

Hypertree Decompositions and Tractable Queries

Georg Gottlob, Nicola Leone, Francesco Scarcello

Institut für Informationssysteme
Technische Universität Wien
Paniglgasse 16, 1040 Wien

E-mail: {gottlob,leone,scarcell}@dbai.tuwien.ac.at

November 1998

DBAI-TR 98/21

HYPERTREE DECOMPOSITIONS AND TRACTABLE QUERIES

Georg Gottlob Nicola Leone Francesco Scarcello *

Institut für Informationssysteme,
Technische Universität Wien
A-1040 Wien, Paniglgasse 16, Austria
{gottlob,leone,scarcell}@dbai.tuwien.ac.at

DBAI-TR 98/21, November 1998

Abstract

Several important decision problems on conjunctive queries (CQs) are NP-complete in general but become tractable, and actually highly parallelizable, if restricted to acyclic or nearly acyclic queries. Examples are the evaluation of Boolean CQs and query containment. These problems were shown tractable for conjunctive queries of bounded treewidth [7], and of bounded degree of cyclicity [18, 17]. The so far most general concept of nearly acyclic queries was the notion of queries of bounded query-width introduced by Chekuri and Rajaraman [7]. While CQs of bounded query width are tractable, it remained unclear whether such queries are efficiently recognizable. Chekuri and Rajaraman [7] stated as an open problem whether for each constant k it can be determined in polynomial time if a query has query width $\leq k$. We give a negative answer by proving this problem NP-complete (specifically, for $k = 4$). In order to circumvent this difficulty, we introduce the new concept of hypertree decomposition of a query and the corresponding notion of hypertree width. We prove: (a) for each k , the class of queries with query width bounded by k is properly contained in the class of queries whose hypertree width is bounded by k ; (b) unlike query width, constant hypertree-width is efficiently recognizable; (c) Boolean queries of constant hypertree width can be efficiently evaluated.

*Partially supported by the *Istituto per la Sistemistica e l'Informatica* of the Italian National Research Council (*ISI-CNR*), under grant n.224.07.5

Contents

1 Introduction and Overview of Results

1.1 Conjunctive Queries

One of the simplest but also one of the most important classes of database queries is the class of *conjunctive queries* (CQs). In this paper we adopt the logical representation of a relational database [29, 1], where data tuples are identified with logical ground atoms, and conjunctive queries are represented as datalog rules. We will, in the first place, deal with *Boolean* conjunctive queries (BCQs) represented by rules whose heads are variable-free, i.e., propositional (see Example 1.1 below). From our results on Boolean queries, we are able to derive complexity results on important database problems concerning general (not necessarily Boolean) conjunctive queries.

Example 1.1 Consider a relational database with the following relation schemas:

```
enrolled(Pers#, Course#, Reg_Date)
teaches(Pers#, Course#, Assigned)
parent(Pers1, Pers2)
```

The BCQ Q_1 below checks whether some student is enrolled in a course taught by his/her parent.

$$Q_1 : ans \leftarrow \text{enrolled}(S, C, R) \wedge \text{teaches}(P, C, A) \wedge \text{parent}(P, S).$$

The following query Q_2 asks: Is there a professor who has a child enrolled in some course?

$$Q_2 : ans \leftarrow \text{teaches}(P, C, A) \wedge \text{enrolled}(S, C', R) \wedge \text{parent}(P, S).$$

Decision problems such as the *evaluation problem* of Boolean CQs, the *query-of-tuple* problem (i.e., checking whether a given tuple belongs to a CQ), and the *containment problem* for CQs have been studied intensively. (For recent references, see [20, 7].) These problems – which are all equivalent via simple logspace transformations (see [13]) – are NP-complete in the general setting but are polynomially solvable for a number of syntactically restricted subclasses.

1.2 Acyclic Queries and Join Trees

Most prominent among the polynomial cases is the class of *acyclic queries* or *tree queries* [32, 3, 12, 33, 8, 9, 10, 22]. A query Q is acyclic if its associated hypergraph $H(Q)$ is acyclic, otherwise Q is cyclic. The vertices of $H(Q)$ are the variables occurring in Q . Denote by $atoms(Q)$ the set of atoms in the body of Q , and by $var(A)$ the variables occurring in any atom $A \in atoms(Q)$. The hyperedges of $H(Q)$ consist of all sets $var(A)$, such that $A \in atoms(Q)$. We refer to the standard notion of cyclicity/acyclicity in hypergraphs used in database theory [21, 29, 1].

A *join tree* $JT(Q)$ for a conjunctive query Q is a tree whose vertices are the atoms in the body of Q such that whenever the same variable X occurs in two atoms A_1 and A_2 , then A_1 and A_2 are connected in $JT(Q)$, and X occurs in each atom on the unique path linking A_1 and A_2 . In other words, the set of nodes in which X occurs induces a (connected) subtree of $JT(Q)$. We will refer to this condition as the *Connectedness Condition* of join trees.

Acyclic queries can be characterized in terms of join trees: A query Q is *acyclic* iff it has a join tree [3, 2].

Example 1.2 While query Q_1 of example 1.1 is cyclic and admits no join tree, query Q_2 is acyclic. A join tree for Q_2 is shown in Figure 1.

Acyclic conjunctive queries have highly desirable computational properties:




Figure 1: A join tree of Q_2

(1) The problem BCQ of evaluating a Boolean conjunctive query can be efficiently solved if the input query is acyclic. Yannakakis provided a (sequential) polynomial time algorithm solving BCQ on acyclic conjunctive queries¹ [32]. The authors of the present paper have recently shown that BCQ is highly parallelizable on acyclic queries, as it is complete for the low complexity class LOGCFL [13]. (2) Acyclicity is efficiently recognizable, and a join tree of an acyclic query is efficiently computable. A linear-time algorithm for computing a join tree is shown in [28]; an L^{SL} method has been provided in [13]. (3) The result of a (non-Boolean) acyclic conjunctive query Q can be *computed* in time polynomial in the combined size of the input instance and of the output relation [32].

Intuitively, the efficient behaviour of Boolean acyclic queries is due to the fact that they can be evaluated by processing the join tree bottom-up by performing upward semijoins, thus keeping small the size of the intermediate relations (that could become exponential if regular join were performed). This method is the Boolean version of Yannakakis evaluation algorithm for general conjunctive queries [32].

Acyclicity is a key-property responsible for the polynomial solvability of problems that are in general NP-hard such as BCQ [6] and other equivalent problems such as Conjunctive Query Containment [23, 7], Clause Subsumption, and Constraint Satisfaction [20, 13]. (For a survey and detailed treatment see [13].)

1.3 Queries of Bounded Width

The tremendous speed-up obtainable in the evaluation of acyclic queries stimulated several research efforts towards the identification of wider classes of queries having the same desirable properties as acyclic queries. These studies identified a number of relevant classes of cyclic queries which are close to acyclic queries, because they can be decomposed via low width decompositions to acyclic queries. The main classes of polynomially solvable bounded-width queries considered in database theory and in artificial intelligence are:

- **The queries of bounded treewidth** [7] (see also [20, 13]). These are queries, whose variable-atom incidence graph has treewidth bounded by a constant.² The treewidth of a graph is a well-known measure of its tree-likeness introduced by Robertson and Seymour in their work on graph minors [24]. This notion plays a central role in algorithmic graph theory as well as in many subdisciplines of Computer Science. We omit a formal definition. It is well-known that checking that a graph has treewidth $\leq k$ for a fixed constant k , and in the positive case, computing a k -width tree decomposition is feasible in linear time [4].
- **Queries of bounded degree of cyclicity** [18, 17]. This is an interesting class of queries which also encompasses the class of acyclic queries. For space reasons, we omit a formal definition. For each constant k , checking whether a query has degree of cyclicity $\leq k$ is feasible in polynomial time [18, 17].

¹Note that, since both the database \mathbf{DB} and the query Q are part of an input-instance of BCQ, what we are considering is the *combined complexity* of the query [31].

²As pointed out in [20], the notion of treewidth of a query can be equivalently based on the Gaifman graph of a query, i.e., the graph linking two variables by an edge if they occur together in a query-atom.




Figure 2: A 2-width query decomposition of query Q_1

- **Queries of bounded query-width.** [7]. This notion of bounded query-width is based on the concept of *query decomposition* [7]. Roughly, a query decomposition of a query Q consists of a tree each vertex of which is labelled by a set of atoms and/or variables. Each variable and atom induces a connected subtree (*connectedness condition*). Each atom occurs in at least one label. The width of a query decomposition is the maximum of the cardinalities of its vertices. The *query width* $qw(Q)$ of Q is the minimum width over all its query decompositions. A formal definition is given in Section 3.1; Figure 2 shows a 2-width query-decomposition for the cyclic query Q_1 of Example 1.1. This class is the widest of the three classes: Each query of bounded treewidth or of bounded degree of cyclicity k has also bounded query width k , but for some queries the converse does not hold [7, 13]. There are even classes of queries with bounded query width but unbounded treewidth. Note, however, that no polynomial algorithm for checking whether a query has width $\leq k$ was known.

All these concepts are true generalizations of the basic concept of acyclicity. For example, a query is acyclic iff it has query width 1.

Intuitively, a vertex of a k -width query decomposition stands for the natural join of (the relations of) its elements – the size of this join is $O(n^k)$, where n is the size of the input database. Once these joins have been done, the query decomposition can be treated exactly like a join tree of an acyclic query, and permits to evaluate the query in time polynomial in n^k [7].

The problem BCQ (evaluation of Boolean conjunctive queries) and the bounded query-width versions of all mentioned equivalent problems, e.g. query-containment $Q_1 \subseteq Q_2$, where the query width of Q_2 is bounded, can be efficiently solved *if a k -width query decomposition of the query is given as (additional) input*. Chekuri and Rajaraman provided a polynomial time algorithm for this problem [7]; Gottlob et al. [13] later pinpointed the precise complexity of the problem by proving it LOGCFL-complete.

1.4 A negative Result

Unfortunately, unlike for acyclicity or for bounded treewidth, or for bounded degree of cyclicity, no efficient method for checking bounded query-width is known, and a k -width query decomposition, which is required for the efficient evaluation of a bounded-width query, is not known to be polynomial time computable.

Chekuri and Rajaraman [7] state this as an open problem. This problem is the first question we address in the present paper.

The fact that treewidth k can be checked in linear time suggests that an analogous algorithm may work for query width, too. Chekuri and Rajaraman [7] express some optimism by writing “*it would be useful to have an efficient algorithm that produces query decompositions of small width, analogous to the algorithm of Bodlaender [4] for decompositions of small treewidth.*” Kolaitis and Vardi [20] write somewhat more pessimistically: “*there is an important advantage of the concept of bounded treewidth over the concept of bounded querywidth. Specifically, as seen above, the classes of structures of bounded treewidth are polynomially recognizable, whereas it is not known whether the same holds true for the classes of queries of bounded querywidth.*”

Our first main result is bad news:

Determining whether the query width of a conjunctive query is at most 4 is NP-complete.

The NP-completeness proof is rather involved and will be given in Section 3.2. We give a rough intuition in this section. The proof led us to a better intuition about (i) why the problem is NP-complete, and (ii) how this could be redressed by suitably modifying the notion of query width. Very roughly, the source of NP-hardness can be pinpointed as follows. In order to obtain a query-decomposition of width bounded by k , it can be seen that it is implicitly required to proceed as follows. At any step, the decomposition is guided by a set C of variables that still needs to be processed. Initially, i.e., at the root of the decomposition, C consists of *all* variables that occur in the query. We then choose as root of the decomposition tree a hypernode R of $\leq k$ query atoms. By fixing this hypernode, we eliminate a set of variables, namely those which occur in the atoms of R . The remaining set of variables disintegrates into connected components. We now expand the decomposition tree by attaching children, and thus, in the long run, subtrees to R . It can be seen that each subtree rooted in R must correspond to one or more of the connected components and each component occurs in exactly one subtree (otherwise the connectedness condition would be violated). In particular, since each atom should be eventually covered, each remaining atom must be covered by some subtree, i.e., must occur in some subtree. This process goes on until all variables are eliminated and until all query atoms are eventually covered. The definition of query width requires that this covering be *exact*, i.e., that each atom containing a variable of a certain component C occurs only in the subtree corresponding to C . (Again, the requirement of exact covering is due to the connectedness condition.) The NP hardness is due to this requirement of *exact* covering. In our NP-completeness proof, we thus tried to reduce the problem of EXACT COVERING BY 3SETS to the query width problem. And this attempt was successful.

1.5 Hypertree Decompositions: Positive Results

To circumvent the high complexity of query-decompositions, we introduce a new concept of decomposition, which we call *hypertree decomposition*. The definition of hypertree decomposition (see Section 4) uses a more liberal notion of covering. When choosing a set SC of components to be processed, the variables of SC are no longer required to exactly *coincide* with the variables occurring in the labels of the decomposition-subtree corresponding to SC , but it is sufficient that the former be a subset of the latter. Based on this more liberal notion of decomposition, we define the corresponding notion of *hypertree width* in analogy to the concept of query width.

We denote the query width of a query by $qw(Q)$ and its hypertree width by $hw(Q)$. We are able to prove the following results:

1. For each conjunctive query Q it holds that $hw(Q) \leq qw(Q)$.
2. There exist queries Q such that $hw(Q) < qw(Q)$.
3. For each fixed constant k , the problems of determining whether $hw(Q) \leq k$ and of computing (in the positive case) a hypertree decomposition of width $\leq k$ are feasible in polynomial time.
4. For fixed k , evaluating a Boolean conjunctive query Q with $hw(Q) \leq k$ is feasible in polynomial time.
5. The result of a (non-Boolean) conjunctive query Q of bounded hypertree-width can be computed in time polynomial in the combined size of the input instance and of the output relation.
6. Tasks 3 and 4 are not only polynomial, but are highly parallelizable. In particular, for fixed k , checking whether $hw(Q) \leq k$ is in the parallel complexity class LOGCFL; computing a hypertree decomposition of width k (if any) is in functional LOGCFL, i.e., is feasible by a logspace transducer that uses an oracle in LOGCFL; evaluating Q where $hw(Q) \leq k$ on a database is complete for LOGCFL under logspace reductions.

