

Approximate Lifted Inference with Probabilistic Databases

Wolfgang Gatterbauer
Carnegie Mellon University
gatt@cmu.edu

Dan Suciu
University of Washington
suciu@cs.washington.edu

ABSTRACT

This paper proposes a new approach for approximate evaluation of #P-hard queries with probabilistic databases. In our approach, every query is evaluated entirely in the database engine by evaluating a fixed number of query plans, each providing an upper bound on the true probability, then taking their minimum. We provide an algorithm that takes into account important schema information to enumerate only the minimal necessary plans among all possible plans. Importantly, this algorithm is a strict generalization of all known results of PTIME self-join-free conjunctive queries: A query is safe if and only if our algorithm returns one single plan. We also apply three relational query optimization techniques to evaluate all minimal safe plans very fast. We give a detailed experimental evaluation of our approach and, in the process, provide a new way of thinking about the value of probabilistic methods over non-probabilistic methods for ranking query answers.

1. INTRODUCTION

Probabilistic inference over large data sets is becoming a central data management problem. Recent large knowledge bases, such as Yago [27], Nell [5], DeepDive [9], or Google’s Knowledge Vault [14], have millions to billions of uncertain tuples. Data sets with missing values are often “completed” using inference in graphical models [6, 52] or sophisticated low rank matrix factorization techniques [15, 51], which ultimately results in a large, probabilistic database. Data sets that use crowdsourcing are also uncertain [1]. And, very recently, probabilistic databases have been applied to bootstrapping over samples of data [55].

However, probabilistic inference is known to be #P-hard in the size of the database, even for some very simple queries [7]. Today’s state of the art inference engines use either sampling-based methods or are based on some variant of the DPLL algorithm for Weighted Model Counting. For example, Tuffy [36], a popular implementation of Markov Logic Networks (MLN) over relational databases, uses Markov Chain Monte Carlo methods (MCMC). Gibbs sampling can be significantly improved by adapting some classical relational optimization techniques [56]. For another example, MayBMS [3] and its successor Sprout [39] use query plans

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/3.0/>. Obtain permission prior to any use beyond those covered by the license. Contact copyright holder by emailing info@vldb.org. Articles from this volume were invited to present their results at the 41st International Conference on Very Large Data Bases, August 31st - September 4th 2015, Kohala Coast, Hawaii.

Proceedings of the VLDB Endowment, Vol. 8, No. 5
Copyright 2015 VLDB Endowment 2150-8097/15/01.

to guide a DPLL-based algorithm for Weighted Model Counting [25]. While both approaches deploy some advanced relational optimization techniques, at their core they are based on general purpose probabilistic inference techniques, which either run in exponential time (DPLL-based algorithms have been proven recently to take exponential time even for queries computable in polynomial time [4]), or require many iterations until convergence.

In this paper, we propose a different approach to query evaluation with probabilistic databases. In our approach, *every query is evaluated entirely in the database engine*. Probability computation is done at query time, using simple arithmetic operations and aggregates. Thus, probabilistic inference is entirely reduced to a standard query evaluation problem with aggregates. There are no iterations and no exponential blowups. All benefits of relational engines (such as cost-based optimizations, multi-core query processing, shared-nothing parallelization) are directly available to queries over probabilistic databases. To achieve this, we compute approximate rather than exact probabilities, with a one-sided guarantee: The probabilities are guaranteed to be upper bounds to the true probabilities, which we show is *sufficient to rank the top query answers with high precision*. Our approach consists of approximating the true query probability by evaluating a fixed number of “safe queries” (the number depends on the query), each providing an upper bound on the true probability, then taking their minimum.

We briefly review “safe queries,” which are queries whose data complexity is in PTIME. They can be evaluated using safe query plans [7, 17, 53], which are related to a technique called *lifted inference* in the AI literature [12, 28]; the entire computation is pushed inside the database engine and is thus efficient. For example, the query $q_1(z) :- R(z, x), S(x, y), K(x, y)$ has the safe query plan $P_1 = \pi_z(R \bowtie_x (\pi_x(S \bowtie_{x,y} K)))$, where every join operator multiplies the probabilities, and every projection with duplicate elimination treats probabilistic events as independent. The literature describes several classes of safe queries [8, 17] and shows that they can be evaluated very efficiently. However, most queries are “unsafe.” They are provably #P-hard and do not admit safe plans.

In this paper, we prove that every conjunctive query without self-joins can be approximated by a fixed number of safe queries, called “*safe dissociations*” of the original query. Every safe dissociation is guaranteed to return an upper bound on the true probability and can be evaluated in PTIME data complexity. The number of safe dissociations depends only on the query and not the data. Moreover, we show how to find “*minimal safe dissociations*” which are sufficient to find the best approximation to the given query. For example, the unsafe query $q_2(z) :- R(z, x), S(x, y), T(y)$ has two minimal safe dissociations, $q'_2(z) :- R(z, x), S(x, y), T'(x, y)$ and $q''_2(z) :- R'(z, x, y), S(x, y), T(y)$. Both queries are safe and, by setting the probability of every tuple $R'(z, x, y)$ equal to that of

$R(z, x)$ and similarly for T' , they return an upper bound for the probabilities of each answer tuple from $q_2(z)$. One benefit of our approach is that, if the query happens to be safe, then it has a unique minimal safe dissociation, and our algorithm finds it.

Contributions. (1) We show that there exists a 1-to-1 correspondence between the safe dissociations of a self-join-free conjunctive query and its query plans. One simple consequence is that every query plan computes an upper bound of the true probability. For example, the two safe dissociations above correspond to the plans $P'_2 = \pi_x(R \bowtie_x (\pi_x(S \bowtie_{x,y} T)))$, and $P''_2 = \pi_z((\pi_{zy}(R \bowtie_x S)) \bowtie_y T)$. We give an intuitive system R-style algorithm [48] for enumerating all minimal safe dissociations of a query q . Our algorithm takes into account important schema-level information: functional dependencies and whether a relation is deterministic or probabilistic. We prove that our algorithm has several desirable properties that make it a *strict generalization of previous algorithms* described in the literature: If q is safe then the algorithm returns only one safe plan that computes q exactly; and if q happens to be safe on the particular database instance (e.g., the data happens to satisfy a functional dependency), then one of the minimal safe dissociations will compute the query exactly. (2) We use relational optimization techniques to compute all minimal safe dissociations of a query efficiently in the database engine. Some queries may have a large number of dissociations; e.g., a 8-chain query has 4279 safe dissociations, of which 429 are minimal. Computing 429 queries sequentially in the database engine would still be prohibitively expensive. Instead, we tailor three relational query optimization techniques to dissociation: (i) combining all minimal plans into *one single query*, (ii) reusing *common subexpressions* with views, and (iii) performing *deterministic semi-join reductions*. (3) We conduct an experimental validation of our technique, showing that, with all our optimizations enabled, computing hard queries over probabilistic databases incurs only a modest penalty over computing the same query on a deterministic database: For example, the 8-chain query runs only a factor of < 10 slower than on a deterministic database. We also show that the dissociation-based technique has high precision for ranking query answers based on their output probabilities.

In summary, our three main contributions are:

- (1) We describe an efficient algorithm for finding all minimal safe dissociations for self-join-free conjunctive queries in the presence of schema knowledge. If the query is safe, then our algorithm returns a single minimal plan, which is the safe plan for the query (Section 3).
- (2) We show how to apply three traditional query optimization techniques to dramatically improve the performance of the dissociation (Section 4).
- (3) We perform a detailed experimental validation of our approach, showing both its effectiveness in terms of query performance, and the quality of returned rankings. Our experiments also include a novel comparison between deterministic and probabilistic ranking approaches (Section 5).

All proofs for this submission together with additional illustrating examples are available in our technical report on arXiv [21].

2. BACKGROUND

Probabilistic Databases. We fix a relational vocabulary $\sigma = (R_1, \dots, R_m)$. A probabilistic database D is a database plus a function $p(t) \in [0, 1]$ associating a probability to each tuple $t \in D$. A *possible world* is a subset of D generated by independently including each tuple t in the world with probability $p(t)$. Thus, the database D is *tuple-independent*. We use bold notation (e.g., \mathbf{x}) to denote sets or tuples. A *self-join-free conjunctive query* is a first-order formula $q(\mathbf{y}) = \exists x_1 \dots \exists x_k. (a_1 \wedge \dots \wedge a_m)$ where each

atom a_i represents a relation $R_i(\mathbf{x}_i)$ ¹, the variables x_1, \dots, x_k are called *existential variables*, and \mathbf{y} are called the *head variables* (or free variables). The term “self-join-free” means that the atoms refer to distinct relational symbols. We assume therefore w.l.o.g. that every relational symbol R_1, \dots, R_m occurs exactly once in the query. Unless otherwise stated, a *query* in this paper denotes a self-join-free conjunctive query. As usual, we abbreviate the query by $q(\mathbf{y}) := a_1, \dots, a_m$, and write $\text{HVar}(q) = \mathbf{y}$, $\text{EVar}(q) = \{x_1, \dots, x_k\}$ and $\text{Var}(q) = \text{HVar}(q) \cup \text{EVar}(q)$ for the set of head variables, existential variables, and all variables of q . If $\text{HVar}(q) = \emptyset$ then q is called a *Boolean query*. We also write $\text{Var}(a_i)$ for the variables in atom a_i and $\text{at}(x)$ for the set of atoms that contain variable x . The *active domain* of a variable x_i is denoted ADom_{x_i} ², and the active domain of the entire database is $\text{ADom} = \bigcup_i \text{ADom}_{x_i}$. The focus of probabilistic query evaluation is to compute $\mathbb{P}(q)$; i.e. the probability that the query is true in a randomly chosen world.

Safe queries, safe plans. It is known that the data complexity of any query q is either in PTIME or #P-hard. The former are called *safe queries* and are characterized precisely by a syntactic property called *hierarchical queries* [7]. We briefly review these results:

DEFINITION 1 (HIERARCHICAL QUERY). *Query q is called hierarchical iff for any $x, y \in \text{EVar}(q)$, one of the following three conditions hold: $\text{at}(x) \subseteq \text{at}(y)$, $\text{at}(x) \cap \text{at}(y) = \emptyset$, or $\text{at}(x) \supseteq \text{at}(y)$.*

For example, the query $q_1 := -R(x, y), S(y, z), T(y, z, u)$ is hierarchical, while $q_2 := -R(x, y), S(y, z), T(z, u)$ is not, as neither of the three conditions holds for the variables y and z .

THEOREM 2 (DICHOTOMY [7]). *If q is hierarchical, then $\mathbb{P}(q)$ can be computed in PTIME in the size of D . Otherwise, computing $\mathbb{P}(q)$ is #P-hard in the size of D .*

We next give an equivalent, recursive characterization of hierarchical queries, for which we need a few definitions. We write $\text{SVar}(q)$ for the *separator variables* (or root variables); i.e. the set of existential variables that appear in every atom. q is disconnected if its atoms can be partitioned into two non-empty sets that do not share any existential variables (e.g., $q := -R(x, y), S(z, u), T(u, v)$ is disconnected and has two connected components: “ $R(x, y)$ ” and “ $S(z, u), T(u, v)$ ”). For every set of variables \mathbf{x} , denote $q - \mathbf{x}$ the query obtained by removing all variables \mathbf{x} (and decreasing the arities of the relation symbols that contain variables from \mathbf{x}).

LEMMA 3 (HIERARCHICAL QUERIES). *q is hierarchical iff either: (1) q has a single atom; (2) q has $k \geq 2$ connected components all of which are hierarchical; or (3) q has a separator variable x and $q - x$ is hierarchical.*

DEFINITION 4 (QUERY PLAN). *Let R_1, \dots, R_m be a relational vocabulary. A query plan P is given by the grammar*

$$P ::= R_i(\mathbf{x}) \mid \pi_{\mathbf{x}} P \mid \bowtie [P_1, \dots, P_k]$$

where $R_i(\mathbf{x})$ is a relational atom containing the variables \mathbf{x} and constants, $\pi_{\mathbf{x}}$ is the project operator with duplicate elimination, and $\bowtie [\dots]$ is the natural join in prefix notation, which we allow to be k -ary, for $k \geq 2$. We require that joins and projections alternate in a plan. We do not distinguish between join orders, i.e. $\bowtie [P_1, P_2]$ is the same as $\bowtie [P_2, P_1]$.

