.
- [4] L. Antova, C. Koch, and D. Olteanu. “ 10^{10^6} Worlds and Beyond: Efficient Representation and Processing of Incomplete Information”. In *Proc. ICDE*, 2007.
- [5] O. Benjelloun, A. D. Sarma, A. Halevy, and J. Widom. “ULDBs: Databases with Uncertainty and Lineage”. In *Proc. VLDB*, 2006.
- [6] E. Birnbaum and E. Lozinskii. “The Good Old Davis-Putnam Procedure Helps Counting Models”. *Journal of AI Research*, **10**(6):457–477, 1999.
- [7] R. E. Bryant. Graph-based algorithms for boolean function manipulation. *IEEE Trans. Computers*, **35**(8):677–691, 1986.
- [8] P. Dagum, R. M. Karp, M. Luby, and S. M. Ross. “An Optimal Algorithm for Monte Carlo Estimation”. *SIAM J. Comput.*, **29**(5):1484–1496, 2000.
- [9] N. Dalvi and D. Suciu. “Efficient query evaluation on probabilistic databases”. *VLDB Journal*, **16**(4):523–544, 2007.
- [10] N. Dalvi and D. Suciu. “Management of Probabilistic Data: Foundations and Challenges”. In *Proc. PODS*, 2007.
- [11] A. Darwiche and P. Marquis. “A knowledge compilation map”. *Journal of AI Research*, **17**:229–264, 2002.
- [12] M. Davis and H. Putnam. “A Computing Procedure for Quantification Theory”. *Journal of ACM*, **7**(3):201–215, 1960.
- [13] N. Fuhr and T. Rölleke. “A Probabilistic Relational Algebra for the Integration of Information Retrieval and Database Systems”. *ACM Trans. Inf. Syst.*, **15**(1):32–66, 1997.
- [14] E. Grädel, Y. Gurevich, and C. Hirsch. “The Complexity of Query Reliability”. In *Proc. PODS*, pages 227–234, 1998.
- [15] J. Huang and D. Olteanu. Conjunctive queries with inequalities on probabilistic databases. Technical report, University of Oxford, 2008.
- [16] T. Imielinski and W. Lipski. “Incomplete information in relational databases”. *Journal of ACM*, **31**(4):761–791, 1984.
- [17] R. M. Karp and M. Luby. “Monte-Carlo Algorithms for Enumeration and Reliability Problems”. In *Proc. FOCS*, pages 56–64, 1983.
- [18] R. M. Karp, M. Luby, and N. Madras. “Monte-Carlo Approximation Algorithms for Enumeration Problems”. *J. Algorithms*, **10**(3):429–448, 1989.
- [19] C. Koch. “Approximating Predicates and Expressive Queries on Probabilistic Databases”. In *Proc. PODS*, 2008.
- [20] C. Meinel and T. Theobald. *Algorithms and Data Structures in VLSI Design*. Springer-Verlag, 1998.
- [21] C. Re, N. Dalvi, and D. Suciu. Efficient top-k query evaluation on probabilistic data. In *Proc. ICDE*, pages 886–895, 2007.
- [22] A. D. Sarma, M. Theobald, and J. Widom. “Exploiting Lineage for Confidence Computation in Uncertain and Probabilistic Databases”. In *Proc. ICDE*, 2008.
- [23] P. Sen and A. Deshpande. “Representing and Querying Correlated Tuples in Probabilistic Databases”. In *Proc. ICDE*, pages 596–605, 2007.
- [24] V. V. Vazirani. *Approximation Algorithms*. Springer, 2001.
- [25] M. Wachter and R. Haenni. “Multi-state Directed Acyclic Graphs”. In *Proc. Canadian AI*, pages 464–475, 2007.

APPENDIX

Proof of Theorem 4.4

We prove that the translation from ws-sets to ws-trees is correct. That is, given a ws-set S , $\text{ComputeTree}(S)$ and S represent the same world-set.

We use induction on the structure of ws-trees. In the base case, we map ws-sets representing the empty world-set to \perp , and ws-sets containing the universal ws-descriptor \emptyset (that represents the whole world-set) to \emptyset . We consider now a ws-set S . We have two cases corresponding to the different types of ws-tree inner nodes.