Similar results hold for the equivalent problem of conjunctive query containment $Q_1 \subseteq Q_2$, where $hw(Q_2) \leq k$, and for all other of the aforementioned equivalent problems.

Let us comment on these results. By statements 1 and 2, the concept of hypertree width is a proper generalization of the notion of query width. By statement 3, constant hypertree-width is efficiently checkable, and by statement 4, queries of constant hypertree width can be efficiently evaluated. In summary, this is truly good news. It means that the notion of constant hypertree width not only shares the desirable properties of constant query-width, it also does *not* share the bad properties of the latter, and, in addition, is a more general concept.

It thus turns out that the high complexity of determining constant query width is not, as one would usually expect, the price for the generality of the concept. Rather, it is due to some peculiarity in its definition related to the exact covering paradigm. In the definition of hypertree width we succeeded to eliminate these problems without paying any additional charge, i.e., hypertree width comes as a freebie!

Statement 6 asserts that the main algorithmic tasks related to constant hypertree-width are in the very low complexity class LOGCFL. This class consists of all decision problems that are logspace reducible to a context-free language. An obvious example of a problem complete for LOGCFL is Greibach's hardest context-free language [16]. There is a number of very interesting natural problems known to be LOGCFL-complete (see, e.g. [27, 26, 13]). The relationship between LOGCFL and other well-known complexity classes is summarized in the following chain of inclusions:

$$AC^0 \subseteq NC^1 \subseteq L \subseteq SL \subseteq NL \subseteq LOGCFL \subseteq AC^1 \subseteq NC^2 \subseteq P \subseteq NP$$

Here L denotes logspace, AC^i and NC^i are *logspace-uniform* classes based on the corresponding types of Boolean circuits, SL denotes symmetric logspace, NL denotes nondeterministic logspace, P is polynomial time, and NP is nondeterministic polynomial time. For the definitions of all these classes, and for references concerning their mutual relationships, see [19].

Since $LOGCFL \subseteq AC^1 \subseteq NC^2$, the problems in LOGCFL are all highly parallelizable. In fact, they are solvable in logarithmic time by a CRCW PRAM with a polynomial number of processors, or in \log^2 -time by an EREW PRAM with a polynomial number of processors.

1.6 Structure of the Paper

Basic notions of database and complexity theory are given in Section 2. Section 3 deals with query decompositions and includes the NP-completeness proof for the problem of deciding bounded query-width. The new

notions of hypertree decomposition and hypertree width are formally defined in Section 4, where also some examples are given, and it is shown that bounded hypertree-width queries are efficiently evaluable. In Section 5, the alternating algorithm `k-decomp` that checks whether a query has hypertree width $\leq k$ is presented. This algorithm is shown to run on a logspace ATM having polynomially sized accepting computation trees, thus the problem is actually in LOGCFL. Finally, a short sketch of a deterministic polynomial algorithm (in form of a datalog program) for checking whether a query has hypertree width $\leq k$ is given. In Section 6 the notion of hypertree decomposition is shown to be the most general among the most important related notions, e.g., the notion of query decomposition.

2 Preliminaries

2.1 Databases and Queries

For a background on databases, conjunctive queries, etc., see [29, 1, 21]. We define only the most relevant concepts here.

A relation schema R consists of a name (name of the relation) r and a finite ordered list of attributes. To each attribute A of the schema, a countable domain $Dom(A)$ of atomic values is associated. A *relation instance* (or simply, a *relation*) over schema $R = (A_1, \dots, A_k)$ is a finite subset of the cartesian product $Dom(A_1) \times \dots \times Dom(A_k)$. The elements of relations are called *tuples*. A database schema DS consists of a finite set of relation schemas. A *database instance*, or simply *database*, \mathbf{DB} over database schema $DS = \{R_1, \dots, R_m\}$ consists of relation instances r_1, \dots, r_m for the schemas R_1, \dots, R_m , respectively, and a finite universe $U \subseteq \bigcup_{R_i(A_1^i, \dots, A_{k_i}^i) \in DS} (Dom(A_1^i) \cup \dots \cup Dom(A_{k_i}^i))$ such that all data values occurring in \mathbf{DB} are from U .

In this paper we will adopt the standard convention [1, 29] of identifying a relational database instance with a logical theory consisting of ground facts. Thus, a tuple $\langle a_1, \dots, a_k \rangle$, belonging to relation r , will be identified with the ground atom $r(a_1, \dots, a_k)$. The fact that a tuple $\langle a_1, \dots, a_k \rangle$ belongs to relation r of a database instance \mathbf{DB} is thus simply denoted by $r(a_1, \dots, a_k) \in \mathbf{DB}$.

A (rule based) *conjunctive query* Q on a database schema $DS = \{R_1, \dots, R_m\}$ consists of a rule of the form

$$Q : \text{ans}(\mathbf{u}) \leftarrow r_1(\mathbf{u}_1) \wedge \dots \wedge r_n(\mathbf{u}_n),$$

where $n \geq 0$, r_1, \dots, r_n are relation names (not necessarily distinct) of DS ; ans is a relation name not in DS ; and $\mathbf{u}, \mathbf{u}_1, \dots, \mathbf{u}_n$ are lists of terms (i.e., variables or constants) of appropriate length. The set of variables occurring in Q is denoted by $\text{var}(Q)$. The set of atoms contained in the body of Q is referred to as $\text{atoms}(Q)$. Similarly, for any atom $A \in \text{atoms}(Q)$, $\text{var}(A)$ denotes the set of variables occurring in A ; and for a set of atoms $R \subseteq \text{atoms}(Q)$, define $\text{var}(R) = \bigcup_{A \in R} \text{var}(A)$.

The *answer* of Q on a database instance \mathbf{DB} with associated universe U , consists of a relation ans whose arity is equal to the length of \mathbf{u} , defined as follows. ans contains all tuples $\text{ans}(\mathbf{u})\vartheta$ such that $\vartheta : \text{var}(Q) \rightarrow U$ is a substitution replacing each variable in $\text{var}(Q)$ by a value of U and such that for $1 \leq i \leq n$, $r_i(\mathbf{u}_i)\vartheta \in \mathbf{DB}$. (For an atom A , $A\vartheta$ denotes the atom obtained from A by uniformly substituting $\vartheta(X)$ for each variable X occurring in A .)

The conjunctive query Q is a *Boolean conjunctive query (BCQ)* if its head atom $\text{ans}(\mathbf{u})$ does not contain variables and is thus a purely propositional atom. Q evaluates to *true* if there exists a substitution ϑ such that for $1 \leq i \leq n$, $r_i(\mathbf{u}_i)\vartheta \in \mathbf{DB}$; otherwise the query evaluates to *false*.

The head literal in Boolean conjunctive queries is actually inessential, therefore we may omit it when specifying a Boolean conjunctive query.

Note that conjunctive queries as defined here correspond to conjunctive queries in the more classical setting of relational calculus, as well as to SELECT-PROJECT-JOIN queries in the classical setting of relational

algebra, or to simple SQL queries of the type

$$\text{SELECT } R_{i_1}.A_{j_1}, \dots, R_{i_k}.A_{j_k} \text{ FROM } R_1, \dots, R_n \text{ WHERE } \textit{cond},$$

such that *cond* is a conjunction of conditions of the form $R_i.A = R_j.B$ or $R_i.A = c$, where c is a constant.

A query Q is *acyclic* [2, 3] if its associated hypergraph $H(Q)$ is acyclic, otherwise Q is cyclic. The vertices of $H(Q)$ are the variables occurring in Q . Denote by $\textit{atoms}(Q)$ the set of atoms in the body of Q , and by $\textit{var}(A)$ the variables occurring in any atom $A \in \textit{atoms}(Q)$. The hyperedges of $H(Q)$ consist of all sets $\textit{var}(A)$, such that $A \in \textit{atoms}(Q)$. We refer to the standard notion of cyclicity/acyclicity in hypergraphs used in database theory [21, 29, 1].

A *join tree* $\mathcal{JT}(Q)$ for a conjunctive query Q is a tree whose vertices are the atoms in the body of Q such that whenever the same variable X occurs in two atoms A_1 and A_2 , then A_1 and A_2 are connected in $\mathcal{JT}(Q)$, and X occurs in each atom on the unique path linking A_1 and A_2 . In other words, the set of nodes in which X occurs induces a (connected) subtree of $\mathcal{JT}(Q)$ (*connectedness condition*).

Acyclic queries can be characterized in terms of join trees: A query Q is *acyclic* iff it has a join tree [3, 2].

Example 2.1 While query Q_1 of example 1.1 is cyclic and admits no join tree, query Q_2 is acyclic. A join tree for Q_2 is shown in Figure 1.

Consider the following query Q_3 :

$$\textit{ans} \leftarrow r(Y, Z) \wedge g(X, Y) \wedge s(Y, Z, U) \wedge s(Z, U, W) \wedge t(Y, Z) \wedge t(Z, U)$$

A join tree for Q_3 is shown in Figure 3.




Figure 3: A join tree of Q_3

Acyclic conjunctive queries have highly desirable computational properties:

1. The problem BCQ of evaluating a Boolean conjunctive query can be efficiently solved if the input query is acyclic. Yannakakis provided a (sequential) polynomial time algorithm solving BCQ on acyclic conjunctive queries³ [32]. The authors of the present paper have recently shown that BCQ is highly parallelizable on acyclic queries, as it is complete for the low complexity class LOGCFL [13].
2. Acyclicity is efficiently recognizable, and a join tree of an acyclic query is efficiently computable. A linear-time algorithm for computing a join tree is shown in [28]; an L^{SL} method has been provided in [13].
3. The result of a (non-Boolean) acyclic conjunctive query Q can be *computed* in time polynomial in the combined size of the input instance and of the output relation [32].

³Note that, since both the database **DB** and the query Q are part of an input-instance of BCQ, what we are considering is the *combined complexity* of the query [31].

Intuitively, the efficient behaviour of Boolean acyclic queries is due to the fact that they can be evaluated by processing the join tree bottom-up by performing upward semijoins, thus keeping small the size of the intermediate relations (that could become exponential if regular join were performed). This method is the Boolean version of Yannakakis evaluation algorithm for general conjunctive queries [32].

Acyclicity is a key-property responsible for the polynomial solvability of problems that are in general NP-hard such as BCQ [6] and other equivalent problems such as Conjunctive Query Containment [23, 7], Clause Subsumption, and Constraint Satisfaction [20, 13]. (For a survey and detailed treatment see [13].)

2.2 The class LOGCFL

LOGCFL consists of all decision problems that are logspace reducible to a context-free language. An obvious example of a problem complete for LOGCFL is Greibach’s hardest context-free language [16]. There are a number of very interesting natural problems known to be LOGCFL-complete (see, e.g. [13, 27, 26]). The relationship between LOGCFL and other well-known complexity classes is summarized in the following chain of inclusions:

$$AC^0 \subseteq NC^1 \subseteq L \subseteq SL \subseteq NL \subseteq LOGCFL \subseteq AC^1 \subseteq NC^2 \subseteq P \subseteq NP$$

Here L denotes logspace, AC^i and NC^i are *logspace-uniform* classes based on the corresponding types of Boolean circuits, SL denotes symmetric logspace, NL denotes nondeterministic logspace, P is polynomial time, and NP is nondeterministic polynomial time. For the definitions of all these classes, and for references concerning their mutual relationships, see [19].

Since – as mentioned in the introduction – $LOGCFL \subseteq AC^1 \subseteq NC^2$, the problems in LOGCFL are all highly parallelizable. In fact, they are solvable in logarithmic time by a CRCW PRAM with a polynomial number of processors, or in \log^2 -time by an EREW PRAM with a polynomial number of processors.

In this paper, we will use an important characterization of LOGCFL by Alternating Turing Machines. We assume that the reader is familiar with the *alternating Turing machine (ATM)* computational model introduced by Chandra, Kozen, and Stockmeyer [5]. Here we assume w.l.o.g. that the states of an ATM are partitioned into existential and universal states.

As in [25], we define a *computation tree* of an ATM M on a input string w as a tree whose nodes are labeled with configurations of M on w , such that the descendants of any non-leaf labeled by a universal (existential) configuration include all (resp. one) of the successors of that configuration. A computation tree is *accepting* if the root is labeled with the initial configuration, and all the leaves are accepting configurations.

Thus, an accepting tree yields a certificate that the input is accepted. A complexity measure considered by Ruzzo [25] for the alternating Turing machine is the *tree-size*, i.e. the minimal size of an accepting computation tree.

Definition 2.2 ([25]) *A decision problem \mathcal{P} is solved by an alternating Turing machine M within simultaneous tree-size and space bounds $Z(n)$ and $S(n)$ if, for every “yes” instance w of \mathcal{P} , there is at least one accepting computation tree for M on w of size (number of nodes) $\leq Z(n)$, each node of which represents a configuration using space $\leq S(n)$, where n is the size of w . (Further, for any “no” instance w of \mathcal{P} there is no accepting computation tree for M .)*

Ruzzo [25] proved the following important characterization of LOGCFL :

Proposition 2.3 (Ruzzo [25]) *LOGCFL coincides with the class of all decision problems recognized by ATMs operating simultaneously in tree-size $O(n^{O(1)})$ and space $O(\log n)$.*

3 Query Decompositions




Figure 4: A 2-width query decomposition of query Q_4

3.1 Bounded Query Width and Bounded Query Decompositions

The following definition of query decomposition is a slight modification of the original definition given by Chekuri and Rajaraman [7]. Our definition is a bit more liberal because, for any conjunctive query Q , we do not take care of the atom $head(Q)$, as well as of the constants possibly occurring in Q . However, in this paper, we will only deal with Boolean conjunctive queries without constants, for which the two notions coincide.