We write $\text{HVar}(P)$ for the head variables of P (defined as the variables \mathbf{x} of the top-most projection $\pi_{\mathbf{x}}$, or the union of the top-most

¹We assume w.l.o.g. that \mathbf{x}_i is a tuple of only variables without constants.

²Defined formally as $\text{ADom}_{x_i} = \bigcup_{j: x_j \in \text{Var}(R_j)} \pi_{x_i}(R_j)$.

projections if the last operation is a join). Every plan P represents a query q_P defined by taking all atoms mentioned in P and setting $\text{HVar}(q_P) = \text{HVar}(P)$. For notational convenience, we also use the “project-away” notation, by writing $\pi_{\mathbf{y}}(P)$ instead of $\pi_{\mathbf{x}}(P)$, where \mathbf{y} are the variables being projected away; i.e. $\mathbf{y} = \text{HVar}(P) - \mathbf{x}$.

Given a probabilistic database D and a plan P , each output tuple $t \in P(D)$ has a *score*(t), defined inductively on the structure of P as follows: If $t \in R_i(\mathbf{x})$, then $\text{score}(t) = p(t)$, i.e. its probability in D ; if $t \in \bowtie [P_1(D), \dots, P_k(D)]$ where $t = \bowtie [t_1, \dots, t_k]$, then $\text{score}(t) = \prod_{i=1}^k \text{score}(t_i)$; and if $t \in \pi_{\mathbf{x}}(P(D))$, and $t_1, \dots, t_n \in P(D)$ are all the tuples that project into t , then $\text{score}(t) = 1 - \prod_{i=1}^n (1 - \text{score}(t_i))$. In other words, *score* computes a probability by assuming that all tuples joined by \bowtie are independent, and all duplicates eliminated by π are also independent. If these conditions hold, then *score* is the correct query probability, but in general the score is different from the probability. Therefore, *score* is not equal to the probability, in general, and is also called an extensional semantics [18, 41]. For a Boolean plan P , we get one single score, which we denote $\text{score}(P)$.

The requirement that joins and projections alternate is w.l.o.g. because nested joins like $\bowtie [\bowtie [R_1, R_2], R_3]$ can be rewritten into $\bowtie [R_1, R_2, R_3]$ while keeping the same probability score. For the same reason we do not distinguish between different join orders.

DEFINITION 5 (SAFE PLAN). A plan P is called *safe* iff, for any join operator $\bowtie^P [P_1, \dots, P_k]$, all subplans have the same head variables: $\text{HVar}(P_i) = \text{HVar}(P_j)$ for all $1 \leq i, j \leq k$.

The recursive definition of Lemma 3 gives us immediately a safe plan for a hierarchical query. Conversely, every safe plan defines a hierarchical query. The following summarizes our discussion:

PROPOSITION 6 (SAFETY [7]). (1) Let P be a plan for the query q . Then $\text{score}(P) = \mathbb{P}(q)$ for any probabilistic database iff P is safe. (2) Assuming $\#P \neq \text{PTIME}$, a query q is safe (i.e. $\mathbb{P}(q)$ has PTIME data complexity) iff it has a safe plan P ; in that case the safe plan is unique, and $\mathbb{P}(q) = \text{score}(P)$.

Boolean Formulas. Consider a set of Boolean variables $\mathbf{X} = \{X_1, X_2, \dots\}$ and a probability function $p : \mathbf{X} \rightarrow [0, 1]$. Given a Boolean formula F , denote $\mathbb{P}(F)$ the probability that F is true if each variable X_i is independently true with probability $p(X_i)$. In general, computing $\mathbb{P}(F)$ is $\#P$ -hard in the number of variables \mathbf{X} . If D is a probabilistic database then we interpret every tuple $t \in D$ as a Boolean variable and denote the lineage of a Boolean $q : -a_1, \dots, a_m$ on D as the Boolean DNF formula $F_{q,D} = \bigvee_{\theta: \theta \models q} \theta(a_1) \wedge \dots \wedge \theta(a_m)$, where θ ranges over all assignments of $\text{EVar}(q)$ that satisfy q on D . It is well known that $\mathbb{P}(q) = \mathbb{P}(F_{q,D})$. In other words the probability of a Boolean query is the same as the probability of its lineage formula.

EXAMPLE 7 (LINEAGE). If $F = XY \vee XZ$ then $\mathbb{P}(F) = p(1 - (1-q)(1-r)) = pq + pr - pqr$, where $p = p(X)$, $q = p(Y)$, and $r = p(Z)$. Consider now the query $q : -R(x), S(x, y)$ over the database $D = \{R(1), R(2), S(1, 4), S(1, 5)\}$. Then the lineage formula is $F_{q,D} = R(1) \wedge S(1, 4) \vee R(1) \wedge S(1, 5)$, i.e. same as F , up to variable renaming. It is now easy to see that $\mathbb{P}(q) = \mathbb{P}(F_{q,D})$.

A key technique that we use in this paper is the following result from [22]: Let F, F' be two Boolean formulas with sets of variables \mathbf{X} and \mathbf{X}' , respectively. We say that F' is a *dissociation* of F if there exists a substitution $\theta : \mathbf{X}' \rightarrow \mathbf{X}$ such that $F'[\theta] = F$. If $\theta^{-1}(X) = \{X', X'', \dots\}$ then we say that the variable X dissociates into X', X'', \dots ; if $|\theta^{-1}(X)| = 1$ then we assume w.l.o.g. that $\theta^{-1}(X) = X$ (up to variable renaming) and we say that X

does not dissociate. Given a probability function $p : \mathbf{X} \rightarrow [0, 1]$, we extend it to a probability function $p' : \mathbf{X}' \rightarrow [0, 1]$ by setting $p'(X') = p(\theta(X'))$. Then, we have shown:

THEOREM 8 (OBLIVIOUS DNF BOUNDS [22]). Let F' be a monotone DNF formula that is a dissociation of F through the substitution θ . Assume that for any variable X , no two distinct dissociations X', X'' of X occur in the same prime implicant of F' . Then (1) $\mathbb{P}(F) \leq \mathbb{P}(F')$, and (2) if all dissociated variables $X \in \mathbf{X}$ are deterministic (meaning: $p(X) = 0$ or $p(X) = 1$) then $\mathbb{P}(F) = \mathbb{P}(F')$.

Intuitively, a dissociation F' is obtained from a formula F by taking different occurrences of a variable X and replacing them with fresh variables X', X'', \dots ; in doing this, the probability of F' may be easier to compute, giving us an upper bound for $\mathbb{P}(F)$.

EXAMPLE 9 (EXAMPLE 7 CONT.). $F' = X'Y \vee X''Z$ is a dissociation of $F = XY \vee XZ$, and its probability is $\mathbb{P}(F') = 1 - (1-pq)(1-pr) = pq + pr - p^2qr$. Here, only the variable X dissociates into X', X'' . It is easy to see that $\mathbb{P}(F) \leq \mathbb{P}(F')$. Moreover, if $p = 0$ or 1 , then $\mathbb{P}(F) = \mathbb{P}(F')$. The condition that no two dissociations of the same variable occur in a common prime implicant is necessary: for example, $F' = X'X''$ is a dissociation of $F = X$. However $\mathbb{P}(F) = p$, $\mathbb{P}(F') = p^2$, and we do not have $\mathbb{P}(F) \leq \mathbb{P}(F')$.

3. DISSOCIATION OF QUERIES

This section introduces our main technique for approximate query processing. After defining dissociations (Section 3.1), we show that some of them are in 1-to-1 correspondence with query plans, then derive our first algorithm for approximate query processing (Section 3.2). Finally, we describe two extensions in the presence of deterministic relations or functional dependencies (Section 3.3).

3.1 Query dissociation

DEFINITION 10 (DISSOCIATION). Given a Boolean query $q : -R_1(\mathbf{x}_1), \dots, R_m(\mathbf{x}_m)$ and a probabilistic database D . Let $\Delta = (\mathbf{y}_1, \dots, \mathbf{y}_m)$ be a collection of sets of variables with $\mathbf{y}_i \subseteq \text{Var}(q) - \text{Var}(g_i)$ for every relation R_i . The dissociation defined by Δ has then two components:

- (1) the dissociated query: $q^\Delta : -R_1^{\mathbf{y}_1}(\mathbf{x}_1, \mathbf{y}_1), \dots, R_m^{\mathbf{y}_m}(\mathbf{x}_m, \mathbf{y}_m)$, where each $R_i^{\mathbf{y}_i}(\mathbf{x}_i, \mathbf{y}_i)$ is a new relation of arity $|\mathbf{x}_i| + |\mathbf{y}_i|$.
- (2) the dissociated database instance D^Δ consisting of the tables over the vocabulary σ^Δ obtained by evaluating (deterministically) the following queries over the instance D :

$$R_i^{\mathbf{y}_i}(\mathbf{x}_i, \mathbf{y}_i) : -R_i(\mathbf{x}_i), \text{ADom}_{\mathbf{y}_{i1}}(\mathbf{y}_{i1}), \dots, \text{ADom}_{\mathbf{y}_{ik}}(\mathbf{y}_{ik})$$

where $\mathbf{y}_i = (\mathbf{y}_{i1}, \dots, \mathbf{y}_{ik_i})$. For each tuple $t' \in R_i^{\mathbf{y}_i}$, its probability is defined as $p'(t') = p(\pi_{\mathbf{x}_i}(t'))$, i.e. the probability of t' in the database D .

Thus, a dissociation acts on both the query expression and the database instance: It adds some variables \mathbf{y}_i to each relational symbol R_i of the query expression, and it computes a new instance for each relation $R_i^{\mathbf{y}_i}$ by copying every record $t \in R_i$ once for every tuple in the cartesian product $\text{ADom}_{\mathbf{y}_{i1}} \times \dots \times \text{ADom}_{\mathbf{y}_{ik}}$. When $\mathbf{y}_i = \emptyset$ then we abbreviate R_i^\emptyset with R_i . We give a simple example:

EXAMPLE 11 (EXAMPLE 7 CONT.). Consider $q : -R(x), S(x, y)$. Then $\Delta = (\{y\}, \emptyset)$ defines the following dissociation: $q^\Delta = R^y(x, y), S(x, y)$, and the new relation R^y contains the tuples $R^y(1, 4), R^y(1, 5), R^y(2, 4), R^y(2, 5)$. Notice that the lineage of the dissociated query q^Δ is $F_{q^\Delta, D^\Delta} = R^y(1, 4), S(1, 4) \vee R^y(1, 5), S(1, 5)$ and is the same (up to variable renaming) as the dissociation of the lineage of query q : $F' = X'Y \vee X''Z$.

THEOREM 12 (UPPER QUERY BOUNDS). For every dissociation Δ of q : $\mathbb{P}(q) \leq \mathbb{P}(q^\Delta)$.

PROOF. Theorem 12 follows immediately from Theorem 8 by noting that the lineage F_{q^Δ, D^Δ} is a dissociation of the lineage $F_{q, D}$ through the substitution $\theta : D^\Delta \rightarrow D$ defined as follows: for every tuple $t' \in R_i^{y_i}$, $\theta(t') = \pi_{x_i}(t')$. \square

DEFINITION 13 (SAFE DISSOCIATION). A dissociation Δ of a query q is called safe if the dissociated query q^Δ is safe.

By Theorem 2, a dissociation is safe (i.e. its probability can be evaluated in PTIME) iff q^Δ is hierarchical. Hence, amongst all dissociations, we are interested in those that are easy to evaluate and use them as a technique to approximate the probabilities of queries that are hard to compute. The idea is simple: Find a safe dissociation Δ , compute $\mathbb{P}(q^\Delta)$, and thereby obtain an upper bound on $\mathbb{P}(q)$. In fact, we will consider *all* safe dissociations and take the minimum of their probabilities, since this gives an even better upper bound on $\mathbb{P}(q)$ than that given by a single dissociation. We call this quantity the *propagation score*³ of the query q .

DEFINITION 14 (PROPAGATION). The propagation score $\rho(q)$ for a query q is the minimum score of all safe dissociations: $\rho(q) = \min_{\Delta} \mathbb{P}(q^\Delta)$ with Δ ranging over all safe dissociations.

