# Read-Once Functions and Query Evaluation in Probabilistic Databases 

Prithviraj Sen<br>Yahoo! Labs<br>Bangalore, India<br>sen@yahoo-inc.com

Amol Deshpande Lise Getoor<br>Computer Science Department<br>University of Maryland, College Park, USA<br>amol@cs.umd.edu, getoor@cs.umd.edu


#### Abstract

Probabilistic databases hold promise of being a viable means for large-scale uncertainty management, increasingly needed in a number of real world applications domains. However, query evaluation in probabilistic databases remains a computational challenge. Prior work on efficient exact query evaluation in probabilistic databases has largely concentrated on query-centric formulations (e.g., safe plans, hierarchical queries), in that, they only consider characteristics of the query and not the data in the database. It is easy to construct examples where a supposedly hard query run on an appropriate database gives rise to a tractable query evaluation problem. In this paper, we develop efficient query evaluation techniques that leverage characteristics of both the query and the data in the database. We focus on tuple-independent databases where the query evaluation problem is equivalent to computing marginal probabilities of Boolean formulas associated with the result tuples. This latter task is easy if the Boolean formulas can be factorized into a form that has every variable appearing at most once (called read-once). However, a naive approach that directly uses previously developed Boolean formula factorization algorithms is inefficient, because those algorithms require the input formulas to be in the disjunctive normal form (DNF). We instead develop novel, more efficient factorization algorithms that directly construct the read-once expression for a result tuple Boolean formula (if one exists), for a large subclass of queries (specifically, conjunctive queries without self-joins). We empirically demonstrate that (1) our proposed techniques are orders of magnitude faster than generic inference algorithms for queries where the result Boolean formulas can be factorized into read-once expressions, and (2) for the special case of hierarchical queries, they rival the efficiency of prior techniques specifically designed to handle such queries.


## 1. INTRODUCTION

The rise in uncertain data in a variety of applications has led to much research in the area of probabilistic databases in recent years [10, 1]. Since query evaluation is \#P-complete in probabilistic databases that use the possible world semantics (the dominant semantics on which most systems are based), two broad approaches

[^0]have been suggested for tractable query execution. Either the system opts to compute approximate query results [21, 18, 25], or the query language is restricted to allow for efficient exact query evaluation. The research on hierarchical queries [11, 25] has identified many such classes of queries (reviewed in the next section) that can be evaluated in PTIME on tuple-independent databases (mutual exclusivity correlations may be permitted in some cases). Unfortunately several negative dichotomy and trichotomy results [11, 21, 26] developed in recent years suggest that the class of hierarchical queries is likely to be too restrictive. More importantly, a querycentric approach is pessimistic by definition; it may be the case that a non-hierarchical query can be tractably evaluated on most probabilistic databases encountered in practice.

In this paper, we develop novel techniques for query evaluation in tuple-independent probabilistic databases by drawing connections to the literature on read-once functions. It is well known that in probabilistic databases with independent tuples, every result tuple is associated with a Boolean formula [3,13] (often called lineage), and the query evaluation problem reduces to computing the marginal probabilities for the result tuple Boolean formulas holding true. Further, if a result tuple's Boolean formula can be factorized into a form where every Boolean variable appears at most once, also known as read-once functions [15, 17], then the marginal probability can be computed easily in linear time. Hierarchical queries always admit a query evaluation plan such that the result tuple formulas are generated in read-once form, thus providing a connection to efficient query evaluation in probabilistic databases [22].

Here we further explore and exploit the connection by first showing how to incorporate previously developed Boolean formula factorization algorithms into a probabilistic database query engine. In this naive approach, we run the user-submitted query using any query plan to generate all the result tuples along with their lineages, and then we factorize them into read-once form (if possible) to compute marginal probabilities. This approach by itself allows us to tackle a superset of hierarchical queries more efficiently than has been previously possible. However, it is quite inefficient because prior factorization algorithms require the disjunctive normal form (DNF) of the input Boolean formula; DNF expansion of a formula may result in an exponential blowup (in the size of the query).

We instead develop novel algorithms that allow us to construct the read-once expression for an output tuple (if one exists) on-thefly during query evaluation, for the class of conjunctive queries without self-joins. Our proposed technique subdivides the result tuple lineage into smaller formulas for which we build read-once expressions separately. We then merge these read-once expressions to produce a read-once expression of the result tuple lineage. Both these steps are challenging. In the first step, we have to carefully strategize how to sub-divide the result tuple lineage into smaller

| L: | X | J : | X | Y | R : | Y |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $x_{1}$ | $x_{1}$ | $z_{1}$ | $x_{1}$ | $y_{1}$ | $y_{1}$ | $y_{1}$ |
| $x_{2}$ | $x_{2}$ | $z_{2}$ | $x_{1}$ | $y_{2}$ | $y_{2}$ | $y_{2}$ |
| $x_{3}$ | $x_{3}$ | $z_{3}$ | $x_{2}$ $x_{3}$ | $y_{3}$ | $y_{3}$ | $y_{3}$ |

$$
\begin{aligned}
q() & : L(\mathbf{X}), J(\mathbf{X}, \mathbf{Y}), R(\mathbf{Y}) \\
r & =x_{1} z_{1} y_{1}+x_{1} z_{2} y_{2}+x_{2} z_{3} y_{3}+x_{3} z_{4} y_{3}
\end{aligned}
$$

$\operatorname{Pr}(e)=\left\{\begin{array}{l}\prod_{c \in \operatorname{ch(e)}} \operatorname{Pr}(c) \\ 1-\prod_{c \in c h(e)}(1-\operatorname{Pr}(c)) \\ p_{t}\end{array}\right.$
(b)

(c)

(d)
if $v$ is (1) if $v$ is (1) if $v=x_{t}$

(e)

$y_{1}$
(f)

Figure 1: (a) A query $q$ and its singleton result $r$ along with its Boolean formula. (b) Computing the marginal probability from a co-tree where $c h(e)$ denotes children of $v$. (c) $a b+b c+c d$, the $P_{4}$ structure. (d) $c(a b+d)$, co-occurrence graph is $P_{4}$-free. (e) The co-occurrence graph of $r$ from (a), and (f) its co-tree.
components. The divisions should be such that if a sub-division does not admit a read-once rewriting, then the result tuple does not have a read-once expression either. In the main body of the paper, we show that if we are using deep query evaluation plans (where at least one input to every join operator is a base relation), it is possible to divide the result tuple's lineage into such sub-divisions. In the appendix, we show that the same logic can extended to include bushy query plans. The second step is where we combine the readonce expressions of the various sub-divisions together. Read-once expressions are usually represented using tree-like data structures known as co-trees [6]. Thus, this step is an instance of the tree alignment problem [20] which in general is MAX SNP-hard; we propose a novel polynomial-time algorithm for solving this problem in our setup. We evaluate our approach by experimenting with several datasets, including the TPC-H benchmark. We show that not only are our proposed techniques orders of magnitude faster than generic inference algorithms when evaluating non-hierarchical queries, they rival the performance of techniques specifically developed to handle hierarchical queries.

Two of the state-of-the-art probabilistic database systems that solve hierarchical queries efficiently are heavily dependant on query compilation techniques. MystiQ [4] looks at the query and constructs a plan that allows efficient evaluation. SPROUT [24] takes the query's hierarchical representation and constructs a signature which is used in subsequent operations. These works also propose the use of functional dependencies described in the schema to simplify the query. It is possible that a non-hierachical query results in a hierarchical query when the functional dependencies are taken into account. Our premise, in this paper, is very simple. If the query results in read-once result tuples (either because of the structure in the query or because the data satisfies functional dependencies or for any other reason) then we should be able to compute marginal probabilities easily. In that sense, the techniques we propose in this paper are data-centric and largely disjoint from the techniques employed in SPROUT and MystiQ which are query-centric. This obviates the need for a fair amount of compile time analysis involving the query and functional dependencies. We evaluate against SPROUT and MystiQ extensively in Section 4.

Recently, Olteanu et al. [25] also proposed using factorization algorithms within a probabilistic query engine, to compile lineages into decision diagrams called d-trees. D-trees allow Shannon expansion and hence are a stronger formalism than read-once functions. However, that work assumes that the DNF of the result tuple is provided as input which, as we mentioned earlier, can be expensive to construct. An interesting open question here is whether it is possible to directly compute small d-trees during query execution similar to what we do with read-once expressions in this work. Another issue with the current work on hierarchical queries is that, with some exceptions (e.g., Olteanu et al. [23, 22]), most of those
deal almost exclusively with equality join predicates (e.g., Dalvi et al. [11]). Extending the approach to other operators requires effort, and dealing with queries composed of different operators is even more cumbersome. On the contrary, viewing result tuples as Boolean formulas allows us to restrict our attention to only two operators, viz. $\wedge($ and $)$ and $\vee$ (or). We do not care what kind of join predicate (equality or inequality or anything else) gave rise to the Boolean formula associated with the result tuple. Thus, our techniques are likely to be more widely applicable than prior work.
Outline: In the next section, we discuss preliminaries and show how to incorporate prior factorization algorithms into the query engine. In Section 3, we devise novel and efficient algorithms to construct read-once functions for result tuples produced by conjunctive queries without self-joins. In Section 4, we experimentally demonstrate the efficacy of our proposed techniques. In Section 5, we discuss related work before concluding with Section 6. Full proofs appear in the accompanying appendix due to space constraints.

## 2. PRELIMINARIES

Let $R$ denote a relation defined over a set of attributes $\operatorname{Attr}(R)$. Each tuple $t \in R$ is a mapping from $\operatorname{Attr}(R)$ to values. We also associate a unique (Boolean-valued) random variable with $t$ denoted by $x_{t}$ and a probability of existence $p_{t}$. Often, when it is clear from the context, we will abuse notation and refer to the tuple's random variable by the tuple itself. A (probabilistic) database $D=\left\{R_{1}, \ldots R_{m}\right\}$ is a set of relations and represents a distribution over many possible worlds each obtained by choosing a (sub)set of tuples in $R_{i}$ to be present. If $t$ is present we say $x_{t}$ is assigned true or t and false or f , otherwise. Each possible world $w$ is associated with a probability:

$$
\operatorname{Pr}(w)=\prod_{i=1}^{m} \prod_{t \in R_{i}, x_{t}=\mathrm{t}} p_{t} \prod_{t \in R_{i}, x_{t}=\mathrm{f}}\left(1-p_{t}\right)
$$

Given a query $q$ to be evaluated against database $D$, the result of the query is defined to be the union of results returned by each possible world along with the marginal probabilities of each result tuple [9]. More precisely, the marginal probability of result tuple $t$ is obtained by adding the probabilities of all possible worlds that return $t$ as a result: $\mu(t)=\sum_{t \in q(w)} \operatorname{Pr}(w)$, where $q(w)$ denotes the result of the query on that possible world.

One way to compute the marginal probability of result tuple $t$ produced by (relational algebra) query $q$ is to extend each (relational algebra) operator in $q$ so that it builds a Boolean formula for each (intermediate) tuple generated during query evaluation. We refer to the Boolean formula for $t$ by $\phi_{t}$. Below we show these extended definitions for operators $\sigma, \times$ and $\Pi$.

$$
\begin{array}{ll}
\phi_{t}=x_{t} \forall t \in R, & \phi_{\sigma_{c}(t)}=\text { if } c(t) \text { then } \phi_{t} \text { else } \mathrm{f} \\
\phi_{t \times t^{\prime}}=\phi_{t} \wedge \phi_{t^{\prime}}, & \phi_{\Pi\left(t_{1}, \ldots t_{k}\right)}=\bigvee_{i=1}^{k} \phi_{t_{i}}
\end{array}
$$

The marginal probability of the result tuple can then be obtained by computing the probability of the corresponding Boolean formula holding true. We refer the interested reader to prior work that illustrates the equivalence between query evaluation under possible world semantics and via Boolean formulas [3, 13, 27]. Figure 1(a) shows a three-relation join query which produces a singleton result and the corresponding result tuple's Boolean formula.

### 2.1 Read-Once Functions; Definitions

Although marginal probability computation is \#P-Complete in general, efficient computation is possible in many cases.

Definition 1 (Read-Once Function [17]). A Boolean formula $\phi$ is read-once if there exists a factorization such that each variable appears not more than once.

The read-once factorized form of the Boolean formula is known as its read-once expression. $r$ in Figure 1(a) is a read-once function with the read-once expression $x_{1}\left(z_{1} y_{1}+z_{2} y_{2}\right)+y_{3}\left(x_{2} z_{3}+x_{3} z_{4}\right)$.

THEOREM 1 ([15]). A Boolean formula is read-once iff it is unate, $P_{4}$-free and normal.