Definition 3.1 A *query decomposition* of a conjunctive query Q is a pair $\langle T, \lambda \rangle$, where $T = (N, E)$ is a tree, and λ is a labeling function which associates to each vertex $p \in N$ a set $\lambda(p) \subseteq (atoms(Q) \cup var(Q))$, such that the following conditions are satisfied:

1. for each atom A of Q , there exists $p \in N$ such that $A \in \lambda(p)$;
2. for each atom A of Q , the set $\{p \in N \mid A \in \lambda(p)\}$ induces a (connected) subtree of T ;
3. for each variable $Y \in var(Q)$, the set

$$\{p \in N \mid Y \in \lambda(p)\} \cup \{p \in N \mid Y \text{ occurs in some atom } A \in \lambda(p)\}$$

induces a (connected) subtree of T .

The *width* of the query decomposition $\langle T, \lambda \rangle$ is $\max_{p \in N} |\lambda(p)|$. The *query width* $qw(Q)$ of Q is the minimum width over all its query decompositions. A query decomposition for Q is *pure* if, for each vertex $p \in N$, $\lambda(p) \subseteq atoms(Q)$.

Note that Condition 3 above is the analogue of the connectedness condition of join trees and thus we will refer to it as the *Connectedness Condition*, as well.

Example 3.2 Figure 2 shows a 2-width query decomposition for the cyclic query of Example 1.1.

Consider the following query Q_4 :

$$ans \leftarrow s(Y, Z, U) \wedge g(X, Y) \wedge t(Z, X) \wedge s(Z, W, X) \wedge t(Y, Z)$$

Q_4 is a cyclic query, and its query width equals 2. A 2-width decomposition of Q_4 is shown in Figure 4. Note that this query decomposition is pure.

The next proposition, which is proved elsewhere [13], shows that we can focus our attention on pure query decompositions.

Proposition 3.3 ([13]) *Let Q be a conjunctive query and $\langle T, \lambda \rangle$ a c -width query decomposition of Q . Then,*

1. there exists a pure c -width query decomposition $\langle T, \lambda' \rangle$ of Q ;
2. $\langle T, \lambda' \rangle$ is logspace computable from $\langle T, \lambda \rangle$.

Thus, by Proposition 3.3, for any conjunctive query Q , $qw(Q) \leq k$ if and only if Q has a pure c -width decomposition, for some $c \leq k$.

k -bounded-width queries are queries whose query width is bounded by a fixed constant $k > 0$. The notion of bounded query-width generalizes the notion of acyclicity [7]. Indeed, acyclic queries are exactly the conjunctive queries of query width 1, because any join tree is a query decomposition of width 1.

Bounded-width queries share an important computational property with acyclic queries: BCQ can be efficiently solved on queries of k -bounded query-width, if a k -width query decomposition of the query is given as (additional) input. Chekuri and Rajaraman provided a polynomial time algorithm for this problem [7]; while Gottlob et al. pinpointed that the precise complexity of the problem is LOGCFL.

Unfortunately, different from acyclicity, no efficient method for checking bounded query-width is known. In fact, we next prove that deciding whether a conjunctive query has a bounded-width query decomposition is NP-complete.

3.2 Recognizing bounded query-width is NP-complete

A k -element-vertex of a query decomposition (T, λ) is a vertex v of T such that $|\lambda(v)| = k$.

Lemma 3.4 *Let Q be a query having variable set $var(Q) = \Gamma \cup Rest$, where*

$$\Gamma = \{V_{ij} \mid 1 \leq i < j \leq 8\},$$

and $Rest$ is an arbitrary set of further variables. Assume the set $atoms(Q)$ contains as subset a set $\Pi = \{P_1, \dots, P_8\}$ of 8 atoms, where, for $1 \leq i \leq 8$, $var(P_i) \cap \Gamma = \{V_{1i}, V_{2i}, \dots, V_{i-1i}, V_{i+1i}, \dots, V_{i8}\}$, i.e.,

$$var(P_i) \cap \Gamma = \bigcup_{k < i} \{V_{ki}\} \cup \bigcup_{i < k} \{V_{ik}\}.$$

and $\forall A \in atoms(Q) - \Pi : var(A) \cap \Gamma = \emptyset$.

If Q admits a pure query decomposition (T, λ) of width 4, then there exist two adjacent 4-element-vertices p_1 and p_2 of T such that $\lambda(p_1) \cup \lambda(p_2) = \Pi$.

Proof. Assume this doesn't hold. Let v be any vertex from T such that $\Gamma \cap var(v) \neq \emptyset$. Let $R = \Pi - \lambda(v)$. The atoms of R must occur somewhere in the tree T . By the *connectedness condition*, all atoms of R must occur in the labels of some neighbours of v . Moreover, by our assumption, the atoms of R are not contained in the label of a single neighbour of v . Thus, there exist two neighbours v_1, v_2 of v and two different atoms $P_i, P_j \in R$ such that $P_i \in \lambda(v_1) - \lambda(v_2)$ and $P_j \in \lambda(v_2) - \lambda(v_1)$. Assume, w.l.o.g., that $i < j$. Then $V_{ij} \in var(v_1)$, $V_{ij} \in var(v_2)$, but $V_{ij} \notin var(v)$. This, however, violates the *connectedness condition*. Contradiction. ■

Definition 3.5 Let S be a set of n elements. A 3-partition $\{S_a, S_b, S_c\}$ of S consists of three nonempty subsets $S_a, S_b, S_c \subset S$ such that $S_a \cup S_b \cup S_c = S$, and $S_x \cap S_y = \emptyset$ for $x \neq y$ from $\{a, b, c\}$. The sets S_a, S_b, S_c are referred to as *classes*.

A 3-Partitioning-System (short 3PS) Σ on a base set S is a set of 3-partitions of S :

$$\Sigma = \{ \{S_a^1, S_b^1, S_c^1\}, \{S_a^2, S_b^2, S_c^2\}, \dots, \{S_a^m, S_b^m, S_c^m\} \},$$

where $\forall \sigma, \sigma' \in \Sigma : \sigma \neq \sigma' \Rightarrow \sigma \cap \sigma' = \emptyset$ (i.e., no class occurs in two or more elements of Σ).

We define $classes(\Sigma) := \bigcup_{\sigma \in \Sigma} \sigma$. The base set S of Σ is referred-to as $base(\Sigma)$: $base(\Sigma) = \bigcup_{C \in classes(\Sigma)} C$.

A 3PS is *strict* if for all $S', S'', S''' \in \text{classes}(\Sigma)$ either $\{S', S'', S'''\} = \sigma$ for some $\sigma \in \Sigma$ or $S' \cup S'' \cup S''' \subset S$. In other words: the only way to obtain S as a union of three classes is via the specified 3-partitions of Σ ; any other union of three classes results in a proper subset of S .

A 3PS Σ is referred to as an (m, k) -3PS if $|\Sigma| \geq m$ and $\forall C \in \text{classes}(\Sigma) : |C| \geq k$.

Lemma 3.6 *For each $m > 0$ and $k > 0$ a strict (m, k) -3PS can be computed in polynomial time.*

Proof. Fix m and k . We will construct a set S of the desired cardinality $n \leq 27m^3 + 2m + 3k$ such that there exists an (m, k) -3PS for S .

In order to construct S , first start with a set S_0 of $2m$ elements. It is a trivial combinatorial fact that we can choose at least m different 3-partitions of S_0 (recall that $m > 3$, thus the actual number of 3-partitions we can build is much higher). Thus, let Σ_0 be a (not necessarily strict) 3PS on base set S_0 , such that $|\Sigma_0| = m$. We now basically have to achieve two goals: (1) we have to transform Σ_0 to a *strict* 3PS, and, (2) we have to make sure that all classes have at least k elements.

In order to achieve goal (1), let New be a set of $27m^3$ fresh elements. Enumerate (e.g., in a FOR loop) all combinations of three sets $S_x^i, S_y^j, S_z^\ell \in \text{classes}(\Sigma_0)$ and check for each such triplet, whether it violates the strictness-condition. If so, choose a fresh (i.e., so far unused) element a from New , insert it into S_0 and use it as follows to redress the situation: For any set $\sigma \in \Sigma_0$ in which neither of S_x^i, S_y^j, S_z^ℓ occurs, insert a into exactly one of the three classes (which one is irrelevant). This means that the so augmented σ is a partition of the new S_0 . On the other hand, for each $\sigma \in \Sigma_0$ such that $\sigma \cap \{S_x^i, S_y^j, S_z^\ell\} \neq \emptyset$, choose a class $C \in \sigma$ such that $C \notin \{S_x^i, S_y^j, S_z^\ell\}$ and insert a into C . It follows that the triplet S_x^i, S_y^j, S_z^ℓ no longer violates the strictness condition w.r.t. the new set S_0 (containing a), because a is not in the union $S_x^i \cup S_y^j \cup S_z^\ell$, and thus this union is a strict subset of S_0 . Note that this operation does never invalidate the strictness condition of any other triplet. After we have repeated this procedure for each triplet, we end-up with a strict 3PS for the resulting set S_0 . Call this resulting set S^+ and denote the resulting strict 3PS by Σ^+ . Retain that the set S^+ has less than $2m + (3m)^3 = 2m + 27m^3$ elements.

In order to achieve goal (2), we simply add a set New' of $3k$ further fresh elements to S^+ , obtaining $S^* = S^+ \cup New'$. We furthermore partition New' into three sets O_1, O_2 , and O_3 of equal cardinality k and do the following for each 3-partition $\sigma = \{C_a, C_b, C_c\}$ of Σ^+ , where $\langle C_a, C_b, C_c \rangle$ is any arbitrarily chosen order of the elements of σ : perform $C_a := C_a \cup O_1$ and $C_b := C_b \cup O_2$ and $C_c := C_c \cup O_3$.

The resulting 3PS Σ is strict, has $|\Sigma| = m$, $|\text{base}(\Sigma)| \leq 27m^3 + 2m + 3k$, and each class C of $\text{classes}(\Sigma)$ has $|C| \geq k$. We are done. Note that our method for generating a strict 3PS works in polynomial time. ■

Theorem 3.7 *Deciding whether the query width of a conjunctive query is at most 4 is NP-complete.*

Proof. 1. *Membership.* It is easy to see that if there exists a query decomposition of width bounded by 4, then there also exists one of polynomial size (in fact, by a simple restructuring technique we can always remove identically labeled vertices from a decomposition tree, and thus for any conjunctive query Q only $O(|\text{atoms}(Q) \cup \text{var}(Q)|^4)$ need to be considered). Therefore, a query decomposition of width $\leq k$ can be found by a nondeterministic guess followed by a polynomial correctness check. The problem is thus in NP.

2. *Hardness.* We transform the well-known NP-complete problem EXACT COVER BY 3-SETS (XC3C) [11] to the problem of deciding whether, for a conjunctive query Q , $qw(Q) \leq 4$ holds. An instance of EXACT COVER BY 3-SETS consists of a pair $I = (R, \Delta)$ where R is a set of $r = 3s$ elements, and Δ is a collection of m 3-element subsets of R . The question is whether we can select s subsets out of Δ such that they form a partition of R .

Consider an instance $I = (R, \Delta)$ of XC3S. Let $\Delta = \{D_i | 1 \leq i \leq m\}$ and let $D_i = \{X_a^i, X_b^i, X_c^i\}$ for $1 \leq i \leq m$ (note that for $i \neq j$, some X_a^i and X_b^j may coincide).

Generate a strict $(m + 1, 2)$ 3PS $\Sigma = \{\sigma_0, \sigma_1, \dots, \sigma_m\}$ on some base set $S = \text{base}(\Sigma)$. By Lemma 3.6, this can be done in polynomial time. Let $\sigma_i = \{S_a^i, S_b^i, S_c^i\}$ for $0 \leq i \leq m$.

Identify each element of S with a separate variable and establish a fixed precedence order \prec among the elements (variables) of S . If S' is a subset of S , and $S' = \{Z_1, \dots, Z_l\}$, where $Z_1 \prec Z_2 \dots \prec Z_l$, then we will abbreviate the list of variables Z_1, \dots, Z_l by S' in query atoms. For example, instead of writing $p(a, Z_1, \dots, Z_l, b)$, we write $p(a, S', b)$.

In order to transform the given instance $I = (R, \Delta)$ of XC3S to a conjunctive query Q , let us first define the following sets of variables Γ^ℓ and Π_i^ℓ , which are all taken to be disjoint from the variables in S .

For $0 \leq \ell \leq s$, let

$$\Gamma^\ell = \{V_{ij}^\ell \mid 1 \leq i < j \leq 8\},$$

and for $(1 \leq i \leq 8)$, let

$$\Pi_i^\ell = \{V_{1i}^\ell, V_{2i}^\ell, \dots, V_{i-1i}^\ell, V_{i+1i}^\ell, \dots, V_{8i}^\ell\}.$$

Let S'_a and S''_a be two nonempty sets which partition S_a^0 . (Such a partition exists because S_a^0 contains at least two elements.)

Define, for $0 \leq \ell \leq s$ the following sets of query atoms:

$$BLOCKA^\ell = \{q(\Pi_1^\ell, S'_a, Z_\ell), p_a(\Pi_2^\ell, S''_a), p_b(\Pi_3^\ell, S_b^0), p_c(\Pi_4^\ell, S_c^0)\}$$

$$BLOCKB^\ell = \{q(\Pi_5^\ell, S'_a, Y_\ell), p_a(\Pi_6^\ell, S''_a), p_b(\Pi_7^\ell, S_b^0), p_c(\Pi_8^\ell, S_c^0)\},$$

where the Y_ℓ and Z_ℓ variables are distinct fresh variables not occurring in any previously defined set. We further define

$$BLOCKSA = \bigcup_{0 \leq \ell \leq s} BLOCKA^\ell, \quad BLOCKSB = \bigcup_{0 \leq \ell \leq s} BLOCKB^\ell,$$

$$\text{and } BLOCKS = BLOCKSA \cup BLOCKSB.$$

Define, for $1 \leq \ell \leq s$:

$$LINK_\ell = \{link(Y_{\ell-1}, Z_\ell)\} \quad \text{and} \quad LINKS = \bigcup_{1 \leq \ell \leq s} LINK_\ell.$$

Finally, define for each set $D_i = \{X_a^i, X_b^i, X_c^i\}$ of Δ , $1 \leq i \leq m$, the set of atoms:

$$\Omega[D_i] = \{s(X_a^i, S_a^i), s(X_b^i, S_b^i), s(X_c^i, S_c^i)\}.$$

Let $\Omega = \bigcup_{1 \leq i \leq m} \Omega[D_i]$, and denote by $\Omega(D_i)$ the set of all atoms of Ω in which some variable of D_i occurs, i.e.,

$$\Omega(D_i) = \{s(X, \alpha) \in \Omega \mid X \in D_i\}.$$

Let Q be the query whose atom-set is $BLOCKS \cup LINKS \cup \Omega$.