The difficulty in computing $\rho(q)$ is the number of dissociations that is large even for relatively small queries: If q has k existential variables and m atoms, then q has $2^{|K|}$ possible dissociations with $K = \sum_{i=1}^m (k - |\text{Var}(a_i)|)$ forming a partial order in the shape of a power set lattice (see Fig. 1a). Therefore, our next step is to prune the space of dissociations and to examine only the minimum number necessary. We start by defining a partial order on dissociations:

DEFINITION 15 (PARTIAL DISSOCIATION ORDER). We define the partial order on the dissociations of a query as:

$$\Delta \preceq \Delta' \Leftrightarrow \forall i : \mathbf{y}_i \subseteq \mathbf{y}'_i$$

Whenever $\Delta \preceq \Delta'$, then q^Δ, D^Δ is a dissociation of $q^{\Delta'}, D^{\Delta'}$ (given by $\Delta'' = \Delta' - \Delta$). Therefore, we obtain immediately:

COROLLARY 16 (PARTIAL DISSOCIATION ORDER). If $\Delta \preceq \Delta'$ then $\mathbb{P}(q^\Delta) \leq \mathbb{P}(q^{\Delta'})$.

EXAMPLE 17 (PARTIAL DISSOCIATION ORDER). Consider the query $q : -R(x), S(x), T(x, y), U(y)$. It is unsafe and allows $2^3 = 8$ dissociations which are shown in Fig. 1a with the help of an “augmented incidence matrix”: each row represents one relation and each column one variable: An empty circle (\circ) indicates that a relation contains a variable; a full circle (\bullet) indicates that a relation is dissociated on a variable (the reason for using two separate symbols becomes clear when we later include domain knowledge). Among those 8 dissociations, 5 are safe, shaded in green, and have the hierarchy among variables highlighted. Furthermore, 2 of the 5 safe dissociations are minimal: $q^{\Delta_3} : -R(x), S(x), T(x, y), U^x(x, y)$, and $q^{\Delta_4} : -R^y(x, y), S^y(x, y), T(x, y), U(y)$. To illustrate that these dissociations are upper bounds, consider a database with $R = T = U = \{1, 2\}$, $S = \{(1, 1), (1, 2), (2, 2)\}$, and the probability of all tuples = $\frac{1}{2}$. Then q has probability $\frac{83}{2^9} \approx 0.161$, while q^{Δ_3} has probability $\frac{169}{2^{10}} \approx 0.165$, and q^{Δ_4} has probability $\frac{353}{2^{11}} \approx 0.172$, both of which are upper bounds. The propagation score is the minimum score of all minimal safe dissociations and thus ≈ 0.165 . \blacksquare

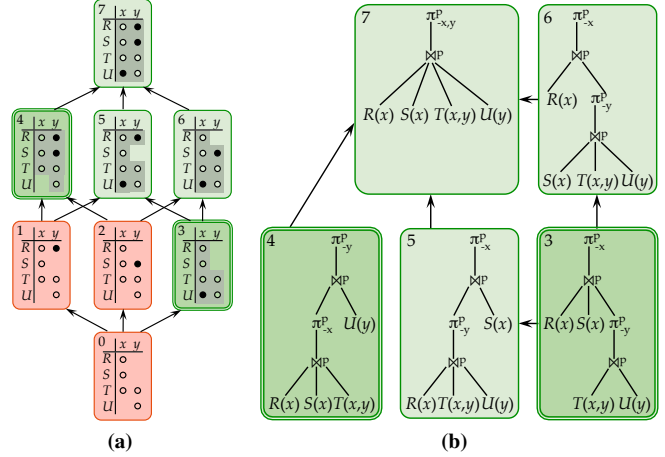


Figure 1: Example 17 (a): Partial dissociation order for $q : -R(x), S(x), T(x, y), U(y)$. Safe dissociations are green and have the hierarchies between variables shown (3 to 7), minimal safe dissociations are dark green and double-lined (3 and 4). (b): All 5 query plans for q and their correspondence to safe dissociations (3 to 7).

In general, the set of dissociations forms a lattice, with the smallest element $\Delta_{\perp} = (\emptyset, \dots, \emptyset)$ ($q^{\Delta_{\perp}} = q$) and the largest element $\Delta_{\top} = (\text{Var}(q) - \text{Var}(a_1), \dots, \text{Var}(q) - \text{Var}(a_m))$ ($q^{\Delta_{\top}}$ is safe, since every atom contains all variables). As we move up in the lattice the probability increases, but the safe/unsafe status may toggle arbitrarily from safe to unsafe and back. For example $q : -R(x), S(x), T(y)$ is safe, its dissociation $q' : -R(x), S^y(x, y), T(y)$ is unsafe, yet the next dissociation $q'' : -R(x), S^y(x, y), T^x(x, y)$ is safe again.

This suggests the following naive algorithm for computing $\rho(q)$: Enumerate all dissociations $\Delta_1, \Delta_2, \dots$ by traversing the lattice breadth-first, bottom up (i.e. whenever $\Delta_i \prec \Delta_j$ then $i < j$). For each dissociation Δ_i , check if q^{Δ_i} is safe. If so, then first update $\rho \leftarrow \min(\rho, \mathbb{P}(q^{\Delta_i}))$, then remove from the list all dissociations $\Delta_j \succ \Delta_i$. However, this algorithm is inefficient for practical purposes for two reasons: (i) we need to iterate over many dissociations in order to discover those that are safe; and (ii) computing $\mathbb{P}(q^{\Delta_i})$ requires computing a new database instance D^{Δ_i} for each safe dissociation Δ_i . We show in the next section how to avoid both sources of inefficiency by exploiting the lattice structure and by iterating over query plans instead of safe dissociations.

3.2 Dissociations and Plans

We prove here that the safe dissociations q^Δ are in 1-to-1 correspondence with query plans of the original query q . This allows us to (i) efficiently find safe dissociations (by iterating over query plans instead of all dissociations), and to (ii) compute $\mathbb{P}(q^\Delta)$ without having to materialize the dissociated database D^Δ .

We next describe the 1-to-1 mapping. Consider a safe dissociation q^Δ and denote its corresponding unique safe plan P^Δ . This plan uses dissociated relations, hence each relation $R_i^{y_i}(x_i, y_i)$ has extraneous variables \mathbf{y}_i . Drop all variables \mathbf{y}_i from the relations and all operators using them: This transforms P^Δ into a regular, generally unsafe plan P for q . For a trivial example, the plan corresponding to the top dissociation Δ_{\top} of a query q is $\pi_{\text{Var}(q)}(\bowtie [P_1, \dots, P_k])$: It performs all joins first, followed by all projections.

³We chose the name “propagation” for our method because of similarities with efficient belief propagation algorithms in graphical models. See [21] for a discussion on how query dissociation generalizes relevance propagation from graphs to hypergraphs, and [19] for a recent approach for speeding up belief propagation even further.

Conversely, consider any plan P for q . We define its corresponding safe dissociation Δ^P as follows. For each join operation $\bowtie^P [P_1, \dots, P_k]$, let its *join variables* JVar be the union of the head variables of all subplans: $\text{JVar} = \bigcup_j \text{HVar}(P_j)$. For every relation R_i occurring in P_j , add the missing variables $\text{JVar} - \text{HVar}(P_j)$ to \mathbf{y}_i . For example, consider $\bowtie^P [R(x), T(x, y), U(y)]$ (this is the lower join in query plan 5 of Fig. 1b). Here, $\text{JVar} = \{x, y\}$, and the corresponding safe dissociation of this subplan is $q^\Delta(x, y) :- R^y(x, y), T(x, y), U^x(x, y)$. Note that while there is a one-to-one mapping between safe dissociations and query plans, unsafe dissociations do not correspond to plans.

THEOREM 18 (SAFE DISSOCIATION). *Let q be a conjunctive query without self-joins. (1) The mappings $\Delta \mapsto P^\Delta$ and $P \mapsto \Delta^P$ are inverses of each other. (2) For every safe dissociation Δ , $\mathbb{P}(q^\Delta) = \text{score}(P^\Delta)$.*

COROLLARY 19 (UPPER BOUNDS). *Let P be any plan for a Boolean query q . Then $\mathbb{P}(q) \leq \text{score}(P)$.*

The proof follows immediately from $\mathbb{P}(q) \leq \mathbb{P}(q^{\Delta^P})$ (Theorem 12) and $\mathbb{P}(q^{\Delta^P}) = \text{score}(P)$ (Theorem 18). In other words, any plan for q computes a probability score that is guaranteed to be an upper bound on the correct probability $\mathbb{P}(q)$.

Theorem 18 suggests the following improved algorithm for computing the propagation score $\rho(q)$ of a query: Iterate over all plans P , compute their scores, and retain the minimum score $\min_P [\text{score}(P)]$. Each plan P is evaluated directly on the original probabilistic database, and there no need to materialize the dissociated database instance. However, this approach is still inefficient because it computes several plans that correspond to non-minimal dissociations. For example, in Fig. 1 plans 5, 6, 7 correspond to non-minimal dissociations, since plan 3 is safe and below them.

Enumerating minimal safe dissociations. Call a plan P *minimal* if Δ^P is minimal in the set of safe dissociations. For example, in Example 17, the minimal plans are 3 and 4. The propagation score is thus the minimum of the scores of the two minimal plans: $\rho(q) = \min_{i \in \{3, 4\}} [\text{score}(P^{(i)})]$. Our improved algorithm will iterate only over minimal plans, by relying on a connection between plans and sets of variables that disconnect a query: A *cut-set* is a set of existential variables $\mathbf{x} \in \text{EVar}(q)$ s.t. $q - \mathbf{x}$ is disconnected. A *min-cut-set* (for minimal cut-set) is a cut-set for which no strict subset is a cut-set. We denote $\text{MinCuts}(q)$ the set of all min-cut-sets. Note that q is disconnected iff $\text{MinCuts}(q) = \{\emptyset\}$.

The connection between $\text{MinCuts}(q)$ and query plans is given by two observations: (1) Let P be any plan for q . If q is connected, then the last operator in P is a projection, i.e. $P = \pi_{-\mathbf{x}}(\bowtie [P_1, \dots, P_k])$, and the projection variables \mathbf{x} are the join variables $\mathbf{x} = \text{JVar}$ because q is Boolean so the plan must project away all variables. We claim that \mathbf{x} is a cut-set for q and that $q - \mathbf{x}$ has k connected components corresponding to P_1, \dots, P_k . Indeed, if P_i, P_j share any common variable y , then they must join on y , hence $y \in \text{JVar}$. Thus, *cut-sets are in 1-to-1 correspondence with the top-most projection operator* of a plan. (2) Now suppose that P corresponds to a safe dissociation Δ^P , and let $P' = \pi_{-\mathbf{x}}(\bowtie [P'_1, \dots, P'_k])$ be its unique safe plan. Then $\mathbf{x} = \text{SVar}(q^{\Delta^P})$; i.e. the top-most project operator removes all separator variables.⁴ Furthermore, if $\Delta \succeq \Delta^P$ is a larger dissociation, then $\text{SVar}(q^\Delta) \supseteq \text{SVar}(q^{\Delta^P})$ (because any separator variable of a query continues to be a separator variable in any dissociation of that query). Thus, *minimal plans correspond to min-cut-sets*; in other words, $\text{MinCuts}(q)$ is in 1-to-1 correspondence with the top-most projection operator of *minimal* plans.

⁴This follows from the recursive definition of the unique safe plan of a query in Lemma 3: the top most projection consists precisely of its separator variables.

```

Recursive algorithm: MP (EnumerateMinimalPlans)
Input: Query  $q(\mathbf{x}) :- R_1(\mathbf{x}_1), \dots, R_m(\mathbf{x}_m)$ 
Output: Set of all minimal query plans  $\mathcal{P}$ 
1 if  $m = 1$  then  $\mathcal{P} \leftarrow \{\pi_{\mathbf{x}} R_1(\mathbf{x}_1)\}$ 
2 else
3   Set  $\mathcal{P} \leftarrow \emptyset$ 
4   if  $q$  is disconnected then
5     Let  $q = q_1, \dots, q_k$  be the connected components of  $q - \text{HVar}(q)$ 
6     foreach  $q_i$  do Let  $\text{HVar}(q_i) \leftarrow \text{HVar}(q) \cap \text{Var}(q_i)$ 
7     foreach  $(P_1, \dots, P_k) \in \text{MP}(q_1) \times \dots \times \text{MP}(q_k)$  do
8        $\mathcal{P} \leftarrow \mathcal{P} \cup \{\bowtie^P [P_1, \dots, P_k]\}$ 
9   else
10    foreach  $\mathbf{y} \in \text{MinCuts}(q - \text{HVar}(q))$  do
11      Let  $q' \leftarrow q$  with  $\text{HVar}(q') \leftarrow \text{HVar}(q) \cup \mathbf{y}$ 
12      foreach  $P \in \text{MP}(q')$  do  $\mathcal{P} \leftarrow \mathcal{P} \cup \{\pi_{-\mathbf{y}} P\}$ 

```

Algorithm 1: generates all minimal query plans for a given query q .