A Boolean formula $\phi$ is unate [15] if every variable only appears in either its positive or negated form. E.g., $a b$ and $\bar{a} b+\bar{a} c$ are unate but $\bar{a} b+a c$ is not. For any Boolean formula $\phi$, the co-occurrence graph $G_{\phi}$ is formed by drawing a vertex for every variable and drawing an edge between two variables if they appear in the same clause in $\phi$ 's disjunctive normal form (DNF). Let $X$ denote a subset of vertices, then the subgraph of $G$ induced by $X$ is the subgraph formed by restricting edges of $G$ to edges with end points in $X$.

The special graph $P_{4}$ denotes a chordless path with 4 vertices and 3 edges (Figure 1(c)). $\phi$ is $P_{4}$-free if no induced subgraph of $G_{\phi}$ forms a $P_{4}$. Figure 1(c) and (d) show two formulas and the former is not read-once because it forms a $P_{4}$. Figure 1(d) is $P_{4}$ free because of the presence of the edge $a-c$. Figure 1(e) shows the co-occurrence graph for the result tuple $r$ from Figure 1(a). A formula $\phi$ is said to be normal if every clique in its co-occurrence graph is contained in some clause in its DNF [15]. For instance, even though both $\phi_{1}=a b c$ and $\phi_{2}=a b+b c+c a$ have the same co-occurrence graph (the triangle), $\phi_{1}$ is normal while $\phi_{2}$ is not.

Traditionally, co-trees [6] have been used to represent read-once expressions. Co-trees are trees where leaves correspond to Boolean variables while internal node (1) represents $\wedge$ and (0) represents $\vee$. A given read-once expression can be represented by many co-trees but there exists a canonical co-tree, where (1) and (0) alternate on every path. See Figure 1(f) for an example. Given the co-tree for a read-once result tuple, the marginal probability can be computed using the simple, bottom-up procedure shown in Figure 1(b).

Our goal in this paper is essentially to construct the co-trees corresponding to all the result tuple lineages efficiently, if they exist.

### 2.2 Hierarchical Queries

Earlier work on query evaluation in probabilistic databases has identified tractable queries, known as hierarchical queries [11], for which probability computation is efficient. For the ensuing discussion, we will assume that all queries are projected onto the empty set (in other words, our goal is simply to compute the probability that there exists at least one result tuple). Queries projected onto a non-empty set of attributes can be handled by replacing these attributes with constants. Let $q$ denote a query in Datalog notation. Let $a$ denote an attribute and $\operatorname{sg}(a)$ denote the set of relations in $q$ that refer to $a$. In Figure $1(\mathrm{a}), \operatorname{sg}(\mathbf{X})=\{L, J\}, s g(\mathbf{Y})=\{J, R\}$.

DEFINITION 2 (Hierarchical Query [11]). A (conjunctive) query $q$ is hierarchical if, for any two attributes $a$ and $b$, either $\operatorname{sg}(a) \subseteq$ $s g(b), s g(b) \subseteq s g(a)$ or $s g(a) \cap s g(b)=\emptyset$.

For instance, the query $q$ in Figure 1(a) is not hierarchical $(\operatorname{sg}(\mathbf{X}) \cap$ $s g(\mathbf{Y})=\{J\} \neq \emptyset, s g(\mathbf{X}) \nsubseteq s g(\mathbf{Y}), s g(\mathbf{Y}) \nsubseteq s g(\mathbf{X}))$ but $q^{\prime}():-S(\mathbf{X}, \mathbf{Y}), T(\mathbf{Y})$ is (because $s g(\mathbf{Y})=\{S, T\} \supset\{S\}=$ $s g(\mathbf{X})$ ). Dalvi et al. [11] (also, Olteanu et al. [22]) illustrated the connection between hierarchical queries and read-once functions:

PROPOSITION 1 ([11, 22]). Hierarchical queries produce result tuples with read-once expressions.

The converse is clearly not true. In Figure 1(a), we show a query that is not hierarchical, but the result tuple it produces has a readonce expression (Figure 1(f)). In this paper, we would like to develop techniques that help us identify such cases efficiently.

### 2.3 A Naive Approach

One viable approach to evaluating queries is to generate Boolean formulas for result tuples (using the extended operators shown earlier) and then determine whether it is a read-once function. If the result tuple lineage is read-once then we compute its probability from its read-once expression's co-tree, else we resort to a generalpurpose inference engine or approximations.

Checking for read-once result tuples is possible in polynomial time. In what follows, let $\phi$ denote a result tuple's Boolean formula in DNF, $|\phi|$ its length (with different occurrences of the same variable counted multiple times) and $\operatorname{Vars}(\phi)$ the set of distinct variables in it. To check for unateness, a linear scan of the formula is sufficient which requires $O(|\phi|)$ time. To check for $P_{4}$ 's in $\phi$ 's co-occurrence graph $G_{\phi}$, a handful of algorithms are available [7, $16,5]$. The common aspect of all of these algorithms is that they all require $G_{\phi}$ to be provided as input and they run in time linear in size of $G_{\phi}(O(|\operatorname{Vars}(\phi)|+|E|)$, where $|E|$ is the number of edges in $G_{\phi}$ ). $G_{\phi}$ can be obtained easily from the formula's DNF. A nice property of these algorithms is that not only do they check for the absence of $P_{4}$ 's, but they also return the co-tree representing the read-once expression which can subsequently be used for probability computation. Checking for normality is also possible in polynomial time. [15] describes a way to do this in $O(|\operatorname{Vars}(\phi) \| \phi|)$ time using the co-tree obtained in the previous step.

### 2.4 Limitations

Even though the above approach to query evaluation is both sound and complete, it can be expensive for large probabilistic databases. Two steps in our query evaluation procedure are of particular concern: (1) checking for normality because of its quadratic time complexity, and (2) checking for $P_{4}$ 's. All the algorithms that check for $P_{4}$ 's that we are aware of expect the co-occurrence graph $\left(G_{\phi}\right)$ as input for which we will likely require the formula to be presented in DNF. Our extended relational algebra operators do not guarantee that a result tuple's formula will be returned in its DNF form (instead they return a more compact representation that we call lineage-trees; see next section). Moreover, converting a formula into its DNF form can result in an exponential blowup. More specifically, if a result tuple produced by query $q$ involves $n$ tuples from each relation and there are $k$ relations being joined in $q$, then the corresponding formula's DNF form could require $O\left(n^{k}\right)$ units of space and time to compute.

## 3. READ-ONCE EXPRESSIONS FOR CONJUNCTIVE QUERIES

In this section, we present our key technical contribution: an efficient algorithm for directly constructing co-trees corresponding to result tuple lineages (if they exist) for conjunctive queries without self-joins, executed using a deep query evaluation plan. Conjunctive queries form a fairly large subset of database queries, and a
number of prior works in probabilistic databases have focused their attention to (subclasses of) conjunctive queries [9, 22].

We begin with formally defining conjunctive queries, and observing a key property of the result tuple lineages generated by a conjunctive query. This property enables us to eliminate the normality-checking step, resulting in significant efficiency gains (Section 3.1). We then show how to sub-divide the result tuple lineage so that co-trees can be built for resulting sub-formulas independently of each other; if any of the resulting sub-formulas is not $P_{4}$-free, then the result tuple lineage is guaranteed to contain a $P_{4}$ (Section 3.2). We then develop an algorithm for merging the co-trees corresponding to the sub-formulas (Section 3.3). We end the section with a discussion of further optimizations (Section 3.4).

### 3.1 Properties of Conjunctive Queries

Let $A$ denote an attribute. An atomic predicate is of the form $A$ op $B$ where $B$ is either an attribute or a constant conforming to the type of $A$ and $o p$ is any binary operator conforming to the same type such as $=,>,<, \neq$ etc. A conjunctive query $q$ involving the set of relations $\operatorname{Rels}(q)=\left\{R_{1}, \ldots R_{k}\right\}$ is a relational algebra query that involves the three operators $\sigma, \bowtie$ and $\Pi$ such that all selection and join predicates are either atomic predicates or conjunctions of atomic predicates. If $\operatorname{Rels}(q)$ does not contain any repeated relations then $q$ is a conjunctive query without self-joins.

Since conjunctive queries do not allow negations, the result tuples' Boolean formulas will be unate (variables appear only in their positive form). We next show that for result tuples we will be dealing with, the normality check is also not required. In the ensuing discussion, let $\phi_{D N F}=C_{1}+C_{2}+\cdots$ denote the DNF of Boolean formula $\phi$ where each $C_{i}$ denotes a clause or a conjunction of Boolean variables $x_{1} x_{2} \ldots$. Often we will treat clauses as sets, and use $x \in C$ to denote a variable $x$ present in the clause $C$.

For result tuple $\phi$ produced by conjunctive query $q$ without selfjoins, two properties easily follow. First, every clause $C$ in $\phi_{D N F}$ consists of exactly one tuple existence random variable from each relation in $\operatorname{Rels}(q)$. This is often referred to as $\phi_{D N F}$ being a kmonotone DNF. Second, given a collection $T$ of $k(=|\operatorname{Rels}(q)|)$ tuples, one each from a different relation, if $T$ satisfies all the selection and join predicates in $q$ and agrees with $\phi$ on the final set of projected attributes, then $T$ must form a clause in $\phi_{D N F}$. The second property can be restated in terms of $\phi$ 's co-occurrence graph $G_{\phi}$. If $G_{\phi}$ contains a $k$-sized clique then it must form a clause in $\phi_{D N F}$. We now state the main result of this subsection (proof can be found in the appendix):

Theorem 2. Let $\phi$ denote a result tuple produced by a conjunctive query without self-joins. If $\phi$ is $P_{4}$-free, then it is normal.

Theorem 2, along with the guarantee of unateness, means that whenever we are working with result tuples produced by conjunctive queries without self-joins, we only need to check if they are $P_{4}$-free. As long as they are $P_{4}$-free we can rest assured that a read-once expression exists. Recall that, the normality check was the most expensive step while checking for read-once functions. So this represents a significant gain. We next proceed towards constructing an efficient $P_{4}$-checking algorithm.

### 3.2 Piecewise Building of Co-Trees

Like prior $P_{4}$-checking algorithms, we would prefer an algorithm that not only checks for the absence of a $P_{4}$, but also returns the co-tree representation of the read-once expression (if no $P_{4}$ exists). Recall that, unlike other algorithms, we don't want use the co-occurrence graph or the DNF expression as input to our $P_{4}$ checking algorithm. Instead we use the representation of the result tuple Boolean formula as constructed by the query engine dur-

(a)

Figure 2: Lineage-trees of $r$ from Figure 1(a) using plans: (a) $\prod_{\emptyset}\left(R \bowtie_{\mathbf{Y}} \prod_{\mathbf{Y}}(L \bowtie \mathbf{x} J)\right)$, (b) $\prod_{\emptyset}\left(R \bowtie_{\mathbf{Y}}(L \bowtie \mathbf{x} J)\right)$.
ing evaluation. That way we do not have to make any significant changes to the query engine, and further the query optimizer is free to choose the best query evaluation plan (subject to the constraint that the query plan be deep).

We call such a representation a lineage-tree. A lineage-tree depicts how tuples join and project among one another to produce the result tuple. Going back to the example in Figure 1 (a), Figure 2 (a) shows the lineage-tree of $r$ obtained using the query plan $\prod_{\emptyset}\left(R \bowtie_{\mathbf{Y}} \prod_{\mathbf{Y}}\left(L \bowtie_{\mathbf{X}} J\right)\right)$ for $q$. Note that, lineage-tree is not the same as the co-tree ( $x_{1}$ appears twice in Figure 2 (a)). Further note that, the lineage-tree of the result tuple depends on the query plan used. Figure 2 (b) shows the lineage-tree obtained using the plan $\prod_{\emptyset}\left(R \bowtie_{\mathbf{Y}}(L \bowtie \mathbf{X} J)\right)$ instead. Part of our challenge now is to design an algorithm that constructs the correct read-once expression given any lineage-tree of the result tuple.

In what follows, we need to refer to nodes in the lineage-tree and for this, we need some extra notation. Let $\partial$ denote a left/right deep plan of query $q$. More specifically, let $\partial$ denote a plan such that every join in the plan involves at least one base relation. Without loss of generality, we will assume $\partial$ is of the form $o p_{1}\left(R_{1} \bowtie \bowtie_{\theta_{1}}\right.$ $\left.o p_{2}\left(R_{2} \bowtie_{\theta_{2}} \ldots\right)\right)$ where $\theta_{i}$ denotes a join predicate and $o p_{i}$ denotes a sequence of projections and selections ( $o p_{i}$ could also be identity). Lineage-trees produced by left/right deep plans are such that every join node $(\bowtie)$ always has at least one base tuple's existence variable as a child. In fact, we can use this as a way of identifying nodes in the lineage-tree.