We claim that Q has query width 4 iff $I = (R, \Delta)$ is a positive instance of EXACT COVER BY 3SETS.

Let us first prove the *if* part. Assume that there exist s 3-sets $D^1, \dots, D^s \in \Delta$ which exactly cover R , i.e., which form a partition of R . We describe a query-decomposition (T, λ) of Q .

The root v_{a0} of T is labeled by the set of atoms $BLOCKA^0$. The root has as unique child a vertex v_{b0} labeled by $BLOCKB^0$.

The decomposition tree is continued as follows. For each $1 \leq \ell \leq s$, do the following.

- Create a vertex $v_{c\ell}$ labeled by $LINK_\ell \cup \Omega[D^\ell]$, and attach $v_{c\ell}$ as a child to $v_{b\ell-1}$.
- For each remaining atom A of $\Omega(D^\ell)$, we create a new vertex, label it with $\{A\}$, and attach it as a leaf to $v_{c\ell}$. (Note that these remaining atoms, if any, stem from other elements of Δ , given that a variable may occur in several 3-sets.)

- Then, create a vertex $v_{a\ell}$ of T , label it by the set of atoms $BLOCKA^\ell$, and attach it as a child of $v_{c\ell}$. The vertex $v_{a\ell}$, in turn, has as only child a vertex $v_{b\ell}$ labeled by $BLOCKB^\ell$.

It is not hard to check that (T, λ) is indeed a valid query decomposition.

Let us now prove the *only-if* part. Assume (T, λ) is a width 4 query decomposition of the above defined query Q . By Proposition 3.3, we also assume, w.l.o.g., that (T, λ) is a pure query decomposition. Since Q is connected, also T is connected.

We observe a number of relevant facts and make some assumptions.

FACT 1: By Lemma 3.4 for each $0 \leq \ell \leq s$, there must exist adjacent vertices $v_{a\ell}$ and $v_{b\ell}$ such that $\lambda(v_{a\ell}) \cup \lambda(v_{b\ell}) = BLOCKA^\ell \cup BLOCKB^\ell$.

FACT 2: It holds that $S \subseteq \text{var}(v_{a\ell})$ and $S \subseteq \text{var}(v_{b\ell})$. In fact, if this were not the case, then both vertices would miss variables from S , but since all variables of S occur together in other pairs of adjacent vertices, this would violate the *connectedness condition* and is thus impossible.

FACT 3: From the latter, and from the fact that the sets S'_a, S''_a, S_b^0 , and S_c^0 form a partition of S , it follows that each of the vertices $v_{a\ell}$ and $v_{b\ell}$ contains a q atom, a p_a atom, a p_b atom, and a p_c atom. Without loss of generality, we can thus make the following assumption.

ASSUMPTION: For $0 \leq \ell \leq s$ we have $Z_\ell \in \text{var}(v_{a\ell})$ and $Y_\ell \in \text{var}(v_{b\ell})$.

FACT 4: For $1 \leq \ell \leq s$, there exists a vertex $v_{c\ell}$ that lies on the unique path from $v_{b\ell-1}$ to $v_{a\ell}$ such that $\{Y_{\ell-1}, Z_\ell\} \subseteq \text{var}(v_{c\ell})$. This can be seen as follows. For any variable ϑ , a ϑ -*path* is a path π in T such that the variable ϑ occurs in the label $\lambda(v)$ of any vertex v of π . The atom $\text{link}(Y_{\ell-1}, Z_\ell)$ must belong to the set $\lambda(v'_{c\ell})$ of some vertex $v'_{c\ell}$ of T . Clearly, by the *connectedness* condition, $v'_{c\ell}$ is connected via an $Y_{\ell-1}$ -path π_b to $v_{b\ell-1}$ and by an Z_ℓ -path π_a to $v_{a\ell}$. Let π denote the unique path from $v_{b\ell-1}$ to $v_{a\ell}$. Then π , π_a , and π_b intersect at exactly one vertex. This is the desired vertex $v_{c\ell}$.

FACT 5: For $1 \leq \ell \leq s$, $S \subseteq \text{var}(v_{c\ell})$. Trivial, because $v_{c\ell}$ lies on a path from $v_{b\ell-1}$ to $v_{a\ell}$ and $S \subseteq \text{var}(v_{b\ell-1})$ and $S \subseteq \text{var}(v_{a\ell})$. The fact follows by the *connectedness* condition.

FACT 6: For $1 \leq \ell \leq s$ $\text{link}(Y_{\ell-1}, Z_\ell)$ belongs to $\lambda(v_{c\ell})$ and there exists an i with $1 \leq i \leq m$ such that $\Omega[D_i] \subseteq \lambda(v_{c\ell})$; in summary, $\lambda(v_{c\ell}) = \{\text{link}(Y_{\ell-1}, Z_\ell)\} \cup \Omega[D_i]$. Let us prove this. By FACT 5 we know that all variables in S must be covered by $v_{c\ell}$. However, it also holds that $\{Y_{\ell-1}, Z_\ell\} \subseteq \text{var}(v_{c\ell})$ (see FACT 4). To cover the latter variables, there are two alternative choices:

1. both atoms $q(\Pi_5^{\ell-1}, S'_a, Y_{\ell-1})$ and $q(\Pi_1^\ell, S'_a, Z_\ell)$ belong to $\lambda(v_{c\ell})$; or
2. the atom $\text{link}(Y_{\ell-1}, Z_\ell)$ belongs to $\lambda(v_{c\ell})$.

Choice 1 is impossible: there exist no two other atoms $A, B \in \text{atoms}(Q)$ such that $\text{var}(A) \cup \text{var}(B) \cup S'_a = S$. We are thus left with Choice 2. Since the atom $\text{link}(Y_{\ell-1}, Z_\ell)$ does not contain any variable from S , there must be three other atoms in $\lambda(v_{c\ell})$ that together cover S . An inspection of the available atoms shows that the only possibility of covering S by three atoms is via some atom set $\Omega[D_i]$ for $1 \leq i \leq m$. The fact is proved.

FACT 7: For $1 \leq i < j \leq s$ it holds that v_{ai} lies on the unique path in T from v_{ci} to v_{cj} .

Consider the edge $\{v_{ai}, v_{bi}\}$. If we cut this edge from the tree T , then we obtain two disconnected trees T_a (containing v_{ai}) and T_b (containing v_{bi}). Since v_{ci} is connected via a Z_i -path to v_{ai} , but Z_i does not occur in $\text{var}(v_{bi})$, it holds that v_{ci} is contained in T_a . On the other hand, by “iterative” application of Fact 4 and of the *connectedness condition* it follows that there is a path π from v_{bi} to v_{cj} such that for each vertex v of π it holds that $\text{var}(v) \cap \text{Bigvars} \neq \emptyset$, where $\text{Bigvars} = \{Y_h \mid i \leq h < j\} \cup \{Z_h \mid i <$

$h < j\}$. Since $\text{var}(v_{ai}) \cap \text{Bigvars} = \emptyset$, π does not traverse v_{ai} . It follows that v_{ci} belongs to T_b . Therefore, the unique path linking v_{ci} to v_{cj} goes through the edge $\{v_{ai}, v_{bi}\}$, and thus contains the vertex v_{ai} .

FACT 8: For $0 \leq i < j \leq s$ it holds that $\text{var}(v_{ci}) \cap \text{var}(v_{cj}) = S$. By Fact 7, v_{ai} lies on the unique path from v_{ci} to v_{cj} . Therefore by the *connectedness condition* it holds that $\text{var}(v_{ci}) \cap \text{var}(v_{cj}) \subseteq \text{var}(v_{ai})$. Moreover, by Fact 6, no variable from $\text{var}(v_{ai}) - S$ is contained in both $\text{var}(v_{ci})$ and $\text{var}(v_{cj})$. Thus $\text{var}(v_{ci}) \cap \text{var}(v_{cj}) \subseteq S$. On the other hand, by Fact 5, $S \subseteq \text{var}(v_{ci})$ and $S \subseteq \text{var}(v_{cj})$, hence, $S \subseteq \text{var}(v_{ci}) \cap \text{var}(v_{cj})$. In summary, we obtain $\text{var}(v_{ci}) \cap \text{var}(v_{cj}) = S$.

For each $1 \leq \ell \leq s$, denote by D^ℓ the set D_i such that $\Omega[D_i] \subseteq \lambda(v_{c\ell})$ (see Fact 6). By FACT 8 it follows that the sets D^ℓ ($1 \leq \ell \leq s$) are mutually disjoint. But then the union of these sets is of cardinality $3s = r$, and hence the union must coincide with R . Thus s subsets out of Δ cover R and (R, Δ) is a positive instance of EXACT COVER BY 3-SETS. \blacksquare

4 Conjunctive Queries of Bounded Hypertree-Width

4.1 Hypertree Width

Let Q be a (conjunctive) query. A *hypertree* for Q is a triple $\langle T, \chi, \lambda \rangle$, where $T = (N, E)$ is a rooted tree, and χ and λ are labeling functions which associate to each vertex $p \in N$ two sets $\chi(p) \subseteq \text{var}(Q)$ and $\lambda(p) \subseteq \text{atoms}(Q)$. If $T' = (N', E')$ is a subtree of T , we define $\chi(T') = \bigcup_{v \in N'} \chi(v)$. We denote the set of vertices N of T by $\text{vertices}(T)$, and the root of T by $\text{root}(T)$. Moreover, for any $p \in N$, T_p denotes the subtree of T rooted at p .

Definition 4.1 A *hypertree decomposition* of a conjunctive query Q is a hypertree $\langle T, \chi, \lambda \rangle$ for Q which satisfies all the following conditions:

1. for each atom $A \in \text{atoms}(Q)$, there exists $p \in \text{vertices}(T)$ such that $\text{var}(A) \subseteq \chi(p)$;
2. for each variable $Y \in \text{var}(Q)$, the set $\{p \in \text{vertices}(T) \mid Y \in \chi(p)\}$ induces a (connected) subtree of T ;
3. for each vertex $p \in \text{vertices}(T)$, $\chi(p) \subseteq \text{var}(\lambda(p))$;
4. for each vertex $p \in \text{vertices}(T)$, $\text{var}(\lambda(p)) \cap \chi(T_p) \subseteq \chi(p)$.

A hypertree decomposition $\langle T, \chi, \lambda \rangle$ of Q is a *complete decomposition* of Q if, for each atom $A \in \text{atoms}(Q)$, there exists $p \in \text{vertices}(T)$ such that $\text{var}(A) \subseteq \chi(p)$ and $A \in \lambda(p)$.

The *width* of the hypertree decomposition $\langle T, \chi, \lambda \rangle$ is $\max_{p \in \text{vertices}(T)} |\lambda(p)|$. The *hypertree width* $hw(Q)$ of Q is the minimum width over all its hypertree decompositions.

In analogy to join trees and query decompositions, we will refer to Condition 2 above as the *Connectedness Condition*. Note that, by Condition 1, $\chi(T) = \text{var}(Q)$. Hence Condition 4 entails that, for $s_0 = \text{root}(T)$, $\text{var}(\lambda(s_0)) = \chi(s_0)$.

Intuitively, the χ labeling selects the set of variables to be fixed in order to split the cycles and achieve acyclicity; $\lambda(p)$ ‘‘covers’’ the variables of $\chi(p)$ by a set of atoms. Thus, the relations associated to the atoms of $\lambda(p)$ restrict the range of the variables of $\chi(p)$. For the evaluation of query Q , each vertex p of the decomposition is replaced by a new atom whose associated database relation is the projection on $\chi(p)$ of the join of the relations in $\lambda(p)$. This way, we obtain a join tree JT of an acyclic query Q' over database \mathbf{DB}' of size $O(n^k)$, where n is the input size and k is the width of the hypertree decomposition. All the efficient techniques available for acyclic queries can be then employed for the evaluation of Q' .

More technically, Condition 1 and Condition 2 above extend the notion of tree decomposition [24] from graphs to hypergraphs (the hypergraph of a query Q groups the variables of the same atom in one hyperedge [2]). Thus, the pair $\langle T, \chi \rangle$ of a hypertree decomposition $\langle T, \chi, \lambda \rangle$ of a conjunctive query Q , can be seen as the correspondent of a tree decomposition on the query hypergraph. However, the treewidth of $\langle T, \chi \rangle$ (i.e., the maximum cardinality of the χ -labels of the vertices of T) is not an appropriate measure of the width of the hypertree decomposition, because a set of m variables appearing in the same atom should count 1 rather than m for the width. Thus, $\lambda(p)$ provides a set of atoms which “covers” $\chi(p)$ and its cardinality gives the measure of the width of vertex p . It is worthwhile noting that $\langle T, \lambda \rangle$ may violate the classical connectedness condition usually imposed on the variables of the join trees, as it is allowed that a variable X appears in both $\lambda(p)$ and $\lambda(q)$ while it does not appear in $\lambda(s)$, for some vertex s on the path from p to q in T . However, this violation is not a problem, as the variables in $var(\lambda(p)) - \chi(p)$ are meaningless and can be projected out before starting the query evaluation process, because the role of $\lambda(p)$ is just that of providing a binding for the variables of $\chi(p)$.