Our discussion leads immediately to Algorithm 1 for computing the propagation score $\rho(q)$. It also applies to non-Boolean queries by treating the head variables as constants, hence ignoring them when computing connected components. The algorithm proceeds recursively. If q is a single atom then it is safe and we return its unique safe plan. If the query has more than one atom, then we consider two cases, when $q - \text{HVar}(q)$ is disconnected or connected. In the first case, every minimal plan is a join, where the subplans are minimal plans of the connected components. In the second case, a minimal plan results from a projection over min-cut-sets. Notice that recursive calls of the algorithm will alternate between these two cases, until they reach a single atom.

THEOREM 20 (ALGORITHM 1). *Algorithm 1 computes the set of all minimal query plans.*

Conservativity. Some probabilistic database systems first check if a query q is safe, and in that case compute the exact probability using the safe plan, otherwise use some approximation technique. We show that Algorithm 1 is conservative, in the sense that, if q is safe, then $\rho(q) = \mathbb{P}(q)$. Indeed, in that case $\text{MP}(q)$ returns a single plan, namely the safe P for q , because the empty dissociation, $\Delta_\perp = (\emptyset, \dots, \emptyset)$, is safe, and it is the bottom of the dissociation lattice, making it the unique minimal safe dissociation.

Score Quality. We show here that the approximation of $\mathbb{P}(q)$ by $\rho(q)$ becomes tighter as the input probabilities in D decrease. Thus, the smaller the probabilities in the database, the closer does the ranking based on the propagation score approximate the ranking by the actual probabilities.

PROPOSITION 21 (SMALL PROBABILITIES). *Given a query q and database D . Consider the operation of scaling down the probabilities of all tuples in D with a factor $f < 1$. Then the relative error of approximation of $\mathbb{P}(q)$ by the propagation score $\rho(q)$ decreases as f goes to 0: $\lim_{f \rightarrow 0} \frac{\rho(q) - \mathbb{P}(q)}{\mathbb{P}(q)} \rightarrow 0$.*

Number of Dissociations. While the number of minimal safe dissociations is exponential in the size of the query, recall that it is independent of the size of the database. Figure 2 gives an overview of the number of minimal query plans, total query plans, and all dissociations for k -star and k -chain queries (which are later used in Section 5). Later Section 4 gives optimizations that allow us to evaluate a large number of plans efficiently.

3.3 Minimal plans with schema knowledge

Next, we show how knowledge of *deterministic relations* (i.e. all tuples have probability = 1), and *functional dependencies* can reduce the number of plans needed to calculate the propagation score.

k -star query				k -chain query			
k	#MP	#P	# Δ	k	#MP	#P	# Δ
1	1	1	1	2	1	1	1
2	2	3	4	3	2	3	4
3	6	13	64	4	5	11	64
4	24	75	4096	5	14	45	4096
5	120	541	$> 10^6$	6	42	197	$> 10^6$
6	720	4683	$> 10^9$	7	132	903	$> 10^9$
7	5040	47293	$> 10^{12}$	8	429	4279	$> 10^{12}$
seq	$k!$	A000670	$2^{k(k-1)}$	seq	A000108	A001003	$2^{(k+1)k}$

Figure 2: Number of minimal plans, total plans, and total dissociations for star and chain queries (A are OEIS sequence numbers [37]).

3.3.1 Deterministic relations (DRs)

Notice that we can treat deterministic relations (DRs) just like probabilistic relations, and Corollary 19 with $\mathbb{P}(q) \leq \text{score}(P)$ still holds for any plan P . Just as before, our goal is to find a minimum number of plans that compute the minimal score of *all plans*: $\rho(q) = \min_P \text{score}(P)$. It is known that an unsafe query q can become safe (i.e., $\mathbb{P}(q)$ can be calculated in PTIME with one single plan) if we consider DRs. Thus, in particular, we would still like an improved algorithm that returns one single plan if a query with DRs is safe. The following lemma will help us achieve this goal:

LEMMA 22 (DISSOCIATION AND DRs). *Dissociating a deterministic relation does not change the probability.*

PROOF. Lemma 22 follows immediately from Theorem 8 (2) and noting that dissociating tuples in DRs corresponds exactly to dissociating variables \mathbf{X} with $p(X_i) = 1$. \square

We thus define a new *probabilistic dissociation preorder* \preceq^p by:

$$\Delta \preceq^p \Delta' \Leftrightarrow \forall i, R_i \text{ probabilistic} : \mathbf{y}_i \subseteq \mathbf{y}'_i$$

In other words, $\Delta \preceq^p \Delta'$ still implies $\mathbb{P}(q^\Delta) \leq \mathbb{P}(q^{\Delta'})$, but \preceq^p is defined on probabilistic relations only. Notice, that for queries without DRs, the relations \preceq^p and \preceq coincide. However, for queries with DRs, \preceq^p is a preorder, not an order. Therefore, there exist distinct dissociations Δ, Δ' that are equivalent under \preceq^p (written as $\Delta \equiv^p \Delta'$), and thus have the same probability: $\mathbb{P}(q^\Delta) = \mathbb{P}(q^{\Delta'})$. As a consequence, using \preceq^p instead of \preceq , allows us to further reduce the number of minimal safe dissociations.

EXAMPLE 23 (DRs). *Consider $q: -R(x), S(x, y), T^d(y)$ where a d -exponent indicates a DR. This query is known to be safe. We thus expect our definition of $\rho(q)$ to find that $\rho(q) = \mathbb{P}(q)$. Ignore that T^d is deterministic, then \preceq has two minimal plans: $q^{\Delta_1}: -R^y(x, y), S(x, y), T^d(y)$, and $q^{\Delta_2}: -R(x), S(x, y), T^{dx}(x, y)$. Since Δ_2 dissociates only T^d , we now know from Lemma 22 that $\mathbb{P}(q) = \mathbb{P}(q^{\Delta_2})$. Thus, by using \preceq as before, we still get the correct answer. However, evaluating the plan P^{Δ_1} is always unnecessary since $\Delta_2 \preceq^p \Delta_1$. In contrast, without information about DRs, $\Delta_2 \not\preceq^p \Delta_1$, and we would thus have to evaluate both plans.*

Figure 3 illustrates this with augmented incidence matrices: dissociated variables in DRs are now marked with empty circles (\circ) instead of full circles (\bullet), and the preorder \preceq^p is determined entirely by full circles (representing dissociated variables in probabilistic relations). However, as before, the correspondence to plans (as implied by the hierarchy between all variables) is still determined by empty and full circles. Figure 3b shows that $\rho(q) = \mathbb{P}(q^{\Delta_2}) = \mathbb{P}(q)$ since $\Delta_0 \equiv^p \Delta_2 \preceq^p \Delta_1 \equiv \Delta_3$. Thus, the query is safe, and it suffices to evaluate only P^{Δ_2} . Notice that q is not hierarchical, but still safe since it is in an equivalence class with a query that is hierarchical:

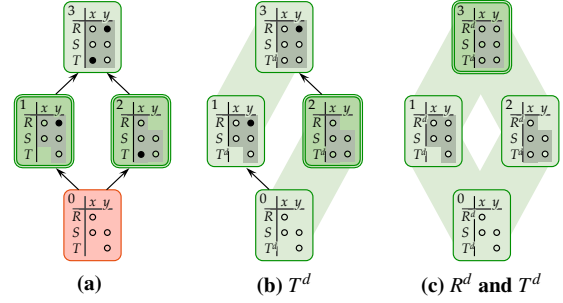


Figure 3: Example 23: The presence of DRs R^d and T^d in (b) and (c) changes the original partial dissociation order for $q: -R(x), S(x, y), T(y)$ in (a): Several dissociations now have the same probability (shown with shaded areas instead of arrows). Our modified algorithm now returns, for each minimal safe equivalence class, the query plan for the top most dissociation (shown in dark green and double-lined).

$\Delta_0 \equiv^p \Delta_2$. Figure 3c shows that, with R^d and T^d being deterministic, all three possible query plans (corresponding to Δ_1, Δ_2 , and Δ_3) form a “minimal equivalence class” in \preceq^p with Δ_0 , and thus give the exact probability. We, therefore, want to modify our algorithm to return just one plan from each “minimal safe equivalence class.” Ideally, we prefer the plan corresponding to Δ_3 (or more generally, the top plan in \preceq for each minimum equivalence class) since P^{Δ_3} least constrains the join order between tables.

We now explain two simple modifications to Algorithm 1 that achieve exactly our desired optimizations described above:

- (1) Denote with $\text{MinPCuts}(q)$ the set of minimal cut-sets that disconnect the query into at least two connected components with probabilistic tables. Replace $\text{MinCuts}(q)$ in line 10 with $\text{MinPCuts}(q)$.
- (2) Denote with m_p the number of probabilistic relations in a query. Replace the stopping condition in line 1 with: **if** $m^p \leq 1$ **then** $\mathcal{P} \leftarrow \{\pi_{\mathbf{x}} \bowtie^p [R_1(\mathbf{x}_1), \dots, R_m(\mathbf{x}_m)]\}$. In other words, if a query has maximal one probabilistic relation, then join all relations followed by projecting on the head variables.

THEOREM 24 (ALGORITHM 1 WITH DRs). *Algorithm 1 with above 2 modifications returns a minimum number of plans to calculate $\rho(q)$ given schema knowledge about DRs.*

For example, for $q: -R(x), S(x, y), T^d(y)$, $\text{MinCuts}(q) = \{\{x\}, \{y\}\}$, while $\text{MinPCuts}(q) = \{\{x\}\}$. Therefore, the modified algorithm returns P^{Δ_2} as single plan. For $q: -R^d(x), S(x, y), T^d(y)$, the stopping condition is reached (also, $\text{MinPCuts}(q) = \{\emptyset\}$) and the algorithm returns P^{Δ_3} as single plan (see Fig. 3c).

3.3.2 Functional dependencies (FDs)

Knowledge of functional dependencies (FDs), such as keys, can also restrict the number of necessary minimal plans. A well known example is the query $q: -R(x), S(x, y), T(y)$ from Example 23; it becomes safe if we know that S satisfies the FD $\Gamma: x \rightarrow y$ and has a unique safe plan that corresponds to dissociation Δ_2 . In other words, we would like our modified algorithm to take Γ into account and to not return the plan corresponding to dissociation Δ_1 .

Let Γ be the set of FDs on $\text{Var}(q)$ consisting of the union of FDs on every atom R_i in q . As usual, denote \mathbf{x}_i^+ the closure of a set of attributes \mathbf{x}_i , and denote $\Delta_\Gamma = (\mathbf{y}_1, \dots, \mathbf{y}_m)$ the dissociation defined as follows: for every atom $R_i(\mathbf{x}_i)$ in q , $\mathbf{y}_i = \mathbf{x}_i^+ \setminus \mathbf{x}_i$. Then we show:

LEMMA 25 (DISSOCIATION AND FDs). *Dissociating a table R_i on any variable $y \in \mathbf{x}_i^+$ does not change the probability.*

This lemma is similar to Lemma 22. We can thus further refine our probabilistic dissociation preorder $\preceq^{p'}$ by:

$$\Delta \preceq^{p'} \Delta' \Leftrightarrow \forall i, R_i \text{ probabilistic} : \mathbf{y}_i \setminus \mathbf{x}_i^+ \subseteq \mathbf{y}'_i \setminus \mathbf{x}'_i^+$$

As a consequence, using $\preceq^{p'}$ instead of \preceq^p , allows us to further reduce the number of minimal safe equivalence classes. We next state a result by [39] in our notation:

PROPOSITION 26 (SAFETY AND FDS [39, PROP. IV.5]). *A query q is safe iff q^{Δ_Γ} is hierarchical.*

This justifies our third modification to Algorithm 1 for computing $\rho(q)$ of a query q over a database that satisfies Γ : First compute Δ_Γ , then run q^{Δ_Γ} on our previously modified Algorithm 1.