Definition 3 (Join Path, Cofactor). Let $\bowtie^{k}$ denote a join node in lineage-tree $\mathcal{L}$. Let $\bowtie^{k-1}, \ldots \bowtie^{1}$ denote all join nodes along the path from $\bowtie^{k}$ to the root of $\mathcal{L}$ such that $\bowtie^{i-1}$ is $\bowtie^{i}$ 's immediate ancestor. Further, let $b_{i}$ denote (one of) the base tuple child of $\bowtie^{i}$. Then we refer to the sequence $\left[b_{1}, b_{2} \ldots b_{k}\right]$ as $a$ join path and $\bowtie^{k}$ as a join node identified by it. Further, we denote the remaining child of $\bowtie^{k}$ as cofactor $\mathcal{L}\left(\left[b_{1}, b_{2} \ldots b_{k}\right]\right)$ or $\operatorname{cof}_{\mathcal{L}}\left(\left[b_{1}, \ldots b_{k}\right]\right)$, in short.

For instance, in $\mathcal{L}$ shown in Figure 2 (a), join path $\left[y_{3}\right]$ identifies the node enclosed in the rectangle, $\left[y_{3}, x_{2}\right]$ identifies the node enclosed in the ellipse and $\operatorname{cof}_{\mathcal{L}}\left(\left[y_{3}, x_{2}\right]\right)$ is $z_{3}$.

In general, a join path identifies a set of nodes in the lineage-tree. For instance, in lineage-tree $\mathcal{L}$ in Figure 2 (b), join path $y_{3}$ identifies the two nodes enclosed in rectangles. In such cases, we define the cofactor to be the set of non-base tuple children of the identified join nodes. Alternatively, we can also represent the cofactor by its Boolean formula, in other words the disjunction of the formulas of the (intermediate) tuples representing the nodes the cofactor contains. For instance, in Figure $2(\mathrm{~b}), \operatorname{cof} f_{\mathcal{L}}\left(\left[y_{3}\right]\right)=x_{3} z_{4}+x_{2} z_{3}$. Another interesting aspect of cofactors is that they can be expressed in a recursive manner. $\operatorname{cof_{\mathcal {L}}([\Gamma ])\text {canbeexpressedintermsof}}$ $\operatorname{cof}_{\mathcal{L}}([\Gamma, b])$, where $b$ denotes a base tuple. For instance, in $\mathcal{L}$ in Figure $2(\mathrm{~b}), \operatorname{cof}_{\mathcal{L}}\left(\left[y_{3}\right]\right)=x_{3} \operatorname{cof}_{\mathcal{L}}\left(\left[y_{3}, x_{3}\right]\right)+x_{2} \operatorname{cof}_{\mathcal{L}}\left(\left[y_{3}, x_{2}\right]\right)$. This also illustrates how $\operatorname{cof}_{\mathcal{L}}([\Gamma, b])$ is a sub-formula of $\operatorname{cof} \mathcal{L}_{\mathcal{L}}([\Gamma])$

$$
\begin{aligned}
& \text { 1. } \mathcal{T}\left(b_{1}, \ldots b_{n}\right)=(0)\left(b_{1}, \ldots b_{n}\right) \text {, s.t. } b_{i} \in R \\
& \text { 2. } \cup\left(o p\left(t_{1}^{1}, \ldots t_{n}^{1}\right), t_{2} \ldots t_{m}\right)=\left\{t_{1}^{1}, \ldots t_{n}^{1}, \cup\left(t_{2}, \ldots t_{m}\right)\right\} \\
& \text { 3. } \mathcal{T}\left(o p\left(t_{1}^{1}, \ldots t_{n}^{1}\right), t_{2}, \ldots t_{m}\right)=\mathcal{T}\left(\cup\left(o p\left(t_{1}^{1}, \ldots t_{n}^{1}\right), t_{2}, \ldots t_{m}\right)\right) \\
& \text { 4. } \mathcal{T}\left(b_{1} \bowtie t_{1}, \ldots b_{n} \bowtie t_{m}\right)=\left(1\left(b_{1}, \mathcal{T}\left(t_{1}^{b_{1}}, \ldots t_{m_{1}}^{b_{1}}\right)\right)\right. \\
& \quad \oplus\left(1 ( b _ { 2 } , \mathcal { T } ( t _ { 1 } ^ { b _ { 2 } } , \ldots t _ { m _ { 2 } } ^ { b _ { 2 } } ) ) \ldots \oplus \left(1\left(b_{n}, \mathcal{T}\left(t_{1}^{b_{n}}, \ldots t_{m_{n}}^{b_{n}}\right)\right)\right.\right.
\end{aligned}
$$

Figure 3: Co-tree building procedure. $o p$ is $\Pi$ or $\sigma$.
and is contained within it.
Just like many other tree-building algorithms, our co-tree building algorithm is a recursive algorithm that builds co-trees for subformulas of the result tuple and combines them to construct the complete read-once expression. However, choosing sub-formulas is tricky - this is because having a $P_{4}$ is not a monotonic property. Let $\phi=\phi_{1}+\phi_{2}$. Then $\phi_{1}$ (or $\phi_{2}$ ) containing a $P_{4}$ does not imply that $\phi$ contains a $P_{4}$. This is illustrated by the formula $a_{1} b_{1} x_{1} y_{1}+a_{1} b_{2} x_{1} y_{2}+a_{2} b_{1} x_{2} y_{1}+a_{2} b_{2} x_{2} y_{2}$ where the sum of any three terms gives us a $P_{4}$ even though the complete formula is read-once $\left(\left(a_{1} x_{1}+a_{2} x_{2}\right)\left(b_{1} y_{1}+b_{2} y_{2}\right)\right)$. Figure 10 in the appendix shows the query and database that produce this result tuple.

In the previous example, essentially, picking any three summands gives us a $P_{4}$ for which the fourth summand provides a chord. For instance, if we set $\phi_{1}=a_{1} b_{1} x_{1} y_{1}+a_{1} b_{2} x_{1} y_{2}+a_{2} b_{1} x_{2} y_{1}$, which contains the $P_{4}: b_{2}-a_{1}-b_{1}-a_{2}$, then the fourth summand $\phi_{2}=a_{2} b_{2} x_{2} y_{2}$ contains the edge $a_{2}-b_{2}$. Recall that, a $P_{4}$ is a chordless path of length 3 edges. Thus, when we consider $\phi_{1}+\phi_{2}$ we no longer have the $P_{4}{ }^{*}$. To express this more concretely we define the notion of interference:

Definition 4 (Interference, Type 1). Two formulas $\phi_{1}$ and $\phi_{2}$ interfere if $\exists a, b \in \operatorname{Vars}\left(\phi_{1}\right), \operatorname{Vars}\left(\phi_{2}\right)$ such that $a-b \notin$ $G_{\phi_{1}}$ and $a-b \in G_{\phi_{2}}$.
where $a-b \notin G$ denotes that $a$ and $b$ do not form an edge in $G$. Interference can also arise due to a triplet of formulas:

Definition 5 (Interference Type 2). Three formulas $\phi_{1}, \phi_{2}$ and $\phi_{3}$ interfere if $\exists a, b$ such that $a \in \operatorname{Vars}\left(\phi_{1}\right), a \notin \operatorname{Vars}\left(\phi_{2}\right), b \notin$ $\operatorname{Vars}\left(\phi_{1}\right), b \in \operatorname{Vars}\left(\phi_{2}\right)$ and $a-b \in G_{\phi_{3}}$.

Essentially, if we consider $\phi_{1}+\phi_{2}$ as one formula then type 2 interference reduces to type 1 interference.

For our task of building read-once expressions, we would like to divide the result tuple into non-interfering sub-formulas; and this is where the lineage-tree, with its cofactors, comes to our rescue:

Lemma 1. Let $\phi$ denote a result tuple, $\mathcal{L}$ its lineage and $\Gamma a$ join path. If $\phi$ is $P_{4}$-free then the set of cofactors $\left\{\operatorname{cof} \mathcal{L}_{\mathcal{L}}([\Gamma, b]) \mid b\right.$ is a base tuple and a valid extension of $\Gamma\}$ is non-interfering.

For instance, in Figure 2(a), the base tuples $y_{1}, y_{2}$ and $y_{3}$ form valid extensions of the empty join path. The lemma says that $\operatorname{cof} \mathcal{L}_{\mathcal{L}}\left(\left[y_{1}\right]\right)=$ $x_{1} z_{1}, \operatorname{cof}_{\mathcal{L}}\left(\left[y_{2}\right]\right)=x_{1} z_{2}$ and $\operatorname{cof} \mathcal{L}_{\mathcal{L}}\left(\left[y_{3}\right]\right)=x_{2} z_{3}+x_{3} z_{4}$ do not interfere with each other if the lineage of $r$ is read-once.

The lemma immediately gives us a way to sub-divide the result tuple lineage properly, to build its co-tree. Figure 3 outlines a recursive procedure $\mathcal{T}$ that takes a number of (base or intermediate) tuples $t_{1}, t_{2}, \cdots, t_{n}$ as arguments and returns the co-tree representing $\phi_{t_{1}} \vee \phi_{t_{2}} \vee \cdots \vee \phi_{t_{n}}$, if possible (recall that $\phi_{t}$ denote the lineage of tuple $t$ ). $\mathcal{T}$ begins at the root of the input lineage-tree and proceeds downwards.

[^1]\[

$$
\begin{aligned}
& \text { 1. } \mathcal{S}(x, y)= \begin{cases}1 & \text { if } x \equiv y \\
0 & \text { otherwise }\end{cases} \\
& \text { 2. } \mathcal{S}\left(T_{1}, T_{2}\right)=0 \text {, if } \neg \mathcal{C}\left(T_{1}\right) \wedge \neg \mathcal{C}\left(T_{2}\right) \\
& \text { 3. } \mathcal{S}\left(\text { (0) }\left(A_{1}, A_{2}, \ldots A_{n}\right), \text { (1) }(\ldots)\right)=\max _{i=1}^{n}\left(\mathcal{S}\left(A_{i}, \text { (1) }(\ldots)\right)\right) \\
& \text { 4. } \mathcal{S}\left(\text { (0) }\left(A_{1}, A_{2}, \ldots A_{n}\right),(0)\left(B_{1}, \ldots B_{m}\right)\right)= \\
& \left\{\begin{array}{l}
\mathcal{S}\left(0\left(A_{2}, \ldots A_{n}\right),(0)\left(B_{1}, B_{2}, \ldots B_{m}\right)\right), \text { if } \neg \mathcal{C}\left(A_{1}\right) \\
\max _{i=1}^{m}\left(\mathcal{S}\left(A_{1}, B_{i}\right)\right) \\
+\mathcal{S}\left(0\left(A_{2}, \ldots A_{n}\right),(0)\left(B_{1}, \ldots B_{m} \backslash B_{i}\right)\right)
\end{array}\right. \\
& \text { 5. } \mathcal{S}\left(1\left(\mathbf{A}_{\mathbf{R}_{1}}, \mathbf{A}_{\mathbf{R}_{2}}, \ldots \mathbf{A}_{\mathbf{R}_{\mathbf{k}}}\right),(1)\left(\mathbf{B}_{\mathbf{R}_{1}}, \mathbf{B}_{\mathbf{R}_{\mathbf{2}}}, \ldots \mathbf{B}_{\mathbf{R}_{\mathbf{k}}}\right)\right)= \\
& \left\{\begin{array}{l}
-\infty, \text { if } \exists \mathbf{A}_{\mathbf{R}}, \mathbf{B}_{\mathbf{R}^{\prime}} \text { s.t. } \mathcal{N}\left(\mathbf{A}_{\mathbf{R}}\right) \wedge \mathcal{N}\left(\mathbf{B}_{\mathbf{R}^{\prime}}\right) \\
\sum_{i=1}^{k} \mathcal{S}\left(\mathbf{A}_{\mathbf{R}_{\mathbf{i}}}, \mathbf{B}_{\mathbf{R}_{\mathbf{i}}}\right)
\end{array}\right.
\end{aligned}
$$
\]

Figure 4: Computing score for $\oplus$ operation. In the last rule bold font denotes sets.

We use prefix notation to denote co-trees. So (0) $\left(b_{1}, \ldots b_{n}\right)$ indicates a (0) node with children $b_{1}, \ldots, b_{n}$, where $b_{i}$ may denote a leaf (base tuple) or another node in the co-tree.

The first rule (the base case) returns a co-tree representing the disjunction of numerous base tuples. The third rule, of which the second rule is a helper, simply goes down a node in the lineage-tree when it encounters a $\sigma$ or a $\prod$ operator. Note that, rule 2 assumes that if the first tuple fed to $\mathcal{T}$ is a tuple produced by a selection (projection) then the rest of the tuples present in the argument list are also selection (projection) tuples. This assumption is guaranteed to hold because lineage-trees are produced by a query (plan) expressed in a high-level language (e.g., relational algebra).