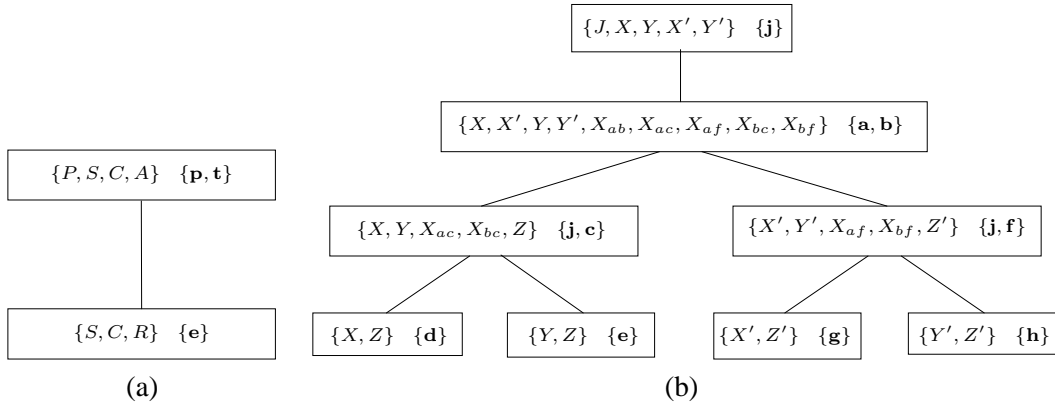


Figure 5: A 2-width hypertree decomposition of (a) query Q_1 ; and (b) query Q_5

Example 4.2 The hypertree width of the cyclic query Q_1 of Example 1.1 is 2; a (complete) 2-width hypertree decomposition of Q_1 is shown in Figure 5.a.

Consider the following conjunctive query Q_5 :

$$ans \leftarrow a(X_{ab}, X, X', X_{ac}, X_{af}) \wedge b(X_{ab}, Y, Y', X_{bc}, X_{bf}) \wedge c(X_{ac}, X_{bc}, Z) \wedge d(X, Z) \wedge e(Y, Z) \wedge f(X_{af}, X_{bf}, Z') \wedge g(X', Z') \wedge h(Y', Z') \wedge j(J, X, Y, X', Y')$$

Q_5 is clearly cyclic, and thus $hw(Q_5) > 1$ (as only acyclic queries have hypertree width equals 1). Figure 5 shows a (complete) hypertree decomposition of Q_5 having width 2, hence $hw(Q_5) = 2$.

Definition 4.1 does not require the presence of all query atoms in a decomposition HD , as it is sufficient that every atom is “covered” by some vertex p of HD (i.e., its variables are included in $\chi(p)$). However, every missing atom can be easily added to complete decompositions.

Lemma 4.3 Given a conjunctive query Q , every k -width hypertree decomposition of Q can be transformed in Logspace into a k -width complete hypertree decomposition of Q .

Proof. Let Q be a conjunctive query and $HD = \langle T, \chi, \lambda \rangle$ a hypertree decomposition of Q . In order to transform HD into a complete decomposition, modify HD as follows. For each atom $A \in atoms(Q)$ such that no vertex $q \in vertices(T)$ satisfies $var(A) \subseteq \chi(q)$ and $A \in \lambda(q)$, create a new vertex v_A with $\lambda(v_A) := \{A\}$ and $\chi(v_A) = var(A)$, and attach v_A as a new child of a vertex $p \in vertex(T)$ s.t. $var(A) \subseteq \chi(p)$. (By Condition 1 of Definition 4.1 such a p must exist.)

This transformation is obviously feasible in Logspace. ■

The acyclic queries are precisely the queries of hypertree width one.

Theorem 4.4 *A conjunctive query Q is acyclic if and only if $hw(Q) = 1$.*

Proof. (*Only if part.*) If Q is an acyclic query, there exists a join tree $JT(Q)$ for Q . Let T be a tree, and f a bijection from $vertices(JT(Q))$ to $vertices(T)$ such that, for any $p, q \in vertices(JT(Q))$, there is an edge between p and q in $JT(Q)$ if and only if there is an edge between $f(p)$ and $f(q)$ in T . Moreover, let λ be the following labeling function: If p is a vertex of $JT(Q)$ and A is the atom of Q associated to p , then $\lambda(f(p)) = \{A\}$. For any vertex $p' \in vertices(T)$ define $\chi(p') = var(\lambda(p'))$. Then, $\langle T, \chi, \lambda \rangle$ is clearly a width 1 hypertree-decomposition of Q .

(*If part.*) Let $HD = \langle T, \chi, \lambda \rangle$ be a width 1 hypertree-decomposition of Q . W.l.o.g., assume that HD is a complete hypertree decomposition. Since HD has width 1, all the λ labels are singletons, i.e., λ associate one atom of Q to each vertex of T .

We next show how to transform HD into a width 1 complete hypertree decomposition of Q such that, for any vertex $p \in vertices(T)$, $\chi(p) = var(A)$, where $\{A\} = \lambda(p)$, and p is the unique vertex labeled with the atom A .

Choose any total ordering \prec of the vertices of T . For any atom $A \in atoms(Q)$, denote by $v(A)$ the \prec -least vertex of T such that $\chi(v(A)) = var(A)$ and $\lambda(v(A)) = \{A\}$. The existence of such a vertex is guaranteed by definition of complete hypertree decomposition and by the hypothesis that every λ label consists of exactly one atom.

For any atom $A \in atoms(Q)$, and for any vertex $p \neq v(A)$ such that $\lambda(p) = \{A\}$, perform the following actions. For any child p' of p , delete the edge between p and p' and let p' be a new child of $v(A)$, hence the subtree $T_{p'}$ is now attached to $v(A)$. Then, delete vertex p . By Condition 3 of Definition 4.1, $\chi(p) \subseteq var(\lambda(p))$. Since $var(\lambda(p)) = var(A) = \chi(v(A))$, we get $\chi(p) \subseteq \chi(v(A))$. Then, it is easy to see that the (transformed) tree T satisfies the connectedness condition.

Eventually, we obtain a new hypertree $H' = \langle T', \chi, \lambda \rangle$ such that $vertices(T') \subseteq vertices(T)$ and H' has the following properties: (i) for any $A \in atoms(Q)$, there exists exactly one vertex $p = v(A)$ of T' such that $\lambda(p) = \{A\}$ and $\chi(p) = var(A)$; (ii) for any vertex p of T' , $p = v(A)$ holds, for some $A \in atoms(Q)$; (iii) H' satisfies the connectedness condition. Thus, H' clearly corresponds to a join tree of Q . ■

4.2 Efficient Query Evaluation

Lemma 4.5 *Let Q be a Boolean conjunctive query over a database \mathbf{DB} , and $HD = \langle T, \chi, \lambda \rangle$ a hypertree decomposition of Q of width k . Then, there exists Q', \mathbf{DB}', JT such that:*

1. Q' is an acyclic (Boolean) conjunctive query answering 'yes' on database \mathbf{DB}' iff the answer of Q on \mathbf{DB} is 'yes'.
2. $||\langle Q', \mathbf{DB}', JT \rangle|| = O(||\langle Q, \mathbf{DB}, HD \rangle||^k)$.
3. JT is a join tree of the query Q' .
4. $\langle Q', \mathbf{DB}', JT \rangle$ is logspace computable from $\langle Q, \mathbf{DB}, HD \rangle$.

Proof. Let Q be a Boolean conjunctive query over a database \mathbf{DB} , and $HD = \langle T, \chi, \lambda \rangle$ a hypertree decomposition of Q of width k . From Lemma 4.3, we can assume that $HD = \langle T, \chi, \lambda \rangle$ is a complete decomposition of Q . W.l.o.g., we also assume Q does not contain any atom A such that $var(A) = \emptyset$.

Note that Q evaluates to *true* on \mathbf{DB} if and only if $\bowtie_{A \in atoms(Q)} rel(A)$ is a non-empty relation, where $rel(A)$ denotes the relation of \mathbf{DB} associated to the atom A , and \bowtie is the natural join operation (with common variables acting as join attributes).

For each vertex $p \in vertices(T)$ define a query $Q(p)$ and a database $\mathbf{DB}(p)$ as follows. For each atom $A \in \lambda(p)$:

- If $var(A) \subseteq \chi(p)$, then A occurs in $Q(p)$ and $rel(A)$ belongs to $\mathbf{DB}(p)$;

- if $(\text{var}(A) \cap \chi(p)) \neq \emptyset$, then $Q(p)$ contains a new atom A' such that $\text{var}(A') = \text{var}(A) \cap \chi(p)$, and $\mathbf{DB}(p)$ contains the corresponding relation $\text{rel}(A')$, which is the projection of $\text{rel}(A)$ on the set of attributes corresponding to the variables in $\text{var}(A')$;

Now, consider the following query \bar{Q} on the database $\bar{\mathbf{DB}} = \bigcup_{p \in \text{vertices}(T)} \mathbf{DB}(p)$.

$$\bar{Q} : \bigwedge_{p \in \text{vertices}(T)} Q(p)$$

By the associative and commutative properties of natural joins, and by the fact that HD is a complete hypertree decomposition, it immediately follows that \bar{Q} on $\bar{\mathbf{DB}}$ is equivalent to Q on \mathbf{DB} .

We build $\langle Q', \mathbf{DB}', JT \rangle$ as follows. JT has exactly the same tree shape of T . For each vertex p , there is precisely one vertex p' in JT , and one relation P' in \mathbf{DB}' . p' is an atom having $\chi(p)$ as arguments and its corresponding relation P' in \mathbf{DB}' is the result of the query $Q(p)$ on $\mathbf{DB}(p)$. Q' is the conjunction of all vertices (atoms) of JT . Q' on \mathbf{DB}' is clearly equivalent to Q on \mathbf{DB} , and JT is a join tree of Q' . Moreover, $\|\mathbf{DB}'\| = O(\|\mathbf{DB}\|^k)$, $\|JT\| = O(\|HD\|)$ and $\|Q'\| = O(\|HD\|)$; thus, $\|\langle Q', \mathbf{DB}', JT \rangle\| = O(\|\langle Q, \mathbf{DB}, HD \rangle\|^k)$.

The transformation is clearly feasible in Logspace. ■

Theorem 4.6 *Given a database \mathbf{DB} , a Boolean conjunctive query Q , and a k -width hypertree decomposition of Q for a fixed constant $k > 0$, deciding whether Q evaluates to true on \mathbf{DB} is LOGCFL-complete.*

Theorem 4.7 *Given a database \mathbf{DB} , a (non-Boolean) conjunctive query Q , and a k -width hypertree decomposition of Q for a fixed constant $k > 0$, the answer of Q on \mathbf{DB} can be computed in time polynomial in the combined size of the input instance and of the output relation.*

Remark. In this section we demonstrated that k -bounded hypertree-width queries are efficiently computable, once a k -width hypertree decomposition of the query is given as (additional) input. In Section 5.2, we will strengthen these results showing that providing the hypertree decomposition in input is unnecessary, as, different from query decompositions, a hypertree decomposition can be computed very efficiently (in L^{LOGCFL} , i.e., in functional LOGCFL).

5 Bounded Hypertree Decompositions are Efficiently Computable

5.1 Normal form

Let $V \subseteq \text{var}(Q)$ be a set of variables, and $X, Y \in \text{var}(Q)$ a pair of variables occurring in Q , then X is $[V]$ -adjacent to Y if there exists an atom $A \in \text{atoms}(Q)$ such that $\{X, Y\} \subseteq (\text{var}(A) - V)$. A $[V]$ -path π from X to Y consists of a sequence $X = X_0, \dots, X_h = Y$ of variables and a sequence of atoms A_0, \dots, A_{h-1} ($h \geq 0$) such that: X_i is $[V]$ -adjacent to X_{i+1} and $\{X_i, X_{i+1}\} \subseteq \text{var}(A_i)$, for each $i \in [0..h-1]$. We denote by $\text{var}(\pi)$ (resp. $\text{atoms}(\pi)$) the set of variables (atoms) occurring in the sequence X_0, \dots, X_h (A_0, \dots, A_{h-1}).

Let $V \subseteq \text{var}(Q)$ be a set of variables occurring in a query Q . A set $W \subseteq \text{var}(Q)$ of variables is $[V]$ -connected if $\forall X, Y \in W$ there is a $[V]$ -path from X to Y . A $[V]$ -component is a maximal $[V]$ -connected non-empty set of variables $W \subseteq (\text{var}(Q) - V)$.

Note that the variables in V do not belong to any $[V]$ -component (i.e., $V \cap C = \emptyset$ for each $[V]$ -component C).

Let C be a $[V]$ -component for some set of variables V . We define:

$$\text{atoms}(C) := \{A \in \text{atoms}(Q) \mid \text{var}(A) \cap C \neq \emptyset\}.$$

Note that, for any set of variables V , and for every atom $A \in \text{atoms}(Q)$ such that $\text{var}(A) \not\subseteq V$, there exists exactly one $[V]$ -component C of Q such that $A \in \text{atoms}(C)$.

Furthermore, let $H = \langle T, \chi, \lambda \rangle$ be a hypertree of Q and $V \subseteq \text{var}(Q)$ a set of variables. We define $\text{vertices}(V, H) = \{p \in \text{vertices}(T) \mid \chi(p) \cap V \neq \emptyset\}$.

For any vertex v of T , we will often use v as a synonym of $\chi(v)$. In particular, $[v]$ -component denotes $[\chi(v)]$ -component; the term $[v]$ -path is a synonym of $[\chi(v)]$ -path; and so on.

Definition 5.1 A hypertree decomposition $HD = \langle T, \chi, \lambda \rangle$ of a conjunctive query Q is in *normal form (NF)* if for each vertex $r \in \text{vertices}(T)$, and for each child s of r , all the following conditions hold:

1. there is (exactly) one $[r]$ -component C_r such that $\chi(T_s) = C_r \cup (\chi(s) \cap \chi(r))$;
2. $\chi(s) \cap C_r \neq \emptyset$, where C_r is the $[r]$ -component satisfying Point 1;
3. $\text{var}(\lambda(s)) \cap \chi(r) \subseteq \chi(s)$.

Note that Condition 2 above entails that, for each vertex $r \in \text{vertices}(T)$, and for each child s of r , $\chi(s) \not\subseteq \chi(r)$. Indeed, $C_r \cap \chi(r) = \emptyset$, and s must contain some variable belonging to the $[r]$ -component C_r .