THEOREM 27 (ALGORITHM 1 WITH FDS). *Algorithm 1 with above 3 modifications returns a minimum number of plans to calculate $\rho(q)$ given schema knowledge about DRs and FDs.*

It is easy to see that our modified algorithm returns one single plan iff the query is safe, taking into account its structure, DRs and FDs. It is thus a *strict generalization of all known safe self-join-free conjunctive queries* [7, 39]. In particular, we can reformulate the known safe query dichotomy [7] in our notation very succinctly:

COROLLARY 28 (DICHOTOMY). $\mathbb{P}(q)$ can be calculated in PTIME iff there exists a dissociation Δ of q that is (i) hierarchical, and (ii) in an equivalence class with q under $\preceq^{p'}$.

To see what the corollary says, assume first that there are no FDs: Then q is in PTIME iff there exists a dissociation Δ of the DRs only, such that q^Δ is hierarchical. If there are FDs, then we first compute the full dissociation Δ_Γ (called “full chase” in [39]), then apply the same criterion to q^{Δ_Γ} .

4. MULTI-QUERY OPTIMIZATIONS

So far, Algorithm 1 enumerates all minimal query plans. We then take the minimum score of those plans in order to calculate the propagation score $\rho(q)$. In this section, we develop three optimizations that can considerably reduce the necessary calculations for evaluating all minimal query plans. Note that these three optimizations and the two optimizations from the previous section are orthogonal and can be arbitrarily combined in the obvious way. We use the following example to illustrate the first two optimizations.

EXAMPLE 29 (OPTIMIZATIONS). *Consider $q: - R(x,z), S(y,u), T(z), U(u), M(x,y,z,u)$. Our default is to evaluate all 6 minimal plans returned by Algorithm 1, then take the minimum score (shown in Fig. 4a). Figure 4b and Fig. 4c illustrate the optimized evaluations after applying Opt. 1, or Opt. 1 and Opt. 2, respectively.* ■

4.1 Opt. 1: One single query plan

Our first optimization creates one single query plan by *pushing the min-operator down into the leaves*. It thus avoids calculations when it is clear that other calculations must have lower bounds. The idea is simple: Instead of creating one query subplan for each top set $\mathbf{y} \in \text{MinCuts}(q)$ in line 12 of Algorithm 1, the adapted Algorithm 2 takes the minimum score over those top sets, for each tuple of the head variables in line 11. It thus creates one single query plan.

4.2 Opt. 2: Re-using common subplans

Our second optimization calculates only once, then *re-uses common subplans shared between the minimal plans*. Thus,

$$\rho(q) = \min \left[\begin{array}{l} \pi_x^p \bowtie^p [T(z)] \pi_x^p \bowtie^p [R(x,z), \pi_u^p \bowtie^p [U(u) \pi_y^p \bowtie^p [S(y,u), M(x,y,z,u)]]] \\ \pi_z^p \bowtie^p [T(z)] \pi_u^p \bowtie^p [U(u), \pi_x^p \bowtie^p [R(x,z) \pi_y^p \bowtie^p [S(y,u), M(x,y,z,u)]]] \\ \pi_x^p \bowtie^p [T(z)] \pi_u^p \bowtie^p [U(u), \pi_y^p \bowtie^p [S(y,u), \pi_z^p \bowtie^p [T(z), \pi_x^p \bowtie^p [R(x,z), M(x,y,z,u)]]]] \\ \pi_u^p \bowtie^p [U(u), \pi_y^p \bowtie^p [S(y,u), \pi_z^p \bowtie^p [T(z), \pi_x^p \bowtie^p [R(x,z), M(x,y,z,u)]]]] \\ \pi_u^p \bowtie^p [U(u), \pi_z^p \bowtie^p [T(z), \pi_x^p \bowtie^p [R(x,z), \pi_y^p \bowtie^p [S(y,u), M(x,y,z,u)]]]] \\ \pi_u^p \bowtie^p [U(u), \pi_z^p \bowtie^p [T(z), \pi_y^p \bowtie^p [S(y,u), \pi_x^p \bowtie^p [R(x,z), M(x,y,z,u)]]]] \end{array} \right]$$

(a) Result from Algorithm 1: six minimal query plans

$$\rho(q) = \min \left[\begin{array}{l} \pi_z^p \bowtie^p [T(z)] \min \left[\begin{array}{l} \pi_x^p \bowtie^p [R(x,z), \pi_u^p \bowtie^p [U(u) \pi_y^p \bowtie^p [S(y,u), M(x,y,z,u)]]] \\ \pi_u^p \bowtie^p [U(u), \min \left[\begin{array}{l} \pi_x^p \bowtie^p [R(x,z), \pi_y^p \bowtie^p [S(y,u), M(x,y,z,u)]] \\ \pi_y^p \bowtie^p [S(y,u), \pi_x^p \bowtie^p [R(x,z), M(x,y,z,u)]] \end{array} \right] \end{array} \right] \\ \pi_u^p \bowtie^p [U(u), \min \left[\begin{array}{l} \pi_y^p \bowtie^p [S(y,u), \pi_z^p \bowtie^p [T(z), \pi_x^p \bowtie^p [R(x,z), M(x,y,z,u)]]] \\ \pi_x^p \bowtie^p [R(x,z), \min \left[\begin{array}{l} \pi_y^p \bowtie^p [S(y,u), M(x,y,z,u)] \\ \pi_z^p \bowtie^p [T(z), \min \left[\begin{array}{l} \pi_x^p \bowtie^p [R(x,z), \pi_y^p \bowtie^p [S(y,u), M(x,y,z,u)]] \\ \pi_y^p \bowtie^p [S(y,u), \pi_x^p \bowtie^p [R(x,z), M(x,y,z,u)]] \end{array} \right] \end{array} \right] \end{array} \right] \end{array} \right] \end{array} \right]$$

(b) Result from Algorithm 4: one single query plan

$$\begin{aligned} V_1(x,z,u) &= \pi_y^p \bowtie^p [S(y,u), M(x,y,z,u)] \\ V_2(y,z,u) &= \pi_x^p \bowtie^p [R(x,z), M(x,y,z,u)] \\ V_3(z,u) &= \min \left[\begin{array}{l} \pi_x^p \bowtie^p [R(x,z), V_1(x,z,u)] \\ \pi_y^p \bowtie^p [S(y,u), V_2(y,z,u)] \end{array} \right] \\ \rho(q) &= \min \left[\begin{array}{l} \pi_z^p \bowtie^p [T(z)] \min \left[\begin{array}{l} \pi_x^p \bowtie^p [R(x,z), \pi_u^p \bowtie^p [U(u), V_1(x,z,u)]] \\ \pi_u^p \bowtie^p [U(u), V_3(z,u)] \end{array} \right] \\ \pi_u^p \bowtie^p [U(u), \min \left[\begin{array}{l} \pi_y^p \bowtie^p [S(y,u), \pi_z^p \bowtie^p [T(z), V_2(y,z,u)]] \\ \pi_z^p \bowtie^p [T(z), V_3(z,u)] \end{array} \right] \end{array} \right] \end{aligned}$$

(c) Result from Algorithm 5: re-using common subplans

Figure 4: Example 29 before and after applying optimizations 1 and 2.

```

Recursive algorithm: SP (SinglePlan)
Input: Query  $q(\mathbf{x}) : - R_1(\mathbf{x}_1), \dots, R_m(\mathbf{x}_m)$ 
Output: Single query plan  $P$ 
1 if  $m = 1$  then  $P \leftarrow \pi_{\mathbf{x}}^p R_1(\mathbf{x}_1)$ 
2 else
3   if  $q$  is disconnected then
4     Let  $q = q_1, \dots, q_k$  be the components connected by  $\text{EVar}(q)$ 
5     Let  $\text{HVar}(q_i) \leftarrow \text{HVar}(q) \cap \text{Var}(q_i)$ 
6      $P \leftarrow \bowtie^p [\text{SP}(q_1), \dots, \text{SP}(q_k)]$ 
7   else
8     Let  $\text{MinCuts}(q) = \{\mathbf{y}_1, \dots, \mathbf{y}_j\}$ 
9     Let  $q'_i \leftarrow q_i$  with  $\text{HVar}(q'_i) \leftarrow \text{HVar}(q) \cup \mathbf{y}_i$ 
10    if  $j = 1$  then  $P \leftarrow \pi_{\mathbf{y}_1}^p \text{SP}(q'_1)$ 
11    else  $P \leftarrow \min [\pi_{\mathbf{y}_1}^p \text{SP}(q'_1), \dots, \pi_{\mathbf{y}_j}^p \text{SP}(q'_j)]$ 

```

Algorithm 2: Optimization 1 recursively pushes the min operator into the leaves and generates one single query plan.

whereas our first optimization reduces computation by combining plans at their roots, the second optimization stores and re-uses common results in the branches. The adapted Algorithm 3 works as follows: It first traverses the whole single query plan (FindingCommonSubplans) and remembers each subplan by the atoms used and its head variables in a HashSet HS (line 14). If it sees a subplan twice (line 13), it creates a new view for this subplan, mapping the subplan to a new view definition. The actual plan (ViewReusingPlan) then uses these views whenever possible (line 18). The order in which the views are created (line 5) assures that the algorithm also discovers and exploits *nested common subexpressions*. Figure 4c illustrates for Example 29, that both the main plan and the view V_3 re-use views V_1 and V_2 .

4.3 Opt. 3: Deterministic semi-join reduction

The most expensive operations in probabilistic query plans are the group-bys for the probabilistic project operations. These are often applied early in the plans to tuples which are later pruned and do not

```

Algorithm: UsingCommonSubplans
Input: Query  $q(\mathbf{x}) :- R_1(\mathbf{x}_1), \dots, R_m(\mathbf{x}_m)$ 
Output: Ordered set of view definitions  $\mathcal{V}$ , final query plan  $P$ 
1 HS  $\leftarrow \emptyset$  // HashSet of all subplans
2 HM  $\leftarrow (\emptyset, \emptyset)$  // HashMap from subplans to unique view names
3  $\mathcal{V} \leftarrow \emptyset$  // Set of view definitions
4 FS( $q$ )
5 foreach  $q_i \in \text{HM.keys}$  in increasing size of HVar( $q_i$ ) and Var( $q_i$ ) do
6    $\mathcal{V} \leftarrow \mathcal{V} \cup \{\text{HM.val} = \text{ViewReusingPlan}(q_i)\}$ 
7  $P = \text{RP}(q)$ 

Recursive function: FS (FindingCommonSubplans)
Input: Query  $q(\mathbf{x}) :- R_1(\mathbf{x}_1), \dots, R_m(\mathbf{x}_m)$ 
8 if  $q$  is disconnected then
9   Let  $q = q_1, \dots, q_k$  be the components connected by EVar( $q$ )
10  foreach  $q_i$  do FS( $q_i(\mathbf{x}_i)$ )
11 else
12   if  $(m = 1 \wedge \mathbf{x} = \mathbf{x}_i) \vee \text{HM}(q) \neq \emptyset$  then return
13   if  $q \in \text{HS} \wedge \text{HM}(q) = \emptyset$  then  $\text{HM}(q) \leftarrow$  new view name
14    $\text{HS} \leftarrow \text{HS} \cup \{q\}$ 
15   foreach  $y \in \text{MinCuts}(q)$  do
16     Let  $q' \leftarrow q$  with HVar( $q'$ )  $\leftarrow$  HVar( $q$ )  $\cup y$ 
17     FS( $q'$ )

Recursive function: RP (ViewReusingPlan)
Input: Query  $q(\mathbf{x}) :- R_1(\mathbf{x}_1), \dots, R_m(\mathbf{x}_m)$ 
Output: Query plan  $P$  that reuses views from HashMap HM
18 if  $\text{HM}(q) \neq \emptyset$  then  $P \leftarrow \text{HM}(q)$ 
19 else
20    $\leftarrow$  Insert here lines 1-11 from Algorithm 2, replacing SP with RP

```

Algorithm 3: Optimizations 1 & 2 together create a query plan which re-uses several previously defined temporary views.

contribute to the final query result. Our third optimization is to first apply a *full semi-join reduction on the input relations* before starting the probabilistic evaluation from these *reduced input relations*. We like to draw here an important connection to [39], which introduces the idea of “lazy plans” and shows orders of magnitude performance improvements for safe plans by computing confidences not after each join and projection, but rather at the very end of the plan. We note that our semi-join reduction *serves the same purpose* with similar performance improvements and also apply for safe queries. The advantage of semi-join reductions, however, is that we do not require any modifications to the query engine.