The critical rule here is the fourth rule (join rule) which is a direct translation of Lemma 1 . We denote by $t^{b}$ a node in the lineage-tree (representing an intermediate tuple) that joins with base tuple $b$. In this case, given a set of join tuples we take each base tuple $b_{i}$ and form its co-tree with the co-tree of its cofactor represented by the set of intermediate tuples it joins with. It is in this rule that the real benefit of Lemma 1 shows up. If a cofactor contains a $P_{4}$ then we know that no other cofactor can provide a chord for it since cofactors are non-interfering. This implies that if a cofactor contains a $P_{4}$ then the result tuple contains a $P_{4}$ and we can bail out immediately. In fact, Lemma 1 guarantees that if we discover a $P_{4}$ during the process of merging co-trees produced out of cofactors then the result tuple is not $P_{4}$-free. Note that, we form a co-tree per base tuple. Once the various $\mathcal{T}$ calls return, we need to combine them and form the disjunction of the various co-trees. Even though the cofactors do not interfere (if the result tuple is read-once), they can still have variables in common which is why we use the $\oplus$ operator to merge these co-trees and form their disjunction. We describe details of the $\oplus$ operator in the next subsection.

Note that, in Figure 3, at no stage do we compute the DNF of any formula. Also, one of the limitations of the approach presented here is that it only applies to left/right deep plans. The example presented earlier in the section to motivate the need for choosing non-interfering formulas is the result of a plan with a bushy join (Figure 10). See Appendix C for ways to handle such plans.

### 3.3 Merging Co-Trees

In this section, we discuss how to merge the co-trees generated in the recursive step above (in other words, the implementation of the $\oplus$ operator from Rule 4). For this purpose, we maintain and use an augmented version of a co-tree where every edge from a parent node $u$ to child node $v$ is annotated with a set of relations $\mathcal{A}(u \rightarrow$ $v$ ). $R \in \mathcal{A}(u \rightarrow v)$ if some tuple from $R$ forms a leaf in the sub-
tree rooted at $v$. Figure 1(f) shows the annotated co-tree for result tuple $r$ in Figure 1(a). Two simple properties about annotated cotrees should be apparent. First, if $\exists R \in \mathcal{A}\left(u \rightarrow v_{1}\right), \mathcal{A}\left(u \rightarrow v_{2}\right)$, where $u$ is a (1) node, then this would imply a self-join so this is not possible (recall that (1) represents an $\wedge$ which corresponds to a join operation). Further, let $u$ denote any (0) node and $T$, the co-tree in which $u$ is present. If $T$ represents a k-monotone DNF, $\mathcal{A}\left(u \rightarrow v_{1}\right)$ has to be equal to $\mathcal{A}\left(u \rightarrow v_{2}\right)$ for any pair of children $v_{1}, v_{2}$.

We implement $\oplus$ as a binary operator using a recursive procedure. Let $T_{1}$ and $T_{2}$ denote the two input co-trees to be merged, then $\oplus\left(T_{1}, T_{2}\right)$ returns a pair consisting of a score and the resulting merged co-tree. Since the merged co-tree must necessarily have every variable as at most a single leaf, one of the main goals of $\oplus$ is to align the variables common to both $T_{1}$ and $T_{2}$. The score returned is simply the number of common variables aligned. Soundness of $\oplus$ follows the semantics of the various operations we subsequently describe. Completeness follows by showing that a $P_{4}$ exists if any of the assumptions we make do not hold. The score computation is described in Figure 4, where we use $\mathcal{S}\left(T_{1}, T_{2}\right)$ to denote the score of merging co-trees $T_{1}$ and $T_{2}$. We describe the creation of the merged co-tree in the text below. We bootstrap the merging process by computing two Boolean quantities for each node in either co-tree: $\mathcal{C}(u)$ is true if the subtree rooted at $u$ contains any variable common to both input co-trees, and $\mathcal{N}(u)$ is true if it contains a variable not present in the other co-tree. At the end of the merging process before returning the result, we check if the returned score matches the number of variables common to both co-trees.

The first rule in Figure 4 is one of the base cases: given a pair of leaves $x$ and $y$, it returns a score of 1 if they correspond to existence variables for the same tuple and 0 otherwise. In case they represent the same tuple, we return $x$, else we return (0) $(x, y)$ as the merged co-tree. The second rule is always the first thing we check for. Given $\mathcal{S}\left(T_{1}, T_{2}\right)$, we return the score 0 if the co-trees don't contain any common variables. The merged co-tree is simply (0) $\left(T_{1}, T_{2}\right)$. We may need to make minor adjustments to the returned co-tree so that it is canonical (0) and (1) nodes alternate along every path).

Rule 3 in Figure 4 computes the score for merging co-trees $T_{1}=$ (0) $\left(A_{1}, \ldots A_{n}\right)$ and $T_{2}=$ (1) $(\ldots)$, where each $A_{i}$ represents a sub-co-tree. Note that, $T_{1}$ and $T_{2}$ must share some common variables otherwise we would have applied rule 2 . Here, we match $T_{2}$ versus each subtree $A_{i}$ and pick the one whose score is the highest. If $A_{k}$ denotes the best match then the returned merged co-tree is (0) $\left(A_{1}, \ldots A_{k-1}, A_{k+1}, \ldots A_{n}, A\right)$ where $A$ denotes the co-tree obtained by merging $A_{k}$ and $T_{2}$. Note that we do not compare $T_{2}$ to sets of subtrees $\left\{A_{i}\right\}$. It can be shown that for a successful merge all common variables present in $T_{2}$ must be present in the same subtree $A_{k}$. In rule 4, we merge two co-trees $T_{1}=$ (0) $\left(A_{1}, \ldots A_{n}\right)$ and $T_{2}=$ (0) $\left(B_{1}, \ldots B_{m}\right)$. In this case also, we follow similar steps. For any $A_{i}$, we first check if $\mathcal{C}\left(A_{i}\right)$ is true. If yes, then we find the best matching subtree in $T_{2}$. We do this for every subtree $A_{i}$. The returned merge result is a co-tree rooted at a (0) node with $A_{i}$ as a subtree, if $\mathcal{C}\left(A_{i}\right)$ is false, otherwise we add the co-tree obtained by merging $A_{i}$ with $B_{k}$ as a subtree, where $B_{k}$ is the subtree from $T_{2} A_{i}$ best matches with. The same lemma, referred to above, applies. We do not need to match sets of subtrees of $T_{1}$ and $T_{2}$.

The last rule merges co-trees of the form $T_{1}=(1)\left(A_{\mathcal{A}_{1}}, \ldots A_{\mathcal{A}_{n}}\right)$ and $T_{2}=(1)\left(B_{\mathcal{A}_{1}}, \ldots B_{\mathcal{A}_{m}}\right)$, where the subscripts denote the annotations on the edges connecting roots of the co-trees to roots of the subtrees. In this case, we first attempt to form disjoint minimal subsets of sub-trees $\left\{A_{\mathcal{A}_{i_{1}}}, \ldots A_{\mathcal{A}_{i_{\mathcal{L}}}}\right\}$ and $\left\{B_{\mathcal{A}_{j_{1}}}, \ldots B_{\mathcal{A}_{j_{l}}}\right\}$ such that $\mathcal{A}_{i_{1}} \cup \ldots \mathcal{A}_{i_{k}}=\mathcal{A}_{j_{1}} \cup \ldots \mathcal{A}_{j_{l}}$. In other words, each pair of minimal subset of sub-trees should contain variables from the same (sub)set of relations. Let $\mathbf{A}_{\mathbf{R}}$ denote the minimal subset
of subtrees from $T_{1}$ involving variables from the set of relations $\mathbf{R}$, similarly, let $\mathbf{B}_{\mathbf{R}}$ denote its counterpart from $T_{2}$. For example, if $T_{1}=(1)\left(A_{R_{1}}, A_{R_{2}, R_{3}}\right)$ and $T_{2}=(1)\left(B_{R_{1}}, B_{R_{2}}, B_{R_{3}}\right)$ denote the input co-trees, then $\mathbf{A}_{\left\{R_{1}\right\}}=\left\{A_{R_{1}}\right\}$ and $\mathbf{B}_{\left\{R_{1}\right\}}=\left\{B_{R_{1}}\right\}$ and $\mathbf{A}_{\left\{R_{2}, R_{3}\right\}}=\left\{A_{R_{2}, R_{3}}\right\}$ and $\mathbf{B}_{\left\{R_{2}, R_{3}\right\}}=\left\{B_{R_{2}}, B_{R_{3}}\right\}$. Subsequently, we attempt to merge each pair $\mathbf{A}_{\mathbf{R}}$ and $\mathbf{B}_{\mathbf{R}}$. The returned score is the sum of scores returned by the recursive calls. If either of $\mathbf{A}_{\mathbf{R}}$ or $\mathbf{B}_{\mathbf{R}}$ is not a singleton set then we first create a dummy (1) root to form a single co-tree out of the non-singleton set before making the recursive merge call. It can be shown that if both $\mathbf{A}_{\mathbf{R}}$ and $\mathbf{B}_{\mathbf{R}}$ are non-singleton then a $P_{4}$ exists. It can also be shown that if $\mathbf{A}_{\mathbf{R}}$ and $\mathbf{B}_{\mathbf{R}^{\prime}}$ are such that $\mathcal{N}\left(\mathbf{A}_{\mathbf{R}}\right) \wedge \mathcal{N}\left(\mathbf{B}_{\mathbf{R}^{\prime}}\right)$ is true then the merge result has a $P_{4}$. Otherwise, the merge result is a co-tree rooted at a (1) node with the co-trees returned from the recursive calls as it child subtrees. An exception to the condition involving $\mathbf{A}_{\mathbf{R}}$ and $\mathbf{B}_{\mathbf{R}^{\prime}}$ described above is when none of $\mathbf{A}_{\mathbf{R}}, \mathbf{B}_{\mathbf{R}}, \mathbf{A}_{\mathbf{R}^{\prime}}$ or $\mathbf{B}_{\mathbf{R}^{\prime}}$ have any common variables. Here, we add (0) (1) $\left.\left(\mathbf{A}_{\mathbf{R}}, \mathbf{A}_{\mathbf{R}^{\prime}}\right),(1)\left(\mathbf{B}_{\mathbf{R}}, \mathbf{B}_{\mathbf{R}^{\prime}}\right)\right)$ as a subtree to the merge result.

Time Complexity: It can be shown that the time complexity of $\oplus\left(T_{1}, T_{2}\right)$ is $O\left(n m k^{2}\right)$, where $n$ and $m$ denote the number of nodes from the two co-trees and $k$ denotes the number of relations the leaves from either co-tree belong to. Note that, the most expensive rule in Figure 4 is rule 5, for which we need to do, at most, $k^{2}+k$ work (dividing the subtrees based on their annotations can be done in $k^{2}$ time). For the simpler case where no new dummy nodes get created (due to rule 5) during the merging process, applying this rule to every pair of nodes from both co-trees incurs a complexity of $O\left(n m k^{2}\right)$. For the more complex case when rule 5 generates dummy nodes also, it is possible to show that this complexity does not increase. However, based on our empirical observations (reported in the next section), this time complexity is rarely, if ever, achieved. In the next subsection, we describe a handful of optimizations that can further reduce the time spent to construct co-trees for result tuples. Coupled with the fact that to construct the co-tree we need simply make one pass over its lineage-tree, means that we have a reasonably efficient approach to computing the marginal probabilities of read-once result tuples produced by conjunctive queries without self-joins.

### 3.4 Further Optimizations

A number of optimizations are possible in the basic algorithms we presented in this section. One reason why our co-tree building algorithm (Figure 3) slows down is because of the join rule wherein we construct one co-tree per base tuple. Let $\mathbf{T}=\left\{b_{1} \bowtie\right.$ $\left.t_{1}, \ldots, b_{n} \bowtie t_{m}\right\}$ denote a set of (intermediate) join tuples. We say that $\mathbf{T}$ is block-structured if it can be rearranged into the form $B_{1} \bowtie T_{1}, \ldots, B_{k} \bowtie T_{j}$ where the $B_{i}$ s and $T_{i}$ s form a disjoint partitioning, or blocks, over the base tuples $\left\{b_{1}, \ldots b_{n}\right\}$ and the (intermediate) tuples $\left\{t_{1}, \ldots t_{m}\right\}$, respectively. A corollary of Lemma 3 is that, unless joins are block-structured, the result tuple contains a $P_{4}$. It is possible to check if a join is block-structured in linear time. If it is, then we call $\mathcal{T}$ on each block $T_{i}$ thus producing one co-tree per block which usually results in a smaller number of co-trees.