Case 1. Assume $S = \bigcup_{i \in I} S_i$ with S_i pairwise independent and $R_i = \text{ComputeTree}(S_i)$. By hypothesis, $\omega(R_i) = \omega(S_i)$. Then, $\text{ComputeTree}(S) = \otimes(R_i)$ and $\omega(\text{ComputeTree}(S)) =$

$$\bigcup_{i \in I} \omega(R_i) = \bigcup_{i \in I} \omega(S_i) = \omega(S).$$

Case 2. Let x be a variable in S and consider the ws-sets $S_{x \mapsto i}$ ($i \in \text{dom}_x$) and T as given by ComputeTree . Because the whole world-set can be represented by $A =$

$\bigcup_{i \in \text{dom}_x} \{x \mapsto i\}$, it holds that $\omega(A) \cap \omega(S) = \omega(S)$. We push the assignments of x in each ws-descriptor of S and obtain

$$\omega(S) = \omega\left(\bigcup_{i \in \text{dom}_x} \{d \cup \{x \mapsto i\} \mid d \in S\}\right).$$

We can remove all inconsistent ws-descriptors in the ws-set of the right-hand side while preserving equivalence:

$$\begin{aligned} & \omega(\{d \cup \{x \mapsto i\} \mid d \in S\}) = \\ & \omega(\{d \cup \{x \mapsto i\} \mid d \in S, \nexists j \in \text{dom}_x : j \neq i, \{x \mapsto j\} \subseteq d\}) = \\ & \omega(\{d \cup \{x \mapsto i\} \mid \{x \mapsto i\} \subseteq d \in S\}) \cup \\ & \omega(\{d \cup \{x \mapsto i\} \mid d \in S, \nexists j \in \text{dom}_x : \{x \mapsto j\} \subseteq d\}) = \\ & \omega(S_{x \mapsto i}) \cup \omega(T) = \omega(S_{x \mapsto i} \cup T) \end{aligned}$$

We now consider all values $i \in \text{dom}_x$ and obtain

$$\begin{aligned} \omega(S) &= \omega\left(\bigcup_{i \in \text{dom}_x} (S_{x \mapsto i} \cup T)\right) \\ &= \omega\left(\bigoplus_{i \in \text{dom}_x} x \mapsto i : (S_{x \mapsto i} \cup T)\right). \end{aligned}$$

Proof of Theorem 5.3

We prove that given a representation of probabilistic database $\mathbf{W} = \{(I_1, p_1), \dots, (I_n, p_n)\}$ and a ws-tree R identifying a nonempty subset of the worlds of \mathbf{W} , the algorithm of Figure 8 computes a representation of probabilistic database

$$\left\{ \left(I_j, \frac{p_j}{c} \right) \mid (I_j, p_j) \in \mathbf{W}, I_j \in \omega(R) \right\}$$

such that the probabilities p_j add up to 1.

The conditioning algorithm computes the probability c of each node of the input ws-tree R as given by our probability computation algorithm of Figure 7. We next consider the correctness of renormalization using induction on the structure of the input ws-tree.

Base case: The ws-tree \emptyset represents the whole world-set and we thus return U unchanged (no conditioning is done).

Induction cases \otimes (independent partitioning) and \oplus (variable elimination). For both node types, we return the union of ws-sets U'_i that are the ws-sets $U_i \subseteq U$ where the variables

encountered at the nodes on the recursion path are replaced by new ones. The ws-sets U_i are the subsets of U consistent with each child of the \oplus or \otimes node. By hypothesis, the ws-sets U_i are conditioned correctly. In case of \otimes -nodes, no further conditioning is done, because no re-weighting takes place. In case of a \oplus -node, we re-weight the assignments of the variable eliminated at that node.

Let $I \subseteq \text{dom}_x$ be the set of alternatives of x present at that node. Since

$$P(R) = P\left(\bigoplus_{i \in I} (x \mapsto i : R_i)\right) = \sum_{i \in I} P(\{x \mapsto i\}) \cdot P(R_i),$$

if we create a new variable x' ,

$$P(\{x' \mapsto i\}) := \frac{P(\{x \mapsto i\}) \cdot P(R_i)}{P(R)}.$$

This guarantees that

$$P\left(\bigoplus_{i \in I} (x' \mapsto i : R_i)\right) = 1.$$

If we ask which tuples of U should be in an instance satisfying R , the answer is of course all those whose ws-descriptors are consistent with one of the ws-descriptors in $x \mapsto i : R_i$ for some $i \in I$. The U -relation tuples in the results of the invocations $\text{cond}(R_i, U_i)$ grant exactly this.