5. EXPERIMENTS

We are interested in both the efficiency (“how fast?”) and the quality (“how good?”) of ranking by dissociation as compared to exact probabilistic inference, Monte Carlo simulation (MC), and standard deterministic query evaluation (“deterministic SQL”).

Ranking quality. We use *mean average precision* (MAP) to evaluate the quality of a ranking by comparing it against the ranking from exact probabilistic inference as ground truth (GT). MAP rewards rankings that place relevant items earlier; the best possible value is 1, and the worst possible 0. We use a variant of “Average Precision at 10” defined as $\text{AP}@10 := \frac{\sum_{k=1}^{10} \text{P}@k}{10}$. Here, $\text{P}@k$ is the precision at the k th answer, i.e., the fraction of top k answers according to GT that are also in the top k answers returned. Averaging over several experiments yields MAP [34]. We use a variant of the analytic method proposed in [35] to calculate AP in the presence of ties. As baseline for no ranking, we assume all tuples have the same score and are thus tied for the same position. We call this baseline “*random average precision*.”

Exact probabilistic inference. Whenever possible, we calculate GT rankings with a tool called SampleSearch [23, 47], which also serves to evaluate the cost of *exact* probabilistic inference. We de-

scribe the method of transforming the lineage DNF into a format that can be read by SampleSearch in [22].

Monte Carlo (MC). We evaluate the MC simulations for different numbers of samples and write $\text{MC}(x)$ for x samples. For example, AP for $\text{MC}(10\text{k})$ is the result of sampling the individual tuple scores 10 000 times from their lineages and then evaluating AP once over the sampled scores. The MAP scores together with the standard deviations are then the average over several repetitions.

Ranking by lineage size. To evaluate the potential of non-probabilistic methods for ranking answers, we also rank the answer tuples by decreasing size of their lineages; i.e. number of terms. Intuitively, a larger lineage size should indicate that an answer tuple has more “support” and should thus be more important.

Setup 1. We use the TPC-H DBGEN data generator [54] to generate a 1GB database to which we add a column P for each table and store it in PostgreSQL 9.2 [43]. We assign to each input tuple i a random probability p_i uniformly chosen from the interval $[0, p_{i\max}]$, resulting in an expected average input probability $\text{avg}[p_i] = p_{i\max}/2$. By using databases with $\text{avg}[p_i] < 0.5$, we can avoid output probabilities close to 1 for queries with very large lineages. We use the following parameterized query:

$$Q(a) :- S(\underline{s}, a), PS(\underline{s}, u), P(\underline{u}, n), s \leq \$1, n \text{ like } \$2$$

select distinct s_nationkey from Supplier, Partsupp, Part
 where s_suppkey = ps_suppkey and ps_partkey = p_partkey
 and s_suppkey <= \$1 and p_name like \$2

Parameters \$1 and \$2 allow us to change the lineage size. Tables Supplier, Partsupp and Part have 10k, 800k and 200k tuples, respectively. There are 25 different numeric attributes for nationkey and our goal is to efficiently rank these 25 nations. As baseline for not ranking, we use random average precision for 25 answers, which leads to $\text{MAP}@10 \approx 0.220$. This query has two minimal query plans and we will compare the speed-up from either evaluating both individually or performing a deterministic semi-join reduction (Optimization 3) on the input tables.

Setup 2. We compare the run times for our three optimizations against evaluation of all plans for k -chain queries and k -star queries over varying database sizes (to evaluate data complexity) and varying query sizes (to evaluate query complexity):

$$\begin{aligned} k\text{-chain: } & q(x_0, x_k) :- R_1(x_0, x_1), R_2(x_1, x_2), \dots, R_k(x_{k-1}, x_k) \\ k\text{-star: } & q(a') :- R_1(a', x_1), R_2(x_2), \dots, R_k(x_k), R_0(x_1, \dots, x_k) \end{aligned}$$

We denote the length of the query with k , the number of tuples per table with n , and the domain size with N . We use integer values which we uniformly draw from the range $\{1, 2, \dots, N\}$. Thus, the parameter N determines the *selectivity* and is varied as to keep the answer cardinality constant around 20-50 for chain queries, or the answer probability between 0.90 and 0.95 for star queries. For the data complexity experiments, we vary the number of tuples n per table between 100 and 10^6 . For the query complexity experiments, we vary k between 2 and 8 for chain queries. For these experiments, the optimized (and often *extremely long*) SQL statements are “calculated” in JAVA and then sent to Microsoft SQL server 2012. To illustrate with numbers, we have to issue 429 query plans in order to evaluate the 8-chain query (see Fig.2). Each of these plans joins 8 tables in a different order. Optimization 1 then merges those plans together into one truly gigantic single query plan.

5.1 Run time experiments

QUESTION 1. *When and how much do our three query optimizations speed up query evaluation?*

Result 1. Combining plans (Opt. 1) and using intermediate views (Opt. 2) almost always speeds up query times. The semi-join reduction (Opt. 3) slows down queries with high selectivities, but considerably speeds up queries with small selectivities.

Figures 5a to 5d show the results on setup 2 for increasing database sizes or query sizes. For example, Fig.5b shows the performance of computing a 7-chain query which has 132 safe dissociations. Evaluating each of these queries separately takes a long time, while our optimization techniques bring evaluation time close to deterministic query evaluation. Especially on larger databases, where the running time is I/O bound, the penalty of the probabilistic inference is only a factor of 2-3 in this example. Notice here the trade-off between optimization 1,2 and optimization 1,2,3: Optimization 3 applies a full semi-join reduction on the input relations before starting the probabilistic plan evaluation from these reduced input relations. This operation imposes a rather large constant overhead, both at the query optimizer and at query execution. For larger databases (but constant selectivity), this overhead is amortized. In practice, this suggests that dissociation allows us a large space of optimizations depending on the query and particular database instance that can conservatively extend the space of optimizations performed today in deterministic query optimizers.

Figures 5e to 5g compare the running times on setup 1 between dissociation with two minimal query plans (“Diss”), dissociation with semi-join reduction (“Diss + Opt3”), exact probabilistic inference (“SampleSearch”), Monte Carlo with 1000 samples (“MC(1k)”), retrieving the lineage only (“Lineage query”), and deterministic query evaluation without ranking (“Standard SQL”). We fixed $\$2 \in \{\text{'red\%green\%'}, \text{'red\%'}, \text{'\%'}\}$ and varied $\$1 \in \{500, 1000, \dots, 10k\}$. Figure 5h combines all three previous plots and shows the times as function of the maximum lineage size (i.e. the size of the lineage for the tuple with the maximum lineage) of a query. We see here again that the semi-join reduction speeds up evaluation considerably for small lineage sizes (Fig. 5e shows speedups of up to 36). For large lineages, however, the semi-join reduction is an unnecessary overhead, as most tuples are participating in the join anyway (Fig. 5f shows overhead of up to 2).

QUESTION 2. *How does dissociation compare against other probabilistic methods and standard query evaluation?*

Result 2. The best evaluation strategy for dissociation takes only a small overhead over standard SQL evaluation and is considerably faster than other probabilistic methods for large lineages.

Figures 5d to 5h show that SampleSearch does not scale to larger lineages as the performance of exact probabilistic inference depends on the tree-width of the Boolean lineage formula, which generally increases with the size of the data. In contrast, dissociation is independent of the treewidth. For example, SampleSearch needed 780 sec for calculating the ground truth for a query with $\max[\text{lin}] = 5.9k$ for which dissociation took 3.0 sec, and MC(1k) took 42 sec for a query with $\max[\text{lin}] = 4.2k$ for which dissociation took 2.4 sec. Dissociation takes only 10.5 sec for our largest query $\$2 = \text{'\%'}$ and $\$1 = 10k$ with $\max[\text{lin}] = 35k$. Retrieving the lineage for that query alone takes 5.8 sec, which implies that any probabilistic method that evaluates the probabilities outside of the database engine needs to issue this query to retrieve the DNF for each answer and would thus have to evaluate lineages of sizes around 35k in only 4.7 (= 10.5 - 5.8) sec to be faster than dissociation.⁵

⁵The time needed for the lineage query thus serves as minimum benchmark for any probabilistic approximation. The reported times for SampleSearch and MC are the sum of time for retrieving the lineage plus the actual calculations, without the time for reading and writing the input and output files for SampleSearch.

5.2 Ranking experiments

For the following experiments, we are limited to those query parameters $\$1$ and $\$2$ for which we can get the ground truth (and results from MC) in acceptable time. We systematically vary $p_{i\max}$ between 0.1 and 1 (and thus $\text{avg}[p_i]$ between 0.05 and 0.5) and evaluate the rankings several times over randomly assigned input tuple probabilities. We only keep data points (i.e. results of individual ranking experiments) for which the output probabilities are not too close to 1 to be meaningful ($\max[p_a] < 0.999999$).

QUESTION 3. *How does ranking quality compare for our three ranking methods and which are the most important factors that determine the quality for each method?*

Result 3. Dissociation performs better than MC which performs better than ranking by lineage size.

Figure 5i shows averaged results of our probabilistic methods for $\$2 = \text{'red\%green\%'}$.⁶ Shaded areas indicate standard deviations and the x-axis shows varying numbers of MC samples. We only used those data points for which $\text{avg}[p_a]$ of the top 10 ranked tuples is between 0.1 and 0.9 according to ground truth ($\approx 6k$ data points for dissociation and lineage, $\approx 60k$ data points for MC, as we repeated each MC simulation 10 times), as this is the best regime for MC, according to Result 4. We also evaluated quality for dissociation and ranking by lineage for more queries by choosing parameter values for $\$2$ from a set of 28 strings, such as $\text{'r\%g\%r\%a\%n\%d\%'}$ and 're\%re\%' . The average MAP over all 28 choices for parameters $\$2$ is 0.997 for ranking by dissociation and 0.520 for ranking by lineage size ($\approx 100k$ data points). Most of those queries have too large of a lineage to evaluate MC. Note that ranking by lineage always returns the same ranking for given parameters $\$1$ and $\$2$, but the GT ranking would change with different input probabilities.

Result 4. Ranking quality of MC increases with the number of samples and decreases when the average probability of the answer tuples $\text{avg}[p_a]$ is close to 0 or 1.

Figure 5j shows the AP as a function of $\text{avg}[p_a]$ of the top 10 ranked tuples according to ground truth by logarithmic scaling of the x-axis (each point in the plot averages AP over ≈ 450 experiments for dissociation and lineage and over $\approx 4.5k$ experiments for MC). We see that MC performs increasingly poor for ranking answer tuples with probabilities close to 0 or 1 and even approach the quality of random ranking (MAP@10 = 0.22). This is so because, for these parameters, the probabilities of the top 10 answers are very close, and MC needs many iterations to distinguish them. Therefore, MC performs increasingly poorly for increasing size of lineage but fixed average input probability $\text{avg}[p_i] \approx 0.5$, as the average answer probabilities $\text{avg}[p_a]$ will be close to 1. In order not to “bias against our competitor,” we compared against MC in its best regime with $0.1 < \text{avg}[p_a] < 0.9$ in Fig. 5i.

Result 5. Ranking by lineage size has good quality only when all input tuples have the same probability.

Figure 5k shows that ranking by lineage is good only when all tuples in the database have the same probability (labeled by $p_i = \text{const}$ as compared to $\text{avg}[p_i] = \text{const}$). This is a consequence of the output probabilities depending mostly on the size of the lineages if all probabilities are equal. Dependence on other parameters, such as overall lineage size and magnitude of input probabilities (here shown for $p_i = 0.1$ and $p_i = 0.5$), seem to matter only slightly.