Lemma 5.2 Let $HD = \langle T, \chi, \lambda \rangle$ be a hypertree decomposition of a conjunctive query Q . Let r be a vertex of T , let s be a child of r , and let C be an $[r]$ -component of Q such that $C \cap \chi(T_s) \neq \emptyset$. Then, $\text{vertices}(C, HD) \subseteq \text{vertices}(T_s)$.

Proof. For any subtree T' of T , let $\text{covered}(T')$ denote the set $\{A \in \text{atoms}(Q) \mid \text{var}(A) \subseteq \chi(v) \text{ for some } v \in \text{vertices}(T')\}$.

Since $C \cap \chi(T_s) \neq \emptyset$, there exists a vertex $p \in \text{vertices}(T_s)$ which also belongs to $\text{vertices}(C, HD)$. We proceed by contradiction. Assume there exists some vertex $q \in \text{vertices}(C, HD)$ such that $q \notin \text{vertices}(T_s)$. By definition of $\text{vertices}(C, HD)$, there exists a pair of variables $\{X, Y\} \subseteq C$ such that $X \in \chi(p)$ and $Y \in \chi(q)$. Since $X, Y \in C$, there exists an $[r]$ -path π from X to Y consisting of a sequence of variables $X = X_0, \dots, X_i, X_{i+1}, \dots, X_\ell = Y$, and a sequence of atoms $A_0, \dots, A_i, A_{i+1}, \dots, A_{\ell-1}$.

Note that $Y \notin \chi(T_s)$. Indeed, $Y \notin \chi(r)$, hence any occurrence of Y in $\chi(v)$, for some vertex v of T_s , would violate Condition 2 of Definition 4.1. Similarly, X only occurs as a variable in $\chi(T_s)$. As a consequence, $A_0 \in \text{covered}(T_s)$ (by Condition 1 of Definition 4.1) and $A_{\ell-1} \notin \text{covered}(T_s)$, hence the $[r]$ -path π leaves T_s , i.e., $\text{atoms}(\pi) \not\subseteq \text{covered}(T_s)$.

Assume w.l.o.g. that the atoms $A_i, A_{i+1} \in \text{atoms}(\pi)$ form the ‘‘frontier’’ of this path w.r.t. T_s , i.e., $A_i \in \text{covered}(T_s)$ and $A_{i+1} \notin \text{covered}(T_s)$, and consider the variable X_{i+1} , which occurs in both A_i and A_{i+1} . X_{i+1} belongs to C , hence it does not occur in $\chi(r)$, and this immediately yields a contradiction to Condition 2 of Definition 4.1. ■

Lemma 5.3 Let $HD = \langle T, \chi, \lambda \rangle$ be a hypertree decomposition of a conjunctive query Q and $r \in \text{vertices}(T)$. If V is an $[r]$ -connected set of variables in $\text{var}(Q) - \chi(r)$, then $\text{vertices}(V, HD)$ induces a (connected) subtree of T .

Proof. We use induction on $|V|$.

Basis. If $|V| = 1$, then V is a singleton, and the statement follows from Condition 2 of Definition 4.1.

Induction Step. Assume the statement is established for set of variables having cardinalities $c \leq h$. Let V be an $[r]$ -connected set of variables such that $|V| = h + 1$, and let $X \in V$ be any variable of V such that $V - \{X\}$ remains $[r]$ -connected. (It is easy to see that such a variable exists.) By the induction hypothesis, $\text{vertices}(V - \{X\}, HD)$ induces a connected subtree of T . Moreover, $\{X\}$ is a singleton, thus $\text{vertices}(\{X\}, HD)$ induces a connected subtree of T , too.

Since $X \in V$, V is $[r]$ -connected, and $|V| > 1$, there exists a variable $Y \in V - \{X\}$ which is $[r]$ -adjacent to X . Hence, there exists an atom $A \in \text{atoms}(Q)$ such that $\{X, Y\} \subseteq \text{var}(A)$. By Condition 1 of Definition 4.1, there exists a vertex $p \in \text{vertices}(T)$ such that $\text{var}(A) \subseteq \chi(p)$. Note that $\text{vertices}(V, HD) = \text{vertices}(V - \{X\}, HD) \cup \text{vertices}(\{X\}, HD)$, and p belongs to both $\text{vertices}(V - \{X\}, HD)$ and $\text{vertices}(\{X\}, HD)$. Then, both sets induce connected subgraphs of T that are, moreover, connected to each other via the vertex p . Thus, $\text{vertices}(V, HD)$ induces a connected subgraph of T , and hence a subtree, because T is a tree. \blacksquare

Theorem 5.4 *For each k -width hypertree decomposition of a conjunctive query Q there exists a k -width hypertree decomposition of Q in normal form.*

Proof. Let $HD = \langle T, \chi, \lambda \rangle$ be any k -width hypertree decomposition of Q . We show how to transform HD into a k -width hypertree decomposition in normal form.

Assume there exist two vertices r and s s.t. s is a child of r , and s violates any condition of Definition 5.1. If s satisfies Condition 1, but violates Condition 2, then $\chi(s) \subseteq \chi(r)$ holds. In this case, simply eliminate vertex s from the tree as shown in Figure 6. It is immediate to see that this transformation is correct.

Assume T_s does not meet Condition 1 of Definition 5.1, and let C_1, \dots, C_h be all the $[r]$ -components containing some variable occurring in $\chi(T_s)$. Hence, $\chi(T_s) \subseteq (\bigcup_{1 \leq i \leq h} C_i \cup \chi(r))$. For each $[r]$ -component C_i ($1 \leq i \leq h$), consider the set of vertices $\text{vertices}(C_i, HD)$. Note that, by Lemma 5.3, $\text{vertices}(C_i, HD)$ induces a subtree of T , and by Lemma 5.2, $\text{vertices}(C_i, HD) \subseteq \text{vertices}(T_s)$, hence $\text{vertices}(C_i, HD)$ induces in fact a subtree of T_s .

For each vertex $v \in \text{vertices}(C_i, HD)$ define a new vertex $\text{new}(v, C_i)$, and let $\lambda(\text{new}(v, C_i)) = \lambda(v)$ and $\chi(\text{new}(v, C_i)) = \chi(v) \cap (C_i \cup \chi(r))$. Note that $\chi(\text{new}(v, C_i)) \neq \emptyset$, because by definition of $\text{vertices}(C_i, HD)$, $\chi(v)$ contains some variable belonging to C_i . Let $N_i = \{\text{new}(v, C_i) \mid v \in \text{vertices}(C_i, HD)\}$. Moreover, for any C_i ($1 \leq i \leq h$), let T_i denote the (directed) graph (N_i, E_i) such that $\text{new}(p, C_i)$ is a child of $\text{new}(q, C_i)$ iff p is a child of q in T . T_i is clearly isomorphic to the subtree of T_s induced by $\text{vertices}(C_i, HD)$, hence T_i is a tree, as well.

Now, transform the hypertree decomposition HD as follows. Delete every vertex in $\text{vertices}(T_s)$ from T , and attach to r every tree T_i for $1 \leq i \leq h$. Intuitively, we replace the subtree T_s by the set of trees $\{T_1, \dots, T_h\}$. By construction, T_i contains a vertex $\text{new}(v, C_i)$ for each vertex v belonging to $\text{vertices}(C_i, HD)$ ($1 \leq i \leq h$). Then, if we let $\text{children}(r)$ denote the set of children of r in the new tree T obtained after the transformation above, it holds that for any $s' \in \text{children}(r)$, there exists an $[r]$ -component C of Q such that $\text{vertices}(T_{s'}) = \text{vertices}(C, HD)$, and $\chi(T_{s'}) \subseteq (C \cup \chi(r))$. Furthermore, it is easy to verify that all the conditions of Definition 4.1 are preserved during this transformation. As a consequence, Condition 2 of Definition 4.1 immediately entails that $(\chi(T_{s'}) \cap \chi(r)) \subseteq \chi(s')$. Hence, $\chi(T_{s'}) = C \cup (\chi(s') \cap \chi(r))$. Thus, any child of r satisfies both Condition 1 and Condition 2 of Definition 5.1.

Now, assume that some vertex $v \in \text{children}(r)$ violates Condition 3 of Definition 5.1. Then, add to the label $\chi(v)$ the set of variables $\text{var}(\lambda(v)) \cap \chi(r)$. Because variables in $\chi(r)$ induce connected subtrees of T , and $\chi(r)$ does not contain any variable occurring in some $[r]$ -component, this further transformation never invalidates any other condition. Moreover, no new vertex is labeled by a set of atoms with cardinality greater than k , then we get in fact a legal k -width hypertree decomposition.

Clearly, $\text{root}(T)$ cannot violate any of the normal form conditions, because it has no parent in T . Moreover, the transformations above never change the parent r of a violating vertex s . Thus, if we apply such a transformation to the children of $\text{root}(T)$, and iterate the process on the new children of $\text{root}(T)$, and so on, we eventually gets a new k -width hypertree decomposition $\langle T', \lambda' \rangle$ of Q in normal form. \blacksquare

If $HD = \langle T, \chi, \lambda \rangle$ is an NF hypertree decomposition of a conjunctive query Q , we can associate a set $\text{treecomp}(s) \subseteq \text{var}(Q)$ to each vertex s of T as follows.

- If $s = \text{root}(T)$, then $\text{treecomp}(s) = \text{var}(Q)$;

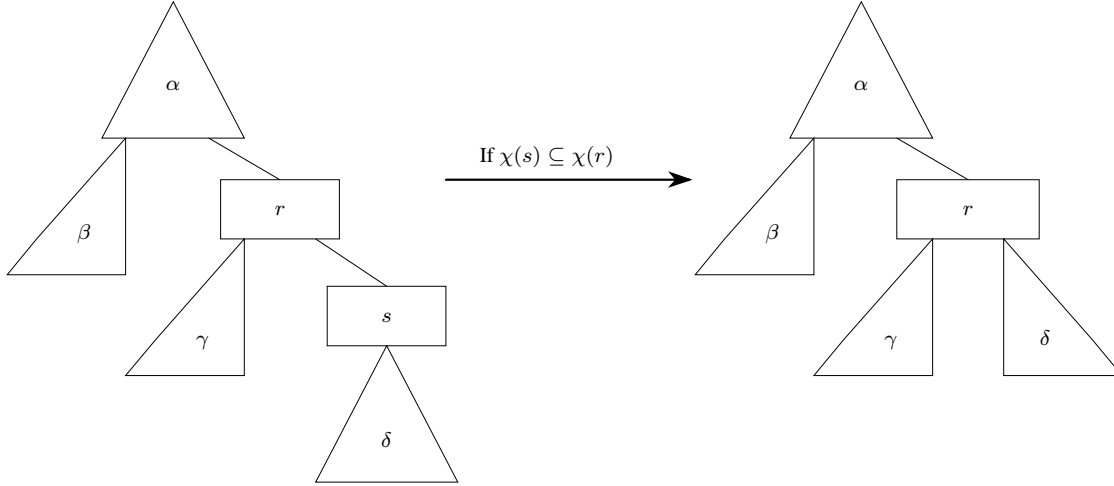


Figure 6: Normalizing a hypertree decomposition

- otherwise, let r be the father of s in T ; then, $treecomps(s)$ is the (unique) $[r]$ -component C such that $\chi(T_s) = C \cup (\chi(s) \cap \chi(r))$.

Note that, since $s \in vertices(T_s)$, also $\chi(T_s) = C \cup \chi(s)$ holds.

Lemma 5.5 Let $HD = \langle T, \chi, \lambda \rangle$ be an NF hypertree decomposition of a conjunctive query Q , v a vertex of T , and $W = treecomps(v) - \chi(v)$. Then, for any $[v]$ -component C such that $(C \cap W) \neq \emptyset$, $C \subseteq W$ holds.

Therefore, the set $\mathcal{C} = \{C' \subseteq var(Q) \mid C' \text{ is a } [v]\text{-component and } C' \subseteq treecomps(v)\}$ is a partition of $treecomps(v) - \chi(v)$.

Proof. Let C be a $[v]$ -component such that $(C \cap W) \neq \emptyset$. We show that $C \subseteq W$. Assume this is not true, i.e., $C - W \neq \emptyset$. By definition of $treecomps(v)$, $\chi(T_v) = treecomps(v) \cup \chi(v)$. Hence, any variable $Y \in (C - W)$ only occurs in the χ label of vertices not belonging to $vertices(T_v)$. However, C is a $[v]$ -component, therefore $C \cap \chi(v) = \emptyset$. As a consequence, $vertices(C, HD)$ induces a disconnected subgraph of T , and thus contradicts Lemma 5.3. ■

Lemma 5.6 Let $HD = \langle T, \chi, \lambda \rangle$ be an NF hypertree decomposition of a conjunctive query Q , and r be a vertex of T . Then, $C = treecomps(s)$ for some child s of r if and only if C is an $[r]$ -component of Q and $C \subseteq treecomps(r)$.

Proof. (If part.) Assume C is an $[r]$ -component of Q and $C \subseteq treecomps(r)$. Let $children(r)$ denote the set of the vertices of T which are children of r . Because $C \subseteq (treecomps(r) - \chi(r))$, C must be included in $\bigcup_{s \in children(r)} \chi(T_s)$. Moreover, for each subtree T_s of T such that $s \in children(r)$, there is a (unique) $[r]$ -component $treecomps(s)$ such that $\chi(T_s) = treecomps(s) \cup (\chi(s) \cap \chi(r))$. Therefore, C necessarily coincides with one of these components, say $treecomps(\bar{s})$ for some $\bar{s} \in children(r)$.

(Only if part.) Assume $C = treecomps(s)$ for some child s of r , and let $C' = treecomps(r)$. By definition of $treecomps(s)$, C is an $[r]$ -component, then $(C \cap \chi(r)) = \emptyset$. Since HD is in normal form, $\chi(T_s) = C \cup (\chi(s) \cap \chi(r))$ and $\chi(T_r) = (C' \cup \chi(r))$. Moreover, s is a child of r , then $vertices(T_s) \subseteq vertices(T_r)$ and thus $\chi(T_s) \subseteq \chi(T_r)$. Therefore, $C \cup (\chi(s) \cap \chi(r)) \subseteq \chi(T_r)$, and hence we immediately get $C \subseteq (C' \cup \chi(r))$. However, $(C \cap \chi(r)) = \emptyset$ and thus $C \subseteq C'$. ■

Lemma 5.7 For any NF hypertree decomposition $HD = \langle T, \chi, \lambda \rangle$ of a query Q , $|vertices(T)| \leq |var(Q)|$ holds.