Another possible optimization lies in the merging process itself. One possible option is to initialize an empty tree and merge each co-tree, produced out of each block of base and intermediate tuples, into it (sequential merge). Another option is to arrange the perblock co-trees into a queue and then dequeue a pair of co-trees, merging them and enqueueing the result, in turn (pairwise merge). It is possible to trace the pairwise merging process by using a tree where the leaves form the co-trees produced by the join rule and the intermediate nodes represent co-trees produced by merging a pair from the queue. Empirically, we noticed pairwise merge to be

(a)

(b)

(c)

Figure 5: Results of evaluating (a) two, (b) three and (c) four relation Boolean conjunctive queries on synthetic databases generated with $d=150, \ldots 250$. Note that, $\mathbf{y}$-axis is log-scale.
significantly faster than sequential merge.
A third optimization is possible in the implementation of $\oplus$. Recall from Figure 4 that, in various rules we check if sub-trees contain common variables before calling $\mathcal{S}$ on them (in Rule 4, for instance). In our implementation, we never call $\mathcal{S}$ on two sub-trees unless they contain the same exact number of common variables. This does not in any way affect the correctness of our algorithm since a prerequisite of a successful alignment of two co-trees is that they contain the same number of common variables. Instead, the count of common variables acts as a much more stringent filter, cutting down on the number of calls made to $\mathcal{S}$. An empirical check to determine the effectiveness of this filter showed that in general, each call to $\oplus$ to merge two co-trees of sizes $n$ and $m$ nodes gave rise to $\min (n, m)$ recursive calls which represents a favourable case (much smaller than $O(n m)$ ).

## 4. EXPERIMENTS

Our aim is to evaluate how our approach compares with state-of-the-art probabilistic databases with respect to query evaluation time and whether there is any merit to the paradigm of per-result tuple evaluation. To demonstrate scalability, we experiment with synthetic databases of varying sizes. We show that, not only do we far outperform current probabilistic databases when evaluating non-hierarchical queries, but that our techniques also compare favourably to query compilation techniques specifically developed to handle hierarchical queries. We also ran experiments using the TPC-H benchmark [28]. Five of the first TPC-H queries are not hierarchical. Of these, one query's result solely comprised of readonce functions and two others generated $\approx 50 \%$ read-once result tuples thus lending credence to the paradigms of database instanceoptimal and per-result tuple query evaluation.

We compare our approach (henceforth referred to as CoTree) against two state-of-the-art probabilistic databases, MystiQ [4] and SPROUT [24], and the naive approach (Section 2.3) referred to as CoOccur using a prior $P_{4}$ checking algorithm [7]. MystiQ contains an implementation of Dalvi et al. [9]'s hierarchical query recognition/safe plan generation technique. For SPROUT, we only consider the $\operatorname{conf}()$ function for comparison since we are mainly interested in exact confidence computation.

To keep our experimental methodology simple, we only experiment with Boolean conjunctive queries (final set of projected attributes is empty) and hence do not perform the normality check for CoOccur. To generate tuples for the synthetic dataset, we take a domain size $d$ as a parameter and randomly generate attribute values from that domain. For TPC-H queries, we use the dbgen program [28] to generate the database. Tuple probabilities are picked randomly. Our code is written in JAVA. All experiments were run on an Intel 2.26 GHz Core 2 Duo machine with 2GB RAM.
Synthetic Dataset: Figures 5(a), (b) and (c) depict the results of


Figure 6: Results on the TPC-H benchmark. "**" indicates the approach ran out of memory. Note that, y -axis is $\log$-scale.
running the four approaches on two, three and four relation Boolean join queries over synthetic databases generated with $d$ varying from $150, \ldots 250$. All three plots report runtimes (in seconds) on the $y$ axis averaged over ten different queries and number of tuples per relation (in thousands) along the x -axis. Partial plots indicate the approach ran out of memory or took too long to run. For the case of two relation joins (Figure 5 (a)), the queries are guaranteed to be hierarchical [11] so MystiQ does very well. The interesting thing in this case is that CoTree matches MystiQ's performance toe-to-toe. CoOccur's performance deteriorates drastically with larger databases because the result tuples produced have larger and denser co-occurrence graphs. Perhaps surprisingly, SPROUT performs worst even though it has specialized algorithms to handle hierarchical queries. The likely reason ${ }^{\dagger}$ for this is that SPROUT prefers plans which defer marginal probability computation till the end (lazy plans [24]) as opposed to MystiQ, which computes marginals as soon as an (intermediate) tuple is produced (eager plans [24]). In this case, the eager plan works better than the lazy plan. Note that, SPROUT is capable of running both lazy and eager plans. For three and four relation joins, Figures 5 (b) and (c) respectively, it is no longer necessary that the queries be hierarchical. However, it is possible to extend the idea behind the running example (Figure 1 (a)) to larger relations and larger number of joins to generate readonce result tuples. For these queries, MystiQ immediately switches to its approximate inference method (a slow MCMC technique). SPROUT fares the worst in these cases also (SPROUT is grazing the y -axis which is in log-scale). CoTree performs the best, providing performance that is orders of magnitude faster than any of the other approaches. Varying $d$ alters performance in expected ways. Upon reducing $d$ (thus increasing join selectivity), the cooccurrence graphs of the result tuples become denser thus deteriorating performance of SPROUT, MystiQ and CoOccur even further, while increasing $d$ reduces the difference among the approaches.
TPC-H: We also experimented with queries Q1 through Q10 in the TPC-H specification on a 0.1 GB database. Since we do not consider running aggregation nor self-joins, we modified the queries slightly. We added random probabilities to all TPC-H relations. This leads to Q2, Q5, Q7, Q8 and Q9 being non-hierarchical. Interestingly, despite the queries being non-hierarchical, all of Q2's result tuples turned out to be read-once, for which SPROUT, CoOccur and CoTree were able to compute marginals easily. Moreover, Q8 and Q9's result comprise $52 \%$ and $43.4 \%$ read-once result tuples, respectively. This illustrates the merit of evaluating on a perresult tuple basis. However, all of these queries still present fairly easy marginal probability computation problems (no query spent more than $15 \%$ of the time computing probabilities). Thus, we picked the queries with more than 1 joining relation and modified them even further by dropping attributes from the SELECT clause and/or predicates from the WHERE clause all the while ensuring that all generated result tuples were read-once. Figure 6 denotes

[^2]the performance of the four approaches on these modified queries where we use bars to denote run-time and Qi' denotes the modified query obtained from TPC-H Query Qi. In Figure 6, except for Q4' (a two relation join query), all other queries were non-hierarchical. MystiQ ran out of memory while running Q5' and Q7', and both MystiQ and CoOccur ran out of memory while running Q9'. The missing bars are denoted by "*". CoTree performs best on all the queries, with CoOccur coming a close second. For five of these queries, CoTree outperforms SPROUT by at least an order of magnitude and on two, CoTree is two orders of maginitude faster.
Discussion: Besides the experiments described above, we extensively evaluated our co-tree merging algorithm and our pairwise approach to merging multiple co-trees. We also conducted an experiment to check if we are repeating work while factorizing each result tuple in isolation by possibly rediscovering the same query plan. Sometimes, by running the correct query plan, the lineagetree of the result tuple itself turns out to be its co-tree (this is the main idea behind MystiQ [9]). To check for this, we took the three relation join Boolean query (Figure 5 (b)) and ran it with all possible query plans (various early projections, various join orders). For each query plan, we measured the hardness of computing the result tuple's marginal probability from the generated lineage-tree by measuring treewidth [2]. If a query plan recovered the co-tree then the treewidth would be 1 . Invariably, the query plans produced very hard inference problems with treewidth sometimes approaching 250 . This implies that to exactly evaluate factorizable result tuples, only employing query rewriting techniques is insufficient and our approach of considering each result tuple separately is justified.

## 5. RELATED WORK

Olteanu and Huang [22, 23] showed that hierarchical queries can be extended to include inequality operators, $\neq,\rangle,<$, in certain cases. These queries may not produce read-once functions. Dalvi et al. [8] propose an approach to efficiently handle disjunctive queries with self-joins. Jha et al. [19] consider unifying hierarchical queries with graphical models-based query evaluation [27]; these techniques however, do not cover the class of read-once functions. Olteanu et al. [25] use a generalized version of cotrees known as $d$-trees but require input in DNF to construct them. Darwiche [12] proposes using Boolean formula factorization algorithms to compile a probabilistic model into a more tractable form. However, he relies on the use of an exponential-sized intermediate representation called multi-linear formula. Our work on merging co-trees is related to tree alignment [20] which comes in two flavours, ordered and unordered. Ordered trees have an ordering defined over the children as opposed to unordered trees. Unordered tree alignment is MAX SNP-hard [20]. So it is noteworthy that, even though our co-trees contain unordered nodes (0) nodes), we can still merge them in polynomial time.

## 6. CONCLUSION

In summary, we considered the problem of efficiently evaluating queries over tuple-level uncertainty probabilistic databases. Earlier approaches have mainly concentrated on query-centric notions of solvability (hierarchical queries). In this paper, we went beyond just looking at the query to decide whether it is PTIME-solvable or not. We first showed how to incorporate Boolean formula factorization techniques into the query engine that exploit structure present in both query and data to efficiently evaluate result tuples' marginal probabilities. We also proposed novel, especially efficient algorithms to evaluate a large class of queries, viz. conjunctive queries without self-joins. We empirically showed that our proposed techniques are much faster than prior Boolean formula fac-
torization techniques, techniques specifically developed for hierarchical queries, and generic inference algorithms.

Even though our discussion mainly involved tuple-independent probabilistic databases, the techniques we proposed are likely to be useful for databases with correlated tuples also. In that case, by plugging in the factorized co-tree generated by our algorithms into the graphical model describing the correlations among the base tuples [27], we should be able to obtain a combined graphical model of lower treewidth. As part of our future work, we intend to extend our techniques to query probabilistic databases with tuple and attribute uncertainty. Also, we would like to explore various generalizations such as P4-tidy graphs [14] to see if these can be advantageously used for querying probabilistic databases.
Acknowledgments: We would like to thank the anonymous reviewers for their suggestions and Dan Olteanu for patiently answering all our questions about the SPROUT codebase. This work was supported in part by NSF grants IIS-0546136 and IIS-0916736.

## 7. REFERENCES

[1] C. C. Aggarwal. Managing and Mining Uncertain Data. Springer, 2009.
[2] S. Arnborg. Efficient algorithms for combinatorial problems on graphs with bounded decomposability - a survey. In BIT, 1985.
[3] O. Benjelloun, A. Das Sarma, A. Halevy, and J. Widom. ULDBs: Databases with uncertainty and lineage. In $V L D B, 2006$.
[4] J. Boulos, N. Dalvi, B. Mandhani, C. Re, S. Mathur, and D. Suciu. Mystiq: A system for finding more answers by using probabilities. In SIGMOD, 2005.
[5] A. Bretscher, D. Corneil, M. Habib, and C. Paul. A simple linear time lexbfs cograph recognition algorithm. SIAM Journal on Discrete Mathematics, 2008.
[6] D. Corneil, H. Lerchs, and L. Burlingham. Complement reducible graphs. Discrete Applied Mathematics, 1981.
[7] D. Corneil, Y. Perl, and L. Stewart. A linear recognition algorithm for cographs. SIAM Journal of Computing, 1985.
[8] N. Dalvi, K. Schnaitter, and D. Suciu. Computing query probability with incidence algebras. In PODS, 2010.
[9] N. Dalvi and D. Suciu. Efficient query eval. on prob. databases. In $V L D B$, ' 04
[10] N. Dalvi and D. Suciu. Management of probabilistic data: foundations and challenges. In PODS, 2007.
[11] N. N. Dalvi and D. Suciu. The dichotomy of conjunctive queries on probabilistic structures. In PODS, 2007.
[12] A. Darwiche. A logical approach to factoring belief networks. In Intl. Conf. on Principles and Know. Rep. and Reasoning, 2002.
[13] N. Fuhr and T. Rolleke. A probabilistic relational algebra for the integration of information retrieval and database systems. ACM Trans. on Info. Systems, 1997
[14] V. Giakoumakis, F. Roussel, and H. Thuillier. On P4-tidy graphs. Discrete Maths and Theoretical Comp. Science, 1997.
[15] M. Golumbic, A. Mintz, and U. Rotics. Factoring and recognition of read-once functions using cographs and normality and the readability of functions associated with partial $k$-trees. Discrete Applied Mathematics, 2006.
[16] M. Habib and C. Paul. A simple linear time algorithm for cograph recognition. Discrete Applied Mathematics, 2005.
[17] J. Hayes. Fanout structure of switching functions. JACM, 1975.
[18] 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, 2008.
[19] A. Jha, D. Suciu, and D. Olteanu. Bridging the gap between intensional and extensional query evaluation in probabilistic databases. In EDBT, 2010.
[20] T. Jiang, L. Wang, and K. Zhang. Alignment of trees - an alternative to tree edit. In Annual Symposium on Combinatorial Pattern Matching, 1994.
[21] C. Koch. Approximating predicates and expressive queries on probabilistic databases. In PODS, 2008.
[22] D. Olteanu and J. Huang. Using OBDDs for efficient query evaluation on probabilistic databases. In SUM, 2008.
[23] D. Olteanu and J. Huang. Secondary-storage conf. computation for conjunctive queries with inequalities. In SIGMOD, 2009.
[24] D. Olteanu, J. Huang, and C. Koch. SPROUT: Lazy vs. eager query plans for tuple-independent probabilistic databases. In ICDE, 2009.
[25] D. Olteanu, J. Huang, and C. Koch. Approximate confidence computation in probabilistic databases. In ICDE, 2010.
[26] C. Ré and D. Suciu. The trichotomy of having queries on a probabilistic database. VLDB J., 18(5):1091-1116, 2009.
[27] P. Sen, A. Deshpande, and L. Getoor. Prdb: managing and exploiting rich correlations in probabilistic databases. VLDB J., 18(5):1065-1090, 2009.
[28] TPC-H Benchmark. http://www.tpc.org/.