⁶Results for MC with other parameters of $\$2$ are similar. However, the evaluation time for the experiments becomes quickly infeasible.

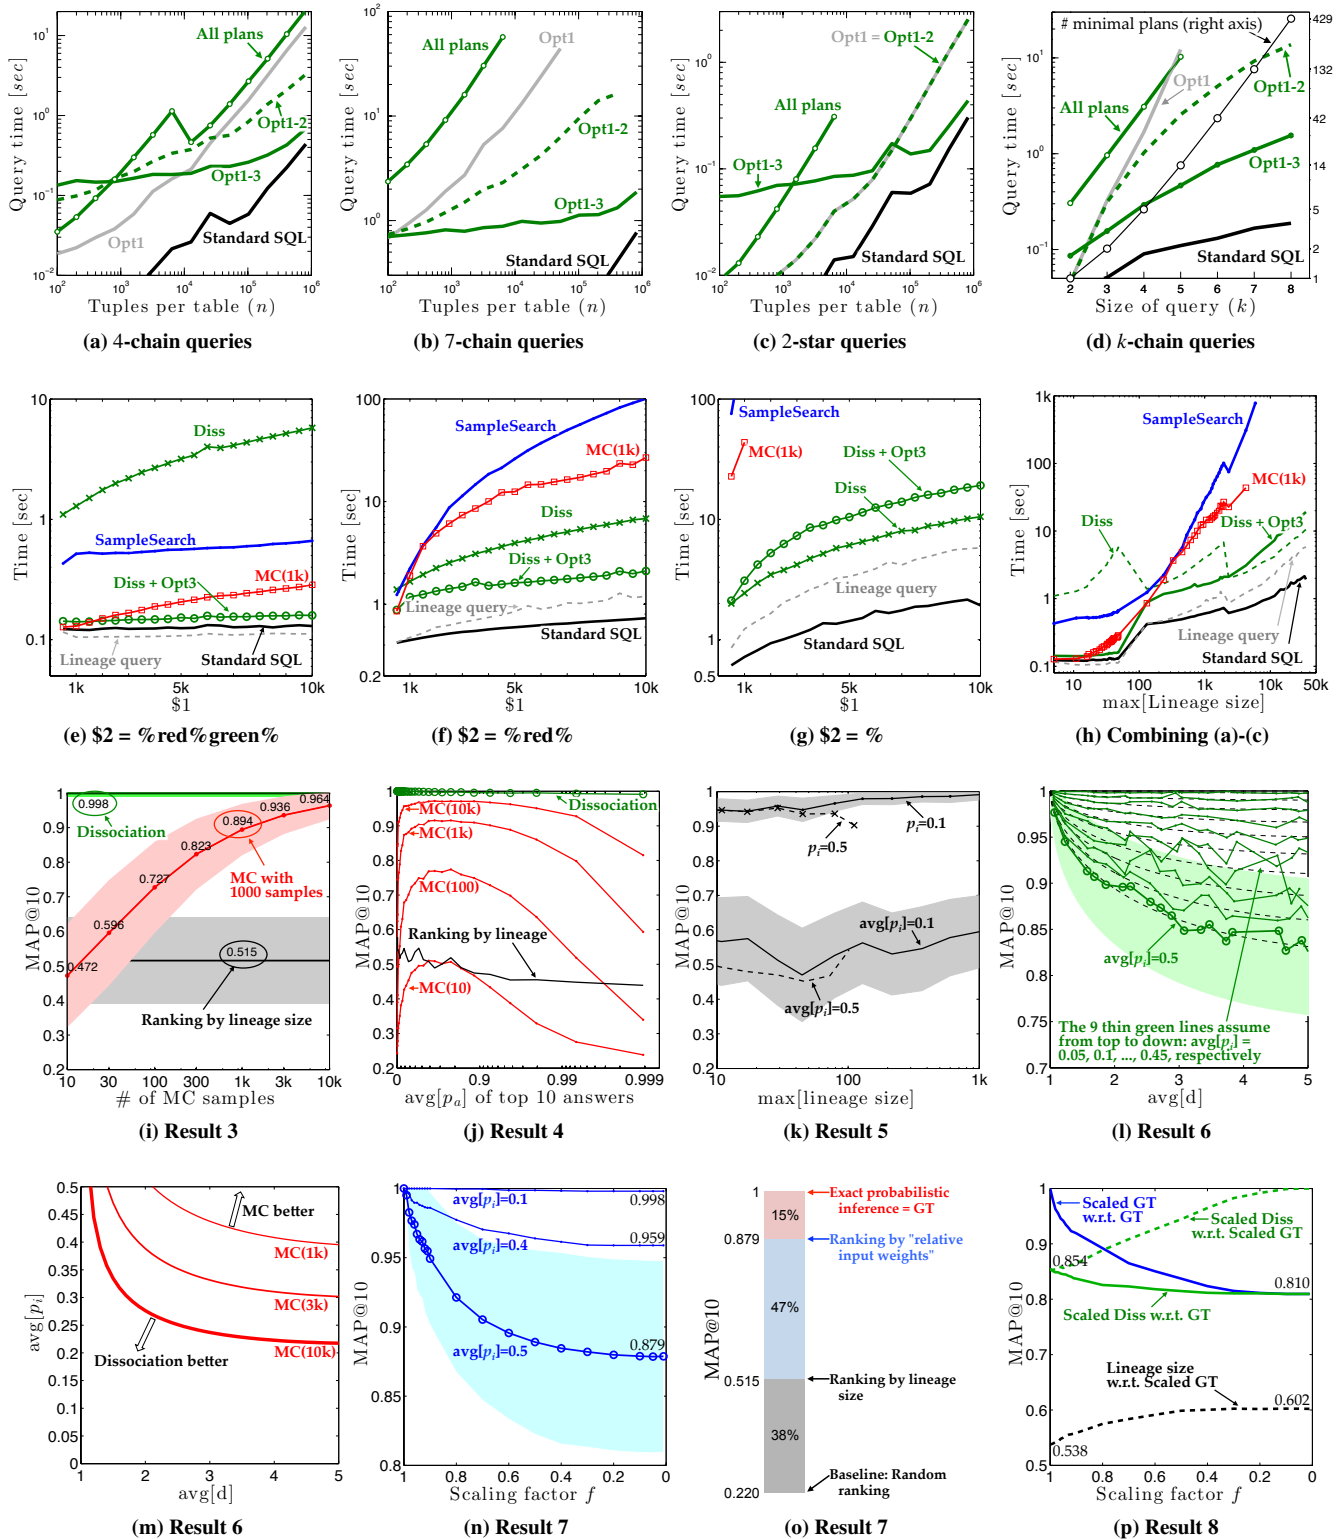


Figure 5: Timing results: (a)-(c) For increasing database sizes and constant cardinalities, our optimizations approach deterministic SQL performance. (d) Our optimizations can even evaluate very large number of minimal plans efficiently (here shown up to 429 for an 8-chain query). (e)-(h) For the TPC-H query, the best evaluation for dissociation is within a factor of 6 of that for deterministic query evaluation. (i)-(p) Ranking experiments on TPC-H: Assumptions for each subfigure and conclusions that can be drawn are described in the main text in the respective result paragraph.

Result 6. The quality of dissociation decreases with the average number of dissociations per tuple $\text{avg}[d]$ and with the average input probabilities $\text{avg}[p_i]$. Dissociation performs very well and notably better than MC(10k) if either $\text{avg}[d]$ or $\text{avg}[p_i]$ are small.

Each answer tuple a gets its score p_a from one of two query plans P_S and P_P that dissociate tuples in tables S and P , respectively. For example, if the lineage size for tuple a is 100 and the lineage contains 20 unique suppliers from table S and 50 unique parts from table P , then P_S dissociates each tuple from S into 5 tuples and P_P each tuple from P into 2 tuples, on average. Most often, P_P will then give the better bounds as it has fewer average dissociations. Let $\text{avg}[d]$ be the mean number of dissociations for each tuple in the dissociated table of its respective optimal query plan, averaged across all top 10 ranked answer tuples. For all our queries (even those with $\$1 = 10k$ and $\$2 = \%$), $\text{avg}[d]$ stays below 1.1 as, for each tuple, there is usually one plan that dissociates few variables. In order to understand the impact of higher numbers of dissociations (increasing $\text{avg}[d]$), we also measured AP for the ranking for each query plan individually. Hence, for each choice of random parameters, we record two new data points – one for ranking all answer tuples by using only P_S and one for using only P_P – together with the values of $\text{avg}[d]$ in the respective table that gets dissociated. This allows us to draw conclusions for a larger set of parameters. Figure 5l plots MAP values as a function of $\text{avg}[d]$ of the top 10 ranked tuples on the horizontal axis, and various values of $\text{avg}[p_i]$ ($\text{avg}[p_i] = 0.05, 0.10, \dots, 0.5$). Each plotted point averages over at least 10 data points (some have 10, other several 1000s). Dashed lines show a fitted parameterized curve to the data points on $\text{avg}[p_i]$ and $\text{avg}[d]$. The figure also shows the standard deviations as shaded areas for $\text{avg}[p_i] = 0.5$. We see that the quality is very dependent on $\text{avg}[p_i]$, as predicted by Prop. 21.

Figure 5m maps the trade-off between dissociation and MC for the two important parameters for the quality of dissociation ($\text{avg}[d]$ and $\text{avg}[p_i]$) and the number of samples for MC. For example, MC(1k) gives a better expected ranking than dissociation only for the small area above the thick red curve marked MC(1k). For MC, we used the test results from Fig. 5i; i.e. assuming $0.1 < \text{avg}[p_a] < 0.9$ for MC. Also recall that for large lineages, having an input probability with $\text{avg}[p_i] = 0.5$ will often lead to answer probabilities close to 1 for which ranking is not possible anymore (recall Fig. 5k). Thus, for large lineages, we need small input probabilities to have meaningful interpretations. And for small input probabilities, dissociation considerably outperforms any other method.

QUESTION 4. *How much would the ranking change according to exact probabilistic inference if we scale down all input tuples?*

Result 7. If the probabilities of all input tuples are already small, then scaling them further down does not affect the ranking much.

Here, we repeatedly evaluated the exact ranking for 7 different parameterized queries over randomly generated databases with one query plan that has $\text{avg}[d] \approx 3$, for two conditions: first on a probabilistic database with $\text{avg}[p_i]$ input probabilities (we defined the resulting ranking as GT); then again on a scaled version, where all input probabilities in the database are multiplied by the same scaling factor $f \in (0, 1)$. We then compared the new ranking against GT. Figure 5n shows that if all input probabilities are already small (and dissociation already works well), then scaling has little effect on the ranking. However, for $\text{avg}[p_i] = 0.5$ (and thus many tuples with p_i close to 1), we have a few tuples with p_i close to 1. These tuples are very influential for the final ranking, but their relative influence decreases if scaled down even slightly. Also note that even

for $\text{avg}[p_i] = 0.5$, scaling a database by a factor $f = 0.01$ instead of $f = 0.2$ does not make a big difference. However, the quality remains well above ranking by lineage size (!). This suggests that the difference between ranking by lineage size (MAP ≈ 0.529) and the ranking on a scaled database for $f \rightarrow 0$ (MAP ≈ 0.879) can be attributed to the relative weights of the input tuples (we thus refer to this as “ranking by relative input weights”). The remaining difference in quality then comes from the *actual probabilities* assigned to each tuple. Using MAP ≈ 0.220 as baseline for random ranking, 38% of the ranking quality can be found by the lineage size alone vs. 85% by the lineage size plus the relative weights of input tuples. The remaining 15% come from the actual probabilities (Fig. 5o).

QUESTION 5. *Does the expected ranking quality of dissociation decrease to random ranking for increasing fractions of dissociation (just like MC does for decreasing number of samples)?*

Result 8. The expected performance of dissociation for increasing $\text{avg}[d]$ for a particular query is lower bounded by the quality of ranking by relative input weights.