Proof. Follows from Lemma 5.6, Lemma 5.5, and Condition 2 of the normal form, which states that, for any $v \in \text{vertices}(T)$, $\chi(v) \cap \text{treecom}(v) \neq \emptyset$. Hence, $\text{treecom}(v) - \chi(v) \subset \text{treecom}(v)$ and thus, for any child s of v in T , $\text{treecom}(s)$ is actually a proper subset of $\text{treecom}(v)$. ■

Lemma 5.8 *Let $HD = \langle T, \chi, \lambda \rangle$ be an NF hypertree decomposition of a query Q , s a vertex of T , and C a set of variables such that $C \subseteq \text{treecom}(s)$. Then, C is an $[s]$ -component if and only if C is a $[\text{var}(\lambda(s))]$ -component.*

Proof. Let $V = \text{var}(\lambda(s))$. By Condition 4 of Definition 4.1, $(V \cap \chi(T_s)) \subseteq \chi(s)$. Since HD is in normal form, V satisfies the following property.

$$(1) \quad (V \cap \text{treecom}(s)) \subseteq \chi(s).$$

(Only if part.) Assume $C \subseteq \text{treecom}(s)$ is an $[s]$ -component. From Property 1 above, $C \cap V = \emptyset$ holds. As a consequence, for any pair of variables $\{X, Y\} \subseteq C$, X $[s]$ -adjacent to Y entails X $[V]$ -adjacent to Y . Hence, C is a $[V]$ -connected set of variables. Moreover, $\chi(s) \subseteq V$. Then, any $[V]$ -connected set which is a maximal $[s]$ -connected set is a maximal $[V]$ -connected set as well, and thus C is a $[V]$ -component.

(If part.) Assume $C \subseteq \text{treecom}(s)$ is a $[V]$ -component. Since $\chi(s) \subseteq V$, C is clearly $[s]$ -connected. Thus, $C \subseteq C'$, where C' is an $[s]$ -component and, by Lemma 5.5, $C' \subseteq (\text{treecom}(s) - \chi(s))$ holds. By the “only if” part of this lemma, C' is a $[V]$ -component, therefore C cannot be a proper subset of C' , and $C' = C$ actually holds. Thus, C is an $[s]$ -component. ■

5.2 A LOGCFL Algorithm Deciding k -bounded Hypertree-Width

Figure 7 shows the algorithm $k\text{-decomp}$, deciding whether a given conjunctive query Q has a k -bounded hypertree-width decomposition. In that figure, we give a high level description of an alternating algorithm, to be run on an alternating Turing machine (ATM). The details of how the algorithm can be effectively implemented on a logspace ATM will be given later (see Lemma 5.14).

To each computation tree τ of $k\text{-decomp}$ on input query Q , we associate a hypertree $\delta(\tau) = \langle T, \chi, \lambda \rangle$, called the *witness tree* of τ , defined as follows: For any existential configuration of τ corresponding to the “guess” of some set $S \subseteq \text{atoms}(Q)$ during the computation of $k\text{-decomposable}(C, R)$, for some $[\text{var}(R)]$ -component C , (i.e., to Step 1 of $k\text{-decomp}$), T contains a vertex s . In particular, the vertex s_0 guessed at the initial call $k\text{-decomposable}(\text{var}(Q), \emptyset)$, is the root of T .

There is an edge between vertices r and s of T , where $s \neq s_0$, if S is guessed at Step 1 during the computation of $k\text{-decomposable}(C, R)$, for some $[\text{var}(R)]$ -component C (S and R are the (guessed) sets of atoms of τ corresponding to s and r in T , respectively). We will denote C by $\text{comp}(s)$, and r by $\text{father}(s)$. Moreover, for the root s_0 of T , we define $\text{comp}(s_0) = \text{var}(Q)$.

The vertices of T are labeled as follows. $\lambda(s) = S$ (i.e., $\lambda(s)$ is the guessed set S of atoms corresponding to s), for any vertex s of T . If $s_0 = \text{root}(T)$, let $\chi(s_0) = \text{var}(\lambda(s_0))$; for any other vertex s , let $\chi(s) = \text{var}(\lambda(s)) \cap (\chi(r) \cup C)$, where $r = \text{father}(s)$ and $C = \text{comp}(s)$.

Lemma 5.9 *For any given query Q such that $\text{hw}(Q) \leq k$, $k\text{-decomp}$ accepts Q . Moreover, for any $c \leq k$, each c -width hypertree-decomposition of Q in normal form is equal to some witness tree for Q .*

Proof. Let $HD = \langle T, \chi, \lambda \rangle$ be a c -width NF hypertree decomposition of a conjunctive query Q , where $c \leq k$. We show that there exists an accepting computation tree τ for $k\text{-decomp}$ on input query Q such that $\delta(\tau) = \langle T', \chi', \lambda' \rangle$ “coincides” with HD . Formally, there exists a bijection $f : \text{vertices}(T) \rightarrow \text{vertices}(T')$ such that, for any pair of vertices $p, q \in T$, p is a child of q in T iff $f(p)$ is a child of $f(q)$ in T' , $\lambda(p) = \lambda'(f(p))$, $\lambda(q) = \lambda'(f(q))$, $\chi(p) = \chi'(f(p))$, and $\chi(q) = \chi'(f(q))$.

To this aim, we impose to $k\text{-decomp}$ on input Q the following choices of sets S in Step 1:


```

ALTERNATING ALGORITHM  $k$ -decomp
Input: A non-empty Query  $Q$ .
Result: “Accept”, if  $Q$  has  $k$ -bounded hypertree width; “Reject”, otherwise.

Procedure  $k$ -decomposable( $C_R$ : SetOfVariables,  $R$ : SetOfAtoms)
begin
1) Guess a set  $S \subseteq \text{atoms}(Q)$  of  $k$  elements at most;
2) Check that all the following conditions hold:
   2.a)  $\forall P \in \text{atoms}(C_R), (\text{var}(P) \cap \text{var}(R)) \subseteq \text{var}(S)$  and
   2.b)  $\text{var}(S) \cap C_R \neq \emptyset$ 
3) If the check above fails Then Halt and Reject; Else
   Let  $\mathcal{C} := \{C \subseteq \text{var}(Q) \mid C \text{ is a } [\text{var}(S)]\text{-component and } C \subseteq C_R\}$ ;
4) If, for each  $C \in \mathcal{C}$ ,  $k$ -decomposable( $C, S$ )
   Then Accept
   Else Reject
end;

begin(* MAIN *)
  Accept if  $k$ -decomposable( $\text{var}(Q), \emptyset$ )
end.

```

Figure 7: A non-deterministic algorithm deciding k -bounded hypertree-width

- a) For the initial call k -decomposable($\text{var}(Q), \emptyset$), the set S chosen in Step 1 is $\lambda(\text{root}(T))$.
- b) Otherwise, for a call k -decomposable(C_R, R), if R is the label $\lambda(r)$ of some vertex r , and if r has a child s such that $\text{treecomp}(s) = C_R$, then choose $S = \lambda(s)$ in Step 1.

We use structural induction on trees to prove that, for any vertex $r \in \text{vertex}(T)$, if we denote $f(r)$ by r' , the following equivalences hold: $\lambda(r) = \lambda'(r')$; $\text{treecomp}(r) = \text{comp}(r')$; and $\chi(r) = \chi'(r')$.

Basis: For $r = \text{root}(T)$, we set $f(\text{root}(T)) := \text{root}(T')$. Thus, by choosing $\lambda'(f(r)) = \lambda(r)$ as described at Point a) above, all the equivalences trivially hold.

Induction Step: Assume that the equivalence holds for some vertex $r \in \text{vertices}(T)$. Then, we will show that the statement also holds for every child of r . Let $r' \in \text{vertices}(T')$ denote $f(r)$, and let s be any child of r in T . By Lemma 5.6 and Lemma 5.8, the $[r]$ -component $\text{treecomp}(s)$ coincides with some $[\text{var}(\lambda'(r'))]$ -component $\text{comp}(s')$ corresponding to the call k -decomposable($\text{comp}(s'), \lambda'(r')$) that generated a child s' of r' , which we define to be the image of s , i.e., we set $f(s) := s'$. Since HD is a k -width hypertree decomposition, and the induction hypothesis holds, it easily follows that, by choosing $\lambda(s) = \lambda'(s')$ as prescribed at Point b) above, no check performed in Step 2 of the call k -decomposable($\text{comp}(s'), \lambda'(r')$) can fail.

Next we show that $\chi(s) = \chi'(s')$. Let $C = \text{comp}(s') = \text{treecomp}(s)$, and $V = \text{var}(\lambda(s)) = \text{var}(\lambda'(s'))$. By Condition 4 of Definition 4.1, $V \cap \chi(T_s) \subseteq \chi(s)$ holds. Since HD is in normal form, we can replace $\chi(T_s)$ by $C \cup \chi(s)$ according to Condition 1 of Definition 5.1, and we get $V \cap (C \cup \chi(s)) \subseteq \chi(s)$. Hence, we obtain the following property

$$(1) \quad V \cap C \subseteq \chi(s).$$

Now, consider $\chi'(s')$. By definition of witness tree, $\chi'(s') = V \cap (\chi'(r') \cup C) = V \cap (\chi(r) \cup C)$. By Property (1) above, $V \cap C \subseteq \chi(s)$. Moreover, HD is in NF, and Condition 3 of the normal form entails that $(V \cap \chi(r)) \subseteq \chi(s)$. As a consequence, $\chi'(s') \subseteq \chi(s)$. We claim that this inclusion cannot be proper. Indeed, by definition of $\chi'(s')$, if $\chi'(s') \subset \chi(s)$, there exists a variable $Y \in \chi(s)$ which belongs neither

to $\chi(r)$, nor to C . However, this entails that Y belongs to some other $[r]$ -component and thus s violates Condition 1 of the normal form.

In summary, $k\text{-decomp}$ accepts Q with the accepting computation tree τ determined by the choices described above, and its witness tree $\delta(\tau)$ is a c -width hypertree decomposition of Q in normal form. \blacksquare

Lemma 5.10 *Assume that $k\text{-decomp}$ accepts an input query Q with an accepting computation tree τ and let $\delta(\tau) = \langle T, \chi, \lambda \rangle$ be the corresponding witness tree. Then, for any vertex s of T :*

- a) *if $s \neq \text{root}(T)$, then $\text{comp}(s)$ is a $[\text{father}(s)]$ -component;*
- b) *for any $C \subseteq \text{comp}(s)$, C is an $[s]$ -component if and only if C is a $[\text{var}(\lambda(s))]$ -component.*

Proof. We use structural induction on the tree T .

Basis: Both parts of the lemma trivially holds if s is the root of T . In fact, in this case, we have $\chi(s) = \text{var}(\lambda(s))$, by definition of witness tree.

Induction Step: Assume that the lemma holds for some vertex $r \in \text{vertices}(T)$. Then, we will show that both parts hold for every child of r . The induction hypothesis states that any $[\text{var}(\lambda(r))]$ -component included in $\text{comp}(r)$ is an $[r]$ -component included in $\text{comp}(r)$, and vice versa. Moreover, if $r \neq \text{root}(T)$, then $\text{comp}(r)$ is a $[\text{father}(r)]$ -component; otherwise, i.e., r is the root, $\text{comp}(r) = \text{var}(Q)$, by definition of witness tree. Let $s \in \text{vertices}(T)$ be a child of r and let $V = \text{var}(\lambda(s))$. We first observe that, by definition of the variable labeling χ of the witness tree, it follows that $\text{var}(\lambda(s)) \cap \text{comp}(s) \subseteq \chi(s)$. Hence, the following holds.

Fact 1: $(V - \chi(s)) \cap \text{comp}(s) = \emptyset$.

(Point a.) Immediately follows by the definition of $\text{comp}(s)$ and by the induction hypothesis. Indeed, r is the father of s and by the induction hypothesis any $[\text{var}(\lambda(r))]$ -component included in $\text{comp}(r)$ is an $[r]$ -component included in $\text{comp}(r)$. Thus, in particular, $\text{comp}(s)$ is an $[r]$ -component.

(Only if part, Point b.) Assume that a set of variables $C \subseteq \text{comp}(s)$ is an $[s]$ -component. By Fact 1, $C \cap (V - \chi(s)) = \emptyset$ holds, and for any pair of variables $\{X, Y\} \subseteq C$, X $[s]$ -adjacent to Y entails X $[V]$ -adjacent to Y . Hence, C is a $[V]$ -connected set of variables. Moreover, $\chi(s) \subseteq V$. Then, any $[V]$ -connected set which is a maximal $[s]$ -connected set is a maximal $[V]$ -connected set as well, and thus C is a $[V]$ -component.