## APPENDIX

## A. PROOFS

Theorem 2 crucially depends on the following lemma where $x-$ $y \in G$ denotes variables $x$ and $y$ form an edge in $G$.

Lemma 2. Let $C_{1}, C_{2}$ denote two clauses in $\phi_{D N F}$ s.t. $C_{1} \cap$ $C_{2}=\mathbf{z}=\left\{z_{1}, \ldots z_{l}\right\}, x \in C_{1}, x \notin C_{2}, y \notin C_{1}, y \in C_{2}$. If $\phi$ is $P_{4}$-free and $x-y \in G_{\phi}$ then $\exists C_{3}$ in $\phi_{D N F}$ s.t. $\left\{x, y, z_{1}, \ldots z_{k}\right\} \subseteq$ $C_{3}$.

Proof. Let $q$ denote the conjunctive query (without self-joins) that produced result tuple $\phi$. Notice that $x, y, \mathbf{z}$ form $l+2$ sized clique in $G_{\phi}$. If $|\operatorname{Rels}(q)|=l+2$ then this clique will form a clause in which case we have found $C_{3}$. We will assume this is not the case and $|\operatorname{Rels}(q)|>l+2$. The basic idea is to show that a selection of the remaining variables from $C_{1}$ and $C_{2}$ along with $x, y$ and $\mathbf{z}$ form a $|\operatorname{Rels}(q)|$-sized clique in $G_{\phi}$. This, in turn will imply the presence of the clause containing $x, y$ and $\mathbf{z}$. Let us first complete the cliques $C_{1}$ and $C_{2}$ :

- $C_{1} \backslash(\{x\} \cup \mathbf{z})=\left\{y^{\prime}, w_{1}, \ldots w_{n}\right\}$
- $C_{2} \backslash(\{y\} \cup \mathbf{z})=\left\{x^{\prime}, v_{1}, \ldots v_{n}\right\}$
where $w_{i} \not \equiv v_{i}$ and $w_{i}, v_{i} \in R_{i} \in \operatorname{Rels}(q) \forall i=1 \ldots n$ (re-index $\operatorname{Rels}(q)$ if necessary). First, notice that one of the edges $x-v_{i}$ or $y-w_{i}$ (or both) has to exist in $G_{\phi} \forall i=1 \ldots n$ if $\phi$ is to be $P_{4}$-free else $w_{i}-x-y-v_{i}$ is a $P_{4}$ ( $w_{i}$ and $v_{i}$ cannot form an edge because they belong to the same relation $R_{i}$ ). This gives us a selection procedure: if $x-v_{i}$ then pick $v_{i}$ else pick $w_{i}$. Further, two variables $v_{i}$ and $w_{j}(i \neq j)$ selected this way must have an edge between them otherwise $\phi$ has a $P_{4}: w_{j}-x-v_{i}-v_{j}\left(v_{j}\right.$ and $x$ do not form an edge otherwise we would have selected $v_{j}$ not $w_{j}$ ). Thus, $x, y, \mathbf{z}$ along with the selected $w$ 's and $v$ 's form a $|\operatorname{Rels}(q)|$-sized clique in $G_{\phi}$.

Theorem 2. Let $\phi$ denote the result tuple produced by a conjunctive query $q$ without self-joins. If $\phi$ is $P_{4}$-free, it is also normal.

Proof. To prove the proposition, we need to show that any clique in $G_{\phi}$ is contained in some clause in $\phi_{D N F}$. Notice that a clique of size greater than $|\operatorname{Rels}(q)|$ is not possible since this would imply a self-join in $q$ and a clique of size equal to $|\operatorname{Rels}(q)|$ directly translates to a clause in $\phi_{D N F}$. Proving the proposition for cliques of size $<|\operatorname{Rels}(q)|$ is slightly trickier. We prove this by induction on the size $k$ of the clique. The base case is when $k=3$ ( $k=2$ is simply an edge which has to be contained in a clause). Let $\{a, b, c\}$ denote the clique. Assume that clause $C_{1}$ contains $a$ and $b$ but not $c$ and $C_{2}$ contains $b$ and $c$ but not $a$. (If either $C_{1}$ contained $c$ or $C_{2}$ contained $a$ then we have found the clause we are looking for and the base case is proved.) Now we invoke Lemma 2 with $x=a, y=c$ and $\mathbf{z}=\{b\}$. Thus, there must exist a clause containing all three variables. For the inductive case, we will assume that any $k-1$ sized clique in $G_{\phi}$ is fully contained in some clause in $\phi_{D N F}$. Now we need to show that given a $k$-sized clique $\left\{a_{1}, \ldots a_{k}\right\}$, a clause containing it exists. Notice that $\left\{a_{1}, \ldots a_{k-1}\right\}$ forms a $k-1$ sized clique and thus there must be a clause $C_{1}$ containing it (by inductive hypothesis). Similarly, there must also be a clause $C_{2}$ containing $\left\{a_{2}, \ldots a_{k}\right\}$. Notice that if $a_{k} \in C_{1}$ or $a_{1} \in C_{2}$ then we have shown what was needed. Assuming that is not true, invoke Lemma 2 with $x=a_{1}, y=a_{k}$ and $\mathbf{z}=\left\{a_{2}, \ldots a_{k-1}\right\}$. Thus, a clause containing $\left\{a_{1}, \ldots a_{k}\right\}$ must exist.

Lemma 1. Let $\phi$ denote a result tuple, $\mathcal{L}$ its lineage-tree and $\Gamma$ a join path. If $\phi$ is $P_{4}$-free then the set of cofactors $\left\{\operatorname{cof}_{\mathcal{L}}([\Gamma, b]) \mid b \in\right.$ $R\}$ is non-interfering.

Proof. It is straightforward to show that type 1 interference is not possible. Let us assume that $\operatorname{cof}_{\mathcal{L}}\left(\left[\Gamma, b_{i}\right]\right)$ and $\operatorname{cof}_{\mathcal{L}}\left(\left[\Gamma, b_{j}\right]\right)$ interfere, where $b_{i}, b_{j} \in R$. In other words, $\exists x, y$ such that $x-$ $y \in G_{c o f_{\mathcal{L}}\left(\left[\Gamma, b_{i}\right]\right)}$ and $x \neq y \in G_{c o f_{\mathcal{L}}\left(\left[\Gamma, b_{j}\right]\right)}$ even though $x, y \in$ $\operatorname{Vars}\left(\operatorname{cof_{\mathcal {L}}}\left(\left[\Gamma, b_{j}\right]\right)\right)$. This implies that the DNF of $\operatorname{cof}_{\mathcal{L}}\left(\left[\Gamma, b_{j}\right]\right)$ must contain two distinct clauses one of which contains $x$ but not $y$ and the other $y$ but not $x$. Coupled with the fact that $\operatorname{cof}_{\mathcal{L}}\left(\left[\Gamma, b_{j}\right]\right)$ joins with $b_{j}$ and tuples in join path $\Gamma$, this implies that $\phi_{D N F}$ contains two distinct clauses one of which contains the set $\left\{\gamma_{1}, \ldots \gamma_{k}, b_{j}, x\right\}$ and the other $\left\{\gamma_{1}, \ldots \gamma_{k}, b_{j}, y\right\}$, where $\gamma_{i}$ denotes a tuple from $\Gamma$. Since $x-y \in G_{c o f_{\mathcal{L}}\left(\left[\Gamma, b_{i}\right]\right)}$, we can now set $\mathbf{z}=\left\{\gamma_{1}, \ldots \gamma_{k}, b_{j}\right\}$ and invoke Lemma 2 to claim the presence of a clause in $\phi_{D N F}$ that contains all of $\gamma_{1}, \ldots \gamma_{k}, b_{j}, x, y$. If that is true, then there should be a clause in DNF of $\operatorname{cof_{\mathcal {L}}}\left(\left[\Gamma, b_{j}\right]\right)$ that contains both $x$ and $y$ and these two variables should form an edge in $G_{c o f_{\mathcal{L}}\left(\left[\Gamma, b_{j}\right]\right)}$ which is not true. Hence we have a contradiction.

For type 2 interference, assume that $\exists x, y$ such that $x-y \in$ $G_{c o f_{\mathcal{L}}\left(\left[\Gamma, b_{h}\right]\right)}, x \in \operatorname{Vars}\left(\operatorname{cof}_{\mathcal{L}}\left(\left[\Gamma, b_{i}\right]\right)\right), x \notin \operatorname{Vars}\left(\operatorname{cof}_{\mathcal{L}}\left(\left[\Gamma, b_{j}\right]\right)\right)$, $y \notin \operatorname{Vars}\left(\operatorname{cof}_{\mathcal{L}}\left(\left[\Gamma, b_{i}\right]\right)\right)$ and $y \in \operatorname{Vars}\left(\operatorname{cof}_{\mathcal{L}}\left(\left[\Gamma, b_{j}\right]\right)\right)$, where $b_{h}, b_{i}, b_{j}$ belong to the same relation $R$. Note that, this gives us a $P_{4}$ in $\phi: b_{i}-x-y-b_{j}$, unless $b_{i}-y$ or $x-b_{j}$ exists in $G_{\phi}$ ( $b_{i}$ and $b_{j}$ belong to the same relation so they cannot form an edge). Without loss of generality, let us assume $b_{i}-y$ exists in $G_{\phi}$ so we do not have the $P_{4}$. Clearly, this edge is not present in $G_{c o f_{\mathcal{L}}\left(\left[\Gamma, b_{i}\right]\right)}$ (since $\operatorname{cof}_{\mathcal{L}}\left(\left[\Gamma, b_{i}\right]\right)$ does not contain $y$ ) so it has to come from some other part of $\mathcal{L}$. We now backtrack along $\Gamma=\gamma_{1}, \ldots \gamma_{k}$ in the reverse direction. Let $k^{\prime} \leq k$ denote the maximum index such that $\exists \gamma_{k^{\prime}}^{\prime} \neq \gamma_{k^{\prime}}$ and $b_{i}-y \in$ $G_{c o f_{\mathcal{L}}\left(\left[\gamma_{1}, \ldots \gamma_{k^{\prime}-1}, \gamma_{k^{\prime}}^{\prime}\right]\right)}$. Note that, $\operatorname{cof_{\mathcal {L}}([\gamma _{1},\ldots \gamma _{k},b_{i}])\text {isasub-}}$ formula of $\operatorname{cof}_{\mathcal{L}}\left(\left[\gamma_{1}, \ldots \gamma_{k^{\prime}-1}, \gamma_{k^{\prime}}\right]\right)$ (due to the recursive nature of cofactors). Let $\Gamma^{\prime}$ denote the join path $\gamma_{1}, \ldots \gamma_{k^{\prime}-1}$ then this implies $\operatorname{cof} f_{\mathcal{L}}\left(\Gamma^{\prime}, \gamma_{k^{\prime}}\right)$, which does not contain an edge between $y$ and $b_{i}$, and $\operatorname{cof}_{\mathcal{L}}\left(\Gamma^{\prime}, \gamma_{k^{\prime}}^{\prime}\right)$, which does, exhibit type 1 interference. As we have already seen, this is not possible. Hence we have a contradiction.