Here, we use a similar setup as before and now compare various rankings against each other: SampleSearch on the original database (“GT”); SampleSearch on the scaled database (“Scaled GT”); dissociation on the scaled database (“Scaled Diss”); and ranking by lineage size (which is unaffected by scaling). From Fig. 5p, we see that the quality of Scaled Diss w.r.t. Scaled GT $\rightarrow 1$ for $f \rightarrow 0$ since dissociation works increasingly well for small $\text{avg}[p_i]$ (recall Prop. 21). We also see that Scaled Diss w.r.t. GT decreases towards Scaled GT w.r.t. GT for $f \rightarrow 0$. Since dissociation can always reproduce the ranking quality of ranking by relative input weights by first downscaling the database (though losing information about the actual probabilities) the expected quality of dissociation for smaller scales does not decrease to random ranking, but rather to ranking by relative weights. Note this result only holds for the expected MAP; any particular ranking can still be very much off.

6. RELATED WORK

Probabilistic databases. Current approaches to query evaluation on probabilistic databases can be classified into three categories: (i) *incomplete approaches* identify tractable cases either at the query-level [7, 8, 17] or the data-level [38, 46, 50]; (ii) *exact approaches* [2, 30, 38, 39, 49] work well on queries with simple lineage expressions, but perform poorly on database instances with complex lineage expressions. (iii) *approximate approaches* either apply general purpose sampling methods [29, 32, 33, 44], or approximate the number of models of the Boolean lineage expression [16, 40, 45]. Our work can be seen as a generalization of several of these techniques: Our algorithm returns the exact score if the query is safe [7, 39] or data-safe [30].

Lifted and approximate inference. Lifted inference was introduced in the AI literature as an approach to probabilistic inference that uses the first-order formula to exploit symmetries at the grounded level [42]. This research evolved independently of that on probabilistic databases, and the two have many analogies: A formula is called *domain liftable* iff its data complexity is in polynomial time [28], which is the same as a *safe query* in probabilistic databases, and the FO-d-DNNF circuits described in [12] correspond to the safe plans discussed in this paper. See [11] for a recent discussion on the similarities and differences.

Representing Correlations. The most popular approach to represent correlations between tuples in a probabilistic database is

by a Markov Logic network (MLN) which is a set of *soft constraints* [13]. Quite remarkably, all complex correlations introduced by an MLN can be rewritten into a query over a tuple-independent probabilistic database [24, 26, 31]. In combination with such rewritings, our techniques can be also applied to MLNs if their rewritings results in conjunctive queries without self-joins.

Dissociation. Dissociation was first introduced in the workshop paper [20], presented as a way to generalize graph propagation algorithms to hypergraphs. Theoretical upper and lower bounds for dissociation of Boolean formulas, including Theorem 8, were proven in [22]. Dissociation is related to a technique called *relaxation* for probabilistic inference in graphical models [10].

7. CONCLUSIONS AND OUTLOOK

This paper proposes to approximate probabilistic query evaluation by evaluating a fixed number of query plans, each providing an upper bound on the true probability, then taking their minimum. We provide an algorithm that takes into account important schema information to enumerate only the minimal necessary plans among all possible plans, and prove it to be a strict generalization of all known results of PTIME self-join-free conjunctive queries. We describe relational query optimization techniques that allow us to evaluate all minimal queries in a single query and very fast: Our experiments show that these optimizations bring approximate probabilistic query evaluation close to standard query evaluation while providing high ranking quality. In future work, we plan to generalize this approach to full first-order queries. We will also make slides illustrating our algorithms available at <http://LaPushDB.com>.

Acknowledgements. This work was supported in part by NSF grants IIS-0513877, IIS-0713576, IIS-0915054, and IIS-1115188. We thank the reviewers for their careful reading of this manuscript and their detailed feedback. WG would also like to thank Manfred Hauswirth for a small comment in 2007 that was crucial for the development of the ideas in this paper.

8. REFERENCES

- [1] A. Amarilli, Y. Amsterdamer, and T. Milo. Uncertainty in crowd data sourcing under structural constraints. In *DASFAA Workshops*, pp. 351–359, 2014.
- [2] L. Antova, T. Jansen, C. Koch, and D. Olteanu. Fast and simple relational processing of uncertain data. In *ICDE*, pp. 983–992, 2008.
- [3] L. Antova, C. Koch, and D. Olteanu. Maybms: Managing incomplete information with probabilistic world-set decompositions. In *ICDE*, pp. 1479–1480, 2007.
- [4] P. Beame, J. Li, S. Roy, and D. Suciu. Model counting of query expressions: Limitations of propositional methods. In *ICDT*, pp. 177–188, 2014.
- [5] A. Carlson, J. Betteridge, B. Kisiel, B. Settles, E. R. H. Jr., and T. M. Mitchell. Toward an architecture for never-ending language learning. In *AAAI*, 2010.
- [6] Y. Chen and D. Z. Wang. Knowledge expansion over probabilistic knowledge bases. In *SIGMOD*, pp. 649–660, 2014.
- [7] N. N. Dalvi and D. Suciu. Efficient query evaluation on probabilistic databases. *VLDB J.*, 16(4):523–544, 2007.
- [8] N. N. Dalvi and D. Suciu. The dichotomy of probabilistic inference for unions of conjunctive queries. *J. ACM*, 59(6):30, 2012.
- [9] DeepDive: <http://deepdive.stanford.edu/>.
- [10] G. V. den Broeck, A. Choi, and A. Darwiche. Lifted relax, compensate and then recover: From approximate to exact lifted probabilistic inference. In *UAI*, pp. 131–141, 2012.
- [11] G. V. den Broeck and D. Suciu. Lifted probabilistic inference in relational models. In *UAI tutorials*, 2014.
- [12] G. V. den Broeck, N. Taghipour, W. Meert, J. Davis, and L. D. Raedt. Lifted probabilistic inference by first-order knowledge compilation. In *IJCAI*, pp. 2178–2185, 2011.
- [13] P. Domingos and D. Lowd. *Markov Logic: An Interface Layer for Artificial Intelligence*. Morgan & Claypool Publishers, 2009.
- [14] X. Dong, E. Gabrilovich, G. Heitz, W. Horn, N. Lao, K. Murphy, T. Strohmann, S. Sun, and W. Zhang. Knowledge vault: A web-scale approach to probabilistic knowledge fusion. In *KDD*, pp. 601–610, 2014.
- [15] B. Ermis and G. Bouchard. Scalable binary tensor factorization. In *UAI*, 2014.
- [16] R. Fink and D. Olteanu. On the optimal approximation of queries using tractable propositional languages. In *ICDT*, pp. 174–185, 2011.
- [17] R. Fink and D. Olteanu. A dichotomy for non-repeating queries with negation in probabilistic databases. In *PODS*, pp. 144–155, 2014.
- [18] 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.
- [19] W. Gatterbauer, S. Günnemann, D. Koutra, and C. Faloutsos. Linearized and single-pass belief propagation. *PVLDB*, 8(5), 2015. (CoRR abs/1406.7288).
- [20] W. Gatterbauer, A. K. Jha, and D. Suciu. Dissociation and propagation for efficient query evaluation over probabilistic databases. In *MUD*, pp. 83–97, 2010.
- [21] W. Gatterbauer and D. Suciu. Dissociation and propagation for efficient query evaluation over probabilistic databases. CoRR abs/1310.6257, 2013.
- [22] W. Gatterbauer and D. Suciu. Oblivious bounds on the probability of Boolean functions. *ACM Trans. Database Syst.*, 39(1):5, 2014. (CoRR abs/1409.6052).
- [23] V. Gogate and P. Domingos. Formula-based probabilistic inference. In *UAI*, pp. 210–219, 2010.
- [24] V. Gogate and P. Domingos. Probabilistic theorem proving. In *UAI*, pp. 256–265, 2011.
- [25] C. P. Gomes, A. Sabharwal, and B. Selman. Model counting. In *Handbook of Satisfiability*, pp. 633–654, 2009.
- [26] A. D. Guy Van den Broeck, Wannes Meert. Skolemization for weighted first-order model counting. In *KR*, 2014.
- [27] J. Hoffart, F. M. Suchanek, K. Berberich, and G. Weikum. Yago2: A spatially and temporally enhanced knowledge base from wikipedia. *Artif. Intell.*, 194:28–61, 2013.
- [28] M. Jaeger and G. V. den Broeck. Liftability of probabilistic inference: Upper and lower bounds. In *StarAI*, 2012.
- [29] R. Jampani, F. Xu, M. Wu, L. L. Perez, C. M. Jermaine, and P. J. Haas. MCDB: a Monte Carlo approach to managing uncertain data. In *SIGMOD*, pp. 687–700, 2008.
- [30] A. Jha, D. Olteanu, and D. Suciu. Bridging the gap between intensional and extensional query evaluation in probabilistic databases. In *EDBT*, pp. 323–334, 2010.
- [31] A. K. Jha and D. Suciu. Probabilistic databases with markovviews. *PVLDB*, 5(11):1160–1171, 2012.
- [32] S. Joshi and C. M. Jermaine. Sampling-based estimators for subset-based queries. *VLDB J.*, 18(1):181–202, 2009.
- [33] O. Kennedy and C. Koch. PIP: A database system for great and small expectations. In *ICDE*, pp. 157–168, 2010.
- [34] C. D. Manning, P. Raghavan, and H. Schütze. *Introduction to Information Retrieval*. Cambridge University Press, 2008.
- [35] F. McSherry and M. Najork. Computing information retrieval performance measures efficiently in the presence of tied scores. In *ECIR*, pp. 414–421, 2008.
- [36] F. Niu, C. Ré, A. Doan, and J. W. Shavlik. Tuffy: Scaling up statistical inference in markov logic networks using an RDBMS. *PVLDB*, 4(6):373–384, 2011.
- [37] OEIS: The on-line encyclopedia of integer sequences: <http://oeis.org/>.
- [38] D. Olteanu and J. Huang. Using OBDDs for efficient query evaluation on probabilistic databases. In *SUM*, pp. 326–340, 2008.
- [39] D. Olteanu, J. Huang, and C. Koch. Sprout: Lazy vs. eager query plans for tuple-independent probabilistic databases. In *ICDE*, pp. 640–651, 2009.
- [40] D. Olteanu, J. Huang, and C. Koch. Approximate confidence computation in probabilistic databases. In *ICDE*, pp. 145–156, 2010.
- [41] J. Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann Publishers, 1988.
- [42] D. Poole. First-order probabilistic inference. In *IJCAI*, pp. 985–991, 2003.
- [43] PostgreSQL 9.2: <http://www.postgresql.org/download/>.
- [44] C. Ré, N. N. Dalvi, and D. Suciu. Efficient top-k query evaluation on probabilistic data. In *ICDE*, pp. 886–895, 2007.
- [45] C. Ré and D. Suciu. Approximate lineage for probabilistic databases. *PVLDB*, 1(1):797–808, 2008.
- [46] S. Roy, V. Perduca, and V. Tannen. Faster query answering in probabilistic databases using read-once functions. In *ICDT*, pp. 232–243, 2011.
- [47] SampleSearch: <http://www.hlt.utdallas.edu/~vgogate/SampleSearch.html>.
- [48] P. G. Selinger, M. M. Astrahan, D. D. Chamberlin, R. A. Lorie, and T. G. Price. Access path selection in a relational database management system. In *SIGMOD*, pp. 23–34, 1979.
- [49] P. Sen and A. Deshpande. Representing and querying correlated tuples in probabilistic databases. In *ICDE*, pp. 596–605, 2007.
- [50] P. Sen, A. Deshpande, and L. Getoor. Read-once functions and query evaluation in probabilistic databases. *PVLDB*, 3(1):1068–1079, 2010.
- [51] A. P. Singh and G. J. Gordon. Relational learning via collective matrix factorization. In *KDD*, pp. 650–658, 2008.
- [52] J. Stoyanovich, S. B. Davidson, T. Milo, and V. Tannen. Deriving probabilistic databases with inference ensembles. In *ICDE*, pp. 303–314, 2011.
- [53] D. Suciu, D. Olteanu, C. Ré, and C. Koch. *Probabilistic Databases*. Morgan & Claypool Publishers, 2011.
- [54] TPC-H benchmark: <http://www.tpc.org/tpch/>.
- [55] K. Zeng, S. Gao, B. Mozafari, and C. Zaniolo. The analytical bootstrap: a new method for fast error estimation in approximate query processing. In *SIGMOD*, pp. 277–288, 2014.
- [56] C. Zhang and C. Ré. Towards high-throughput Gibbs sampling at scale: a study across storage managers. In *SIGMOD*, pp. 397–408, 2013.