Lemma 3 (Block-structured Co-factors). Let $\Gamma$ denote a join path in lineage-tree $\mathcal{L}$ of result tuple $\phi$. If $\phi$ is $P_{4}$ free then $\operatorname{cof}_{\mathcal{L}}([\Gamma])=b_{1} \operatorname{cof}_{\mathcal{L}}\left(\left[\Gamma, b_{1}\right]\right)+\ldots+b_{n} \operatorname{cof}_{\mathcal{L}}\left(\left[\Gamma, b_{n}\right]\right)$, where $b_{i}$ denotes a base tuple, is block-structured. In other words, if $\exists$ (intermediate) tuples $t_{i}, t_{j}$ such that $t_{i} \in \operatorname{cof} f_{\mathcal{L}}\left(\left[\Gamma, b_{i}\right]\right), t_{i} \in$ $\operatorname{cof}_{\mathcal{L}}\left(\left[\Gamma, b_{j}\right]\right)$ but $t_{j} \in \operatorname{cof}_{\mathcal{L}}\left(\left[\Gamma, b_{i}\right]\right), t_{j} \notin \operatorname{cof}_{\mathcal{L}}\left(\left[\Gamma, b_{j}\right]\right)$ then $\phi$ has a $P_{4}$.

Proof. Notice that, since $b_{j}$ joins with $t_{i}, G_{\phi}$ contains an edge between $b_{j}$ and any variable present in the formula of $t_{i}$. Also, for a variable $x$ that is present in $t_{j}$ but not in $t_{i}$ we have a $P_{4}$ in $\phi$ if $b_{j} \nsucc x \in G_{\phi}$ (the $P_{4}$ is $b_{j}-x^{\prime}-b_{i}-x$, where $x^{\prime}$ belongs to the same relation as $x$ and is present in $t_{i}$ ). Thus, unless $b_{j}$ is connected to all variables in $t_{j}$ in $G_{\phi}, \phi$ has a $P_{4}$ and there is nothing left to prove. Now, we go back to $\operatorname{cof} \mathcal{L}_{\mathcal{L}}(\Gamma)$. Since $\Gamma$ joins with $b_{j}$ and $\Gamma$ joins with any clause in DNF of $t_{j}$, we can apply Lemma 2 repeatedly to show that $b_{j}$ joins with $t_{j}$ obtaining a contradiction.

## B. LEMMAS FOR MERGING ALGORITHM

Recall from Section 3.2 that, if two cofactors interfere then the result tuple has a $P_{4}$. Also recall that, in the description of the $\oplus$ operator (Section 3.3) we check to see if the returned score matches the number of common variables. This check is a crucial step because it helps detects $P_{4}$ s in the result tuple. It is possible to show (by tracing $\mathcal{S}$ ) that if $T_{1}$ and $T_{2}$ (co-trees being merged) are such that $x-y \in G_{\phi_{T_{1}}}$ and $x, y \in \operatorname{Vars}\left(T_{2}\right)$ but $x \not f y \in G_{\phi_{T_{2}}}$ then the returned score will always be less than the number of variables


Figure 8: When merging two co-trees rooted at (1) nodes, the child sub-trees must be divisible into sets that cater to the same relations, else the merging has a $P_{4}$. Moreover, at least one such subset should be singleton.


Figure 7: Separated common nodes indicate a $P_{4}$.
common to $T_{1}$ and $T_{2}$. A simple intuitive explanation is by considering least common ancestors (LCA) of the two variables in the co-trees $T_{1}$ and $T_{2}$. One of the basic properties of a co-tree is that given a pair of variables, their ancestor is a (1) node if the variables form an edge and a (0) node, if not [15]. This implies that the LCA of $x$ and $y$ in $T_{1}, L C A_{1}$, is a (1) node and their LCA in $T_{2}, L C A_{2}$, is a (0) node. Now, as per merging rules 3 and 4 (Figure 4), we only match one rooted subtree against another. This implies that whenever we match the subtrees rooted at $L C A_{1}$ and $L C A_{2}$ we are bound to misalign at least one of either $x$ or $y$ thus detecting that the result tuple contains a $P_{4}$. In what follows, we prove three lemmas required to justify our algorithm for $\oplus\left(T_{1}, T_{2}\right)$ presented in Section 3.3. In none of these proofs do we consider cases leading to edge discrepancies where two variables $x, y$ form an edge in $T_{1}$ $\left(T_{2}\right)$ and do not form an edge in $T_{2}\left(T_{1}\right)$ because, as we just showed, these cases lead to $P_{4}$ s and by comparing the returned score we can easily capture such cases.

LEMMA 4. Let $T_{1}=(0)\left(A_{1}, \ldots A_{n}\right)$ such that $x \in \operatorname{Vars}\left(A_{i}\right)$ and $y \in \operatorname{Vars}\left(A_{j}\right)$, where $i \neq j$. Let $T_{2}$ denote another co-tree rooted at (1) node such that $x, y \in \operatorname{Vars}\left(T_{2}\right)$. If $T_{1}$ and $T_{2}$ involve variables from the same set of relations then $\phi_{T_{1}} \vee \phi_{T_{2}}$ contains $a$ $P_{4}$.

Proof. If LCA of $x$ and $y$ in $T_{2}$ is a (1) node then we have an edge discrepancy between $T_{1}$ and $T_{2}$ which implies interference and a $P_{4}$. Assuming LCA of $x$ and $y$ in $T_{2}$ is a (0) node, Figure 7 shows the setup described in the lemma where $L C A$ denotes LCA of $x$ and $y$ in $T_{2}$ and $\alpha$ denotes root of $T_{2}$. Dashed edges indicate ancestor-child (not necessarily parent-child) relationships. Notice that, besides $L C A$, $\alpha$ must have at least one other descendant (oth-
erwise there is no point of having $\alpha$ ). Let $c_{1}$ denote one such descendant leaf and let $R$ denote the relation $c_{1}$ belongs to. Now, the merging rules (Figure 4) ensure that sub-trees being aligned involve variables from the same relations. This implies that variables from $R$ are present in both child sub-trees of $A$. Notice that $c_{1}$ 's presence in either of these sub-trees will lead to edge discrepancies that would indicate the presence of a $P_{4}$ and there would be nothing left to prove. Thus, let $c_{2}\left(\not \equiv c_{1}\right)$ denote the variable from $R$ that is present in the child sub-tree containing $x$, and let $c_{3}\left(\not \equiv c_{1}\right)$ denote the variable from $R$ that is present in the child sub-tree containing $y$. Again, neither $c_{2}$ nor $c_{3}$ can be present in $B$ because if they were then we would have edge discrepancies indicating a $P_{4}$ and there would be nothing left to prove. Now, we have a $P_{5}: c_{2}-x-c_{1}-y-c_{3}$ that contains two $P_{4}$ 's.

The next two lemmas relate to rule 5 in Figure 4. The first one shows that when aligning two co-trees rooted at (1) nodes, the child sub-trees need to be such that we can divide them into minimal subsets such that each pair of subsets involves variables from the same (sub)set of relations (see Section 3.3). This will not be possible if there exists relations $R_{1}, R_{2}, R_{3}$, two child sub-trees in the first cotree rooted at $a_{1}, a_{2}$ and two child sub-trees in the second co-tree rooted at $b_{1}, b_{2}$ such that $R_{1} \in \mathcal{A}\left(a \rightarrow a_{1}\right),\left\{R_{2}, R_{3}\right\} \subseteq \mathcal{A}(a \rightarrow$ $\left.a_{2}\right),\left\{R_{1}, R_{2}\right\} \subseteq \mathcal{A}\left(b \rightarrow b_{1}\right), R_{3} \in \mathcal{A}\left(b \rightarrow b_{2}\right)$ where $a$ and $b$ denote the roots of the two input co-trees, respectively. The proof of the lemma considers various cases that satisfy this property and shows that a $P_{4}$ exists in each case. The second lemma deals with the placement of variables that are not present in both co-trees.

LEMMA 5. Let $a$ and $b$ denote two (1) roots of two co-trees $A$ and $B$, respectively. Let $a_{1}, a_{2}$ and $b_{1}, b_{2}$ denote two children of $a$ and $b$, respectively, such that there is at least one common variable present in the sub-trees rooted at $a_{1}, a_{2}$ and $b_{1}, b_{2}$. If $\exists R_{1}, R_{2}, R_{3}$ s.t. $R_{1} \in \mathcal{A}\left(a \rightarrow a_{1}\right),\left\{R_{2}, R_{3}\right\} \subseteq \mathcal{A}\left(a \rightarrow a_{2}\right)$, $\left\{R_{1}, R_{2}\right\} \subseteq \mathcal{A}\left(b \rightarrow b_{1}\right), R_{3} \in \mathcal{A}\left(b \rightarrow b_{2}\right)$ then the merging of $A$ and $B$ contains a $P_{4}$.

PROOF. The assumption of sub-trees rooted at $a_{1}, a_{2}$ and $b_{1}, b_{2}$ having at least one common variable is important because if this is not the case then we have work arounds. For instance, if the input co-trees rooted at $a$ and $b$ do not have any variable in common then we apply merging rule 2 not 5 (Figure 4 ). Figure 8 explains the proof. In Figure $8, x, x_{1}, x_{2}$ denote variables from $R_{1}, y, y_{1}, y_{2}$ de-

if $l c a$ is (0) $P_{4}: v^{*}-u_{1}-x-u^{*}$
if $l c a$ is (1) $P_{4}: v_{1}-u^{*}-x-v^{*}$

$P_{4}: v_{1}-u_{1}-x-u^{*}$




Figure 9: Proof showing that merging co-trees after application of rule 5 in Figure 4 is simple. If the presence of new variables complicate things then a $P_{4}$ exists in the merge result.
note variables from $R_{2}$ and $z, z_{1}$ denote variables from $R_{3}$. There are three cases depicted in Figure 8:

- In the first case, we assume that $x$ and $z$ are common ( $y$ may or may not be common), i.e., $x$ is present in both sub-trees rooted at $a_{1}$ and $b_{1}$, and $z$ is present in both sub-trees rooted at $a_{2}$ and $b_{2}$. Note that in this case, if $y_{1}$ is present under $b_{1}$ or $y_{2}$ under $a_{2}$ then we will introduce an edge discrepancy, in which case we know there is a $P_{4}$ and there is nothing left to prove. The figure depicts the $P_{4}$ when $y_{1}$ exists only under $a_{2}$ and $y_{2}$ exists only under $b_{1}$.
- In the second case, we assume only one of $x$ or $z$ are common. Since both cases are symmetric, we assume $x$ is common. Notice that in this case, a common variable $y$ can exist in both sub-trees rooted at $a_{2}$ and $b_{1}$ but this does not contradict the existence of $y_{2}$ which cannot exist in the sub-tree rooted at $a_{2}$ without introducing an edge discrepancy. The figure depicts the $P_{4}$ obtained.
- In the last case, we assume $y$ is the only variable common to both co-trees. If either $x$ or $z$ or both are also common then we fall under the previous cases. Since there are not common variables $x$ or $z$ present in sub-trees rooted at $a_{1}, a_{2}$ and $b_{1}, b_{2}, x_{1}, x_{2}, z_{1}$ must exist such that $x_{1}$ is not present in the sub-tree rooted at $b_{1}, x_{2}$ does not exist in the sub-tree rooted at $a_{1}$ and $z_{1}$ does not exist in the sub-tree rooted at $a_{2}$. This gives us a $P_{4}$ (see Figure 8).

The second lemma which deals with merging rule 5 in Figure 4, has implications on how we actually merge co-trees formed by the various recursive calls once they have returned. We first illustrate an example that explains the kind of problem we might run into had the following lemma not been true. Assume that $\left\{x_{1}, x_{2}\right\},\left\{v_{1}, v_{2}\right\}$ and $\left\{u_{1}, u_{2}\right\}$ belong to the same relations. Now, consider merging (1)(©)(1) $\left(x_{1}, v_{1}\right)$, (1) $\left.\left.\left(x_{2}, v_{2}\right)\right), u_{1}\right)$ (or $\left.x_{1} v_{1} u_{1}+x_{2} v_{2} u_{1}\right)$ and (1) $\left(x_{1}, u_{2}, v_{1}\right)$ (or $x_{1} u_{2} v_{1}$ ). When we apply $\mathcal{S}$ on these two, we will need to apply rule 5 . Merging rule 5 will then make two recursive calls $\mathcal{S}\left(\mathbb{( 0 )}(1)\left(x_{1}, v_{1}\right)\right.$, (1) $\left.\left(x_{2}, v_{2}\right)\right)$, (1) $\left.\left(x_{1}, v_{1}\right)\right)$ and $\mathcal{S}\left(u_{1}, u_{2}\right)$. Assume that the co-trees returned, (1)(1) $\left(x_{1}, v_{1}\right)$, (1) $\left.\left(x_{2}, v_{2}\right)\right)$ and (0) $\left(u_{1}, u_{2}\right)$, are correctly formed. Thus, we now need only worry about putting these together. The simple thing that we do is to put a (1) node as root and make the roots of the co-trees returned by the recursive calls its children. But in the case of this example this would not work. This is because in doing so, we would introduce an edge between $v_{2}$ and $u_{2}$ (among other spurious edges that will
also get introduced). Taking a closer look at the correct merge result, $x_{1} v_{1} u_{1}+x_{2} v_{2} u_{1}+x_{1} u_{2} v_{1}$, we should notice that it contains a $P_{4}: v_{2}-u_{1}-x_{1}-u_{2}$ (among others). This kind of a tricky merge situation occurs everytime we produce two separate recursive calls one of which involves variables from the first co-tree not present in the second and the other containing variables from the second co-tree not present in the first (e.g., $v_{2}$ and $u_{2}$, respectively, in the example). In the case of this example, as it turned out we were lucky and the result was not read-once. The question is, is this always the case? The next lemma answers this question in the affirmative.

Lemma 6. Let application of merging rule 5 (Figure 4) produce two recursive calls $\mathcal{S}\left(T_{1}^{A}, T_{1}^{B}\right)$ and $\mathcal{S}\left(T_{2}^{A}, T_{2}^{B}\right)$, where $A$ and $B$ denote the input co-trees, $T_{1}^{A}, T_{2}^{A}$ minimal subsets of subtrees formed from $A$ and $T_{1}^{B}, T_{2}^{B}$ minimal subsets formed from $B$. If $\mathcal{N}\left(T_{1}^{A}\right) \wedge \mathcal{N}\left(T_{2}^{B}\right) \vee \mathcal{N}\left(T_{2}^{A}\right) \wedge \mathcal{N}\left(T_{1}^{B}\right)$ is true then the merge result contains a $P_{4}$.

Proof. Just as in the previous proof of Lemma 5, we will assume that $A$ and $B$ share at least one common variable (otherwise we will not apply merging rule 5 but rule 2 instead). Moreover, if at least one common node is not present in $T_{1}^{A}, T_{2}^{A}$ and $T_{1}^{B}, T_{2}^{B}$ then we still have work arounds that will not allow the $\mathcal{S}$ calls stated in the lemma to occur. We will denote by $x$ the common variable and by $R_{x}$ the relation it belongs to. Let $v^{*}$ and $u^{*}$ denote the new variables from $A$ and $B$, respectively, such that $v^{*}$ is not present in $B$ and $u^{*}$ is not present in $A$. Moreover, in at least one co-tree, either $A$ or $B$, the child sub-tree containing the new variable must contain a common variable in it. Otherwise, we have work arounds that do not let the $\mathcal{S}$ calls stated in the lemma to occur. WLOG, we will assume that the new variable present with the common variable in the same sub-tree is $v^{*}$ (present in co-tree $A$ ) and if the common variable is not $x$ then rename it. Now, we have three cases, they are shown in Figure 9. In Figure 9, $v_{1}$ and $u_{1}$ belong to relations $R_{v}$ and $R_{u}$, respectively, where $R_{u}$ contains $u^{*}$ and $R_{v}$ contains $v^{*}$. The first case contains two subcases, depending on whether the LCA of $x$ and $v^{*}$ is a (0) or (1) node. In either case, we find a $P_{4}$. Also, note that in the first case when $L C A$ is a (0) node, $v^{*}$ could belong to $R_{x}$ so there is no separate $R_{v}$ but the stated $P_{4}$ still exists. The rest of the figure should be self-explanatory.

Time Complexity: Recall that in Section 3.3, we showed that for the simple case when no new dummy (1) nodes are introduced by rule 5 in Figure 4, $\oplus\left(T_{1}, T_{2}\right)$ incurs a time complexity of $O\left(n m k^{2}\right)$, where $n$ and $m$ denote the number of nodes in $T_{1}$ and $T_{2}$, respectively, and $k$ denotes the total number of relations. For the more



Figure 10: Query plan with a bushy join. Finding non-interfering sub-formulas is an issue.
general case when dummy nodes are introduced, we first note that two dummy nodes generated from, say, $T_{1}$ introduced by an invocation of rule 5 never get matched with the same node in $T_{2}$. Also, due to Lemma 5, dummy nodes always get matched against real nodes that were present in the original co-tree. This implies that extra work incurred per (1) node in $T_{1}$ is at most $O\left(m k^{2}\right)$, which implies the total extra work done due to dummy nodes cannot exceed $O\left(n m k^{2}\right)$. So, the total time complexity remains $O\left(n m k^{2}\right)$.

## C. PLANS WITH BUSHY JOINS

Plans with bushy joins can also be handled but they require slightly more work. Let us first consider an example that helps illustrate the issues in this case. Figure 10 shows how a plan with a bushy join between $R_{1} \times R_{2}$ and $R_{3} \times R_{4}$ can cause complications. More specifically, the subtree in the lineage-tree rooted at $\bowtie^{2}$ contains variable $y_{2}$ but not $x_{2}$, the subtree rooted at $\bowtie^{3}$ contains $x_{2}$ but not $y_{2}$ whereas the subtree rooted at $\bowtie^{4}$ contains both variables with an edge between them in the corresponding co-occurrence graph (type 2 interference). Another way of stating the issue is to notice that, assuming we manage to construct co-trees for subtrees rooted at $\bowtie^{1}, \bowtie^{2}, \bowtie^{3}$ and $\bowtie^{4}$, merging any three of these leads to a $P_{4}$ (e.g., merging graphs of $\bowtie^{1}, \bowtie^{2}$ and $\bowtie^{3}$ would give us $P_{4}$ $y_{2}-x_{1}-y_{1}-x_{2}$ ) but the result tuple $r$ in Figure 10 is read-once.

We now describe a technique to handle bushy joins. Let $\mathcal{L}$ denote the lineage-tree of result tuple $\phi$. Let $n_{\bowtie}$ denote an internal node in lineage-tree $\mathcal{L}$ that represents a bushy join (e.g., $\bowtie^{1}$ ). Moreover, let $n_{\bowtie \leftarrow}$ and $n_{\bowtie \rightarrow}$ represent $n_{\bowtie}$ 's left and right child respectively. Let us consider computing $\mathcal{T}\left(n_{\bowtie, 1}, \ldots n_{\bowtie, m}\right)$ where each of $n_{\bowtie, i}$ represents a node in $\mathcal{L}$ representing a tuple produced by a bushy join. Essentially, in $\mathcal{L}$ produced by plans with bushy joins interference can occur but we will try to come up with an ordering over the co-tree merging operations that circumvents this, i.e., does not detect a $P_{4}$ unless the result tuple has one. First consider type 1 interference. Suppose $n_{\bowtie, i \leftarrow}$ contains variable $a \in R$. Repeating arguments presented in Lemma 1 we can show that if $n_{\bowtie, i \rightarrow}$ is such that it contains variables $x, y$ such that $x-y \in G_{\phi}$ but $x \not f y$ in $n_{\bowtie, i \rightarrow}$ then there must exist another node $n_{\bowtie, j}$ such that $a$ is a leaf in $n_{\bowtie, j \leftarrow}$ and $x, y$ form an edge in $n_{\bowtie, j \rightarrow}$. Thus, if we do not separate the tuples represented by $n_{\bowtie, i}$ and $n_{\bowtie, j}$ then the type 1 interference can never cause a problem because the potential $P_{4}$ and the chord are always present together. Again, by repeating arguments presented in Lemma 1, if we avoid type 1 interference then we avoid type 2 interference also. The complete ordering procedure is as follows. Given set of tuples $N=\left\{n_{\bowtie, 1}, \ldots n_{\bowtie, m}\right\}$ to be merged and formed a co-tree for, first pick two relations $R_{\leftarrow}$ and $R_{\rightarrow}$ such that base tuples from $R_{\leftarrow}$ and $R_{\rightarrow}$ form leaves in $N$ 's left children and right children, respectively. Now, let $N_{a, b} \subseteq N$ represent the set of nodes such that $\forall n_{\bowtie, i} \in N_{a, b} a \in R_{\leftarrow}$ is a leaf in $n_{\bowtie, i \leftarrow}$ and $b \in R_{\rightarrow}$ forms a leaf in $n_{\bowtie, i \rightarrow}$. The idea is to take all such pairs $a, b$ and form co-trees for $N_{a, b}$ separately. Subsequently, we merge the various co-trees obtained by merging co-trees for all

| $\mathrm{X}:$ | $\mathbf{A}_{\mathbf{1}}$ | $\mathbf{B}_{\mathbf{1}}$ | $\mathrm{Y}:$ | $\mathbf{A}_{\mathbf{2}}$ |  | $\mathrm{Z}:$ |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: |
|  | $\mathrm{Z}_{1}$ | 3 | 4 | $\mathbf{B}_{\mathbf{2}}$ |  |  |
|  | $y_{1}$ | 1 | $z_{1}$ | 2 |  |  |
| $x_{2}$ | 1 | 2 | $y_{2}$ | 3 | $z_{2}$ | 4 |
|  |  |  |  |  |  |  |

$q():-X\left(\mathbf{A}_{1}, \mathbf{B}_{1}\right), Y\left(\mathbf{A}_{\mathbf{2}}\right), Z\left(\mathbf{B}_{\mathbf{2}}\right), A_{1}=A_{2} \vee B_{1}=B_{2}$
$r=x_{1} y_{1} z_{2}+x_{1} y_{2} z_{1}+x_{1} y_{2} z_{2}+x_{2} y_{1} z_{1}+x_{2} y_{1} z_{2}$
$+x_{2} y_{2} z_{1}$


Figure 11: A disjunctive query producing a result tuple that is not normal.
$N_{a, b}$ with the same $a$ first. To form the co-tree for some $N_{a, b}$ we recurse, i.e., we pick another pair of relations $R_{\leftarrow}^{\prime}$ and $R_{\rightarrow}^{\prime}$ different from $R_{\leftarrow}$ and $R_{\rightarrow}$ and repeat the whole procedure. This merge operation ordering approach is a simple way to extend the approach for plans with left/right deep joins presented in the main body of the paper to the case of plans involving bushy joins. In the case of bushy joins, we no longer have join paths easily determined from the lineage-tree but need to construct them ourselves which we do through recursion.

Going back to the example presented Figure 10, one viable ordering is determined by setting $R_{\leftarrow}=R_{1}$ and $R_{\rightarrow}=R_{3}$. Now we have, $N_{a_{1}, x_{1}}=\left\{\bowtie^{1}, \bowtie^{2}\right\}$ whose co-tree is easily formed, let's call this $T_{a_{1}, x_{1}}$, and $N_{a_{2}, x_{2}}=\left\{\bowtie^{3}, \bowtie^{4}\right\}$ whose co-tree we will denote by $T_{a_{2}, x_{2}}$. Finally, we merge $T_{a_{1}, x_{1}}$ and $T_{a_{2}, x_{2}}$. Another ordering that would also work (besides others) is to set $R_{\leftarrow}=R_{1}$ and $R_{\rightarrow}=R_{4}$, in which case, $N_{a_{1}, y_{1}}=\left\{\bowtie^{1}\right\}, N_{a_{1}, y_{2}}=\left\{\bowtie^{2}\right\}$, $N_{a_{2}, y_{1}}=\left\{\bowtie^{3}\right\}, N_{a_{2}, y_{2}}=\left\{\bowtie^{4}\right\}$. Let $T_{a, y}$ denote the co-tree obtained from $N_{a, y}$. We then perform the remaining merges in the following order: $\left(T_{a_{1}, y_{1}} \oplus T_{a_{1}, y_{2}}\right) \oplus\left(T_{a_{2}, y_{1}} \oplus T_{a_{2}, y_{2}}\right)$. Note that, had we chosen to perform $\left(T_{a_{1}, y_{1}} \oplus T_{a_{2}, y_{1}}\right) \oplus\left(T_{a_{1}, y_{2}} \oplus T_{a_{2}, y_{2}}\right)$ instead, then that would have also worked.

## D. OPEN QUESTIONS

Having considered the case of conjunctive queries in depth, the next obvious question is whether our techniques extend to larger classes of queries. Consider the case of queries with disjunctions, some disjunctions can be allowed without breaking any of our results. But if the query contains a disjunctive join predicate that involves attributes from more than two relations and cannot be broken down into smaller conjunctive predicates then our techniques do not apply. Figure 11 shows one such case with a disjunctive join predicate that leads to a result tuple which is $P_{4}$-free but not normal (e.g., $x_{1}, y_{1}, z_{1}$ is a clique in $G_{\phi_{r}}$ but no single clause contains all three variables). Of course, the case for queries with self-joins also remains open.


[^0]:    Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Articles from this volume were presented at The 36th International Conference on Very Large Data Bases, September 13-17, 2010, Singapore.
    Proceedings of the VLDB Endowment, Vol. 3, No. 1
    Copyright 2010 VLDB Endowment 2150-8097/10/09... \$ 10.00.

[^1]:    ${ }^{*} \phi_{2}$ also has a chord for $\phi_{1}$ 's other $P_{4}, y_{2}-x_{1}-y_{1}-x_{2}$.

[^2]:    ${ }^{\dagger}$ Dan Olteanu, personal communication.