(If part, Point b.) We proceed by contradiction. Assume $C \subseteq \text{comp}(s)$ is a $[V]$ -component, but C is not an $[s]$ -component, i.e., C is not a maximal $[s]$ -connected set of variables. Since $\chi(s) \subseteq V$, C is clearly $[s]$ -connected, then it is not maximal. That is, there exists a pair of variables $X \in C$ and $Y \notin C$ such that X is $[s]$ -adjacent to Y , but X is not $[\text{var}(\lambda(s))]$ -adjacent to Y . Let A be any atom proving their adjacency w.r.t. s , i.e., $\{X, Y\} \subseteq \text{var}(A) - \chi(s)$. Hence, because $X \in C$ and X is not $[V]$ -adjacent to Y , it follows that $Y \in (V - \chi(s))$. By Fact 1, $(V - \chi(s)) \cap \text{comp}(s) = \emptyset$, therefore $Y \notin \text{comp}(s)$. In summary, $X \in \text{comp}(s)$ and $Y \notin \text{comp}(s)$. Moreover, $\text{comp}(s) \subseteq \text{comp}(r)$, by Step 4 of $k\text{-decomp}$. Hence, by induction hypothesis, $\text{comp}(s)$ is an $[r]$ -component and thus X is not $[r]$ -adjacent to Y . Consider again the atom A . We get $\{X, Y\} \not\subseteq \text{var}(A) - \chi(r)$. Since $X \in \text{comp}(s)$, the variable Y must belong to $\chi(r)$. However, by definition of witness tree, $Y \in \chi(r)$ and $Y \in \text{var}(\lambda(s))$ entail that $Y \in \chi(s)$, which is a contradiction. \blacksquare

Lemma 5.11 *Assume that $k\text{-decomp}$ accepts an input query Q with an accepting computation tree τ . Let $\delta(\tau) = \langle T, \chi, \lambda \rangle$ be the corresponding witness tree, and $s \in \text{vertices}(T)$. Then, for each vertex $v \in T_s$:*

$$\begin{aligned} \chi(v) &\subseteq \text{comp}(s) \cup \chi(s) \\ \text{comp}(v) &\subseteq \text{comp}(s) \end{aligned}$$

Proof. We use induction on the distance $d(v, s)$ between any vertex $v \in \text{vertices}(T_s)$ and s . The basis is trivial, since $d(v, s) = 0$ means $v = s$.

Induction Step. Assume both statements hold for distance n . Let $v \in \text{vertices}(T_s)$ be a vertex such that $\text{dist}(v, s) = n + 1$. Let v' be the father of v in T_s . Clearly, $\text{dist}(v', s) = n$, thus

(a) $\chi(v') \subseteq (\text{comp}(s) \cup \chi(s))$; and

(b) $\text{comp}(v') \subseteq \text{comp}(s)$.

v is generated by some call $k\text{-decomposable}(\text{comp}(v), \lambda(v'))$. By the choice of v and the definition of witness tree, it must hold (a') $\chi(v) \subseteq (\text{comp}(v) \cup \chi(v'))$, and by Step 4 of the call $k\text{-decomposable}(\text{comp}(v'), \lambda(\text{father}(v')))$ we get (b') $\text{comp}(v) \subseteq \text{comp}(v')$. By (a') and (b'), we obtain (a'') $\chi(v) \subseteq (\text{comp}(v') \cup \chi(v'))$. By (a''), (b), and (a) we get $\chi(v) \subseteq (\text{comp}(s) \cup \chi(s))$. Moreover, (b) and (b') yield $\text{comp}(v) \subseteq \text{comp}(s)$. ■

Lemma 5.12 *If $k\text{-decomp}$ accepts an input query Q , then $hw(Q) \leq k$. Moreover, each witness tree for Q is a c -width hypertree-decomposition of Q in normal form, where $c \leq k$.*

Proof. Assume that τ is an accepting computation tree of $k\text{-decomp}$ on input query Q . We show that $\delta(\tau) = \langle T, \chi, \lambda \rangle$ is an NF c -width hypertree decomposition of Q , for some $c \leq k$.

First, we will prove that $\delta(\tau)$ fulfils all the properties of Definition 4.1 and is thus a hypertree decomposition of Q .

Property 1 : $\forall A \in \text{atoms}(Q) \exists v \in \text{vertices}(T) \text{ s.t. } \text{var}(A) \subseteq \chi(v)$.

We first prove the following claim.

CLAIM A: Let s be any vertex of T , and let $C_r = \text{comp}(s)$. Then, for each $P \in \text{atoms}(C_r)$ it holds that

(a) $\forall A \in (\text{atoms}(Q) - \text{atoms}(C_r)), (\text{var}(P) \cap \text{var}(A)) \subseteq \chi(s)$; and

(b) either $\text{var}(P) \subseteq \chi(s)$ or there exists an $[s]$ -component $C_s \subseteq C_r$ such that $P \in \text{atoms}(C_s)$.

Proof of Claim A. (Part a). We use structural induction on the tree T .

Basis: Part (a) of the claim trivially holds if s is the root of T . In fact, in this case, we have $C_r = \text{comp}(s) = \text{var}(Q)$ and hence $\text{atoms}(C_r) = \text{atoms}(Q)$.

Induction Step: Assume that Part (a) holds for some vertex $s \in \text{vertices}(T)$. Then, we will show that the statement also holds for every child of s . Let $C_r = \text{comp}(s)$ and $V = \text{var}(\lambda(s))$. The induction hypothesis states that $\forall P \in \text{atoms}(C_r)$ and $\forall A \notin \text{atoms}(C_r)$, $\text{var}(P) \cap \text{var}(A) \subseteq \chi(s)$. By Step 4 of $k\text{-decomp}$, for each $[V]$ -component C s.t. $C \subseteq C_r$, T contains a vertex q such that $\text{comp}(q) = C$ and $\text{father}(q) = s$. Moreover, by Lemma 5.10, C is an $[s]$ -component. Let C_s be an $[s]$ -component s.t. $C_s \subseteq C_r$, and let P' belong to $\text{atoms}(C_s)$. By choice of $C_s \subseteq C_r$, P' also belongs to $\text{atoms}(C_r)$. First note that, $\forall A \notin \text{atoms}(C_s)$, we have

$$(1) \text{var}(P') \cap \text{var}(A) \subseteq \chi(s).$$

Indeed, if $\text{var}(A) \subseteq \chi(s)$ (1) is trivial, and if $A \notin \text{atoms}(C_r)$, it follows from the induction hypothesis. Otherwise, i.e. if A contains some variable belonging to another $[s]$ -component included in C_r , it immediately follows by definition of $[s]$ -component. Now, for the component C_s , Step 1 of $k\text{-decomposable}(C_s, \lambda(s))$ guesses a vertex s' such that, $\forall B \in \text{atoms}(C_s)$, $(\text{var}(B) \cap \text{var}(\lambda(s))) \subseteq \text{var}(\lambda(s'))$. In particular, $(\text{var}(P') \cap \text{var}(\lambda(s))) \subseteq \text{var}(\lambda(s'))$. Because $\chi(s) \subseteq \text{var}(\lambda(s))$, this yields $(\text{var}(P') \cap \chi(s)) \subseteq (\text{var}(\lambda(s')) \cap \chi(s))$. By definition of witness tree, $(\text{var}(\lambda(s')) \cap \chi(s)) \subseteq \chi(s')$, hence we get $(\text{var}(P') \cap \chi(s)) \subseteq \chi(s')$. By combining this result with relationship (1) above, we get that $\forall A \notin \text{atoms}(C_s)$ $(\text{var}(P') \cap \text{var}(A)) \subseteq (\text{var}(P') \cap \chi(s)) \subseteq \chi(s')$. Hence, Part (a) of the claim holds even for s' and thus for every child of s in T .

(Part b). Let s be any vertex of T , let $C_r = \text{comp}(s)$, and let P belong to $\text{atoms}(C_r)$. Assume that $\text{var}(P) \not\subseteq \chi(s)$ and that $P \in \text{atoms}(C'_s)$, where C'_s is an $[s]$ -component not included in C_r , i.e., $C'_s \not\subseteq C_r$. Then, there exists a variable $Y \in C'_s$ s.t. $Y \notin C_r$ and there is an $[s]$ -path π from Y to any variable $X \in (\text{var}(P) - \chi(s))$.

Let A be an atom belonging to both $atoms(C_r)$ and $atoms(C'_s)$. Then, $var(A) - \chi(r) \subseteq C_r$ and $var(A) - \chi(s) \subseteq C_s$ hold. As a consequence, $(var(A) \cap C'_s) \subseteq C_r$. Indeed, if this is not true, there exists a variable $Z \in (var(A) - \chi(s))$ such that $Z \in \chi(r)$. By the definition of the χ labeling of a witness tree, this entails that $Z \in (var(A) \cap var(\lambda(r)))$, but $Z \notin var(\lambda(s))$, which contradicts the fact that A satisfies the condition checked at Step 2.a of k -decomp, because τ is an accepting computation tree.

Therefore, there exist two atoms $\{Q', P'\} \subseteq atoms(\pi)$ belonging to $atoms(C'_s)$ and adjacent in π s.t. $Q' \notin atoms(C_r)$, $P' \in atoms(C_r)$, and $var(Q') \cap var(P') \not\subseteq \chi(s)$. However, this contradicts Part (a) of the claim. \diamond

Note that, by Lemma 5.10, in the Step 4 of any call k -decomposable($C, \lambda(r)$) of an accepting computation of k -decomp, $[s]$ -components included in C and $[var(\lambda(s))]$ -components included in C coincide. Thus, Property 1 follows by inductive application of Part b of Claim A. In fact, Part (b) of the claim applied to the root s_0 of T , states that, $\forall A \in atoms(Q)$, either $var(A) \subseteq \chi(s_0)$, or $A \in atoms(C_S)$ for some $[S]$ -component C_S of Q that will be further treated in Step 4 of the algorithm. Thus, $var(A)$ is covered eventually by some chosen set of atoms S , i.e., there exists some vertex s of T , such that $\lambda(s) = S$, and $var(A) \subseteq \chi(s)$.

Property 2 : For each variable $Y \in var(Q)$ the set $\{v \in vertices(T) \mid Y \in \chi(v)\}$ induces a connected subtree of T .

Assume that Property 2 does not hold. Then, there exists a variable $Y \in var(Q)$ and two vertices v_1 and v_2 of T such that $Y \in (\chi(v_1) \cap \chi(v_2))$ but the unique path from v_1 to v_2 in T contains a vertex w such that $Y \notin \chi(w)$. W.l.o.g, assume that v_1 is adjacent to w and that v_2 is a descendant of w in T , i.e., $v_2 \in vertices(T_w)$. There are two possibilities to consider:

- v_1 is a child of w and v_2 belongs to the subtree T_p of another child p of w . However, this would mean that, by Step 4 of k -decomp and by Lemma 5.11, the variables in sets $V_1 = (\chi(v_1) - \chi(w))$ and $V_2 = (\chi(v_2) - \chi(w))$ belong to distinct $[w]$ -components. But this is not possible, because $Y \in (V_1 \cap V_2)$.
- w is a child of v_1 and v_2 belongs to the subtree T_w of T rooted at w . Then, $\lambda(w)$ was chosen as set S in Step 1 of k -decomposable($C, \lambda(v_1)$), where C is a $[v_1]$ -component. Note that $Y \in \chi(v_1)$ entails $Y \notin C$, by definition of $[v_1]$ -component. Since v_2 belongs to the subtree T_w , by Lemma 5.11 it holds that $\chi(v_2) \subseteq (C \cup \chi(w))$. This is a contradiction, because $Y \in \chi(v_2)$, but Y belongs neither to $\chi(w)$, nor to C .

Property 3 : $\forall p \in vertices(T), \chi(p) \subseteq var(\lambda(p))$.

Follows by definition of the χ labeling of a witness tree.

Property 4 : $\forall p \in vertices(T), var(\lambda(p)) \cap \chi(T_p) \subseteq \chi(p)$.

Let v be any vertex in T_p , and let $V = var(\lambda(p))$. By Lemma 5.11, $\chi(v) \subseteq comp(p) \cup \chi(p)$. Hence, $V \cap \chi(v) \subseteq V \cap comp(p) \cup \chi(p)$, because Property 3 holds for p . However, by definition of witness tree, $(V \cap comp(p)) \subseteq \chi(p)$, and thus $(V \cap \chi(v)) \subseteq \chi(p)$.

Thus, $\delta(\tau)$ is a hypertree decomposition of Q . Let c be the width of $\delta(\tau)$. Since Step 1 of k -decomp only chooses set of atoms having cardinality bounded by k , $c \leq k$ holds.

Moreover, $\delta(\tau)$ is in normal form. Indeed, Condition 2 and Condition 3 of Definition 5.1 hold by Step 2.b of k -decomp, and by definition of the χ labeling of a witness tree. Finally, since $\delta(\tau)$ is a hypertree decomposition, by Lemma 5.2, Lemma 5.11, and the definition of the χ labeling of a witness tree, we get that Condition 1 holds for $\delta(\tau)$, too. \blacksquare

By combining Lemma 5.9 and Lemma 5.12 we get:

Theorem 5.13 k -decomp accepts an input query Q if and only if $hw(Q) \leq k$.

Lemma 5.14 *k -decomp can be implemented on a logspace ATM having polynomially bounded tree-size.*

Proof. Let us refer to logspace ATMs with polynomially bounded tree-size as LOGCFL-ATMs. We will outline how the algorithm k -decomp can be implemented on an LOGCFL-ATM M .

We first describe the data-structures used by M . Instead of manipulating atoms directly, *indices* of atoms will be used in order to meet the logarithmic space bound. Thus the i -th atom occurring in the given representation of the input query Q will be represented by integer i . Sets of at most k atoms, k -sets for short, are represented by k -tuples of integers; since k is fixed, representing such sets requires logarithmic space only. Variables are represented as integers, too.

If R is a k -set, then a $[var(R)]$ -component C is represented by a pair $\langle rep(R), first(C) \rangle$, where $rep(R)$ is the representation of the k -set R , and $first(C)$ is the smallest integer representing a variable of the component C . For example, the \emptyset -component $var(Q)$ is represented by the pair $\langle rep(\emptyset), 1 \rangle$. It is thus clear that $[var(R)]$ -components can be represented in logarithmic space, too.

The main data structures carried with each configuration of M consist of (the representations of):

- a k -set R ,⁴
- a $[var(R)]$ -component C_R ,
- a k -set S , and
- a $[var(S)]$ -component C .

Not all these items will contain useful data in all configurations. We do not describe further auxiliary logspace data structures that may be used for control tasks and for other tasks such as counting or for performing some of the SL subtasks described below.

We are now ready to give a description of the computation M performs on an input query Q .

To facilitate the description, we will specify some subtasks of the computation, that are themselves solvable in LOGCFL, as macro-steps without describing their corresponding computation (sub-)trees. We may imagine a macro-step as a special kind of configuration – termed *oracle configuration* – that acts as an oracle for the subtask to be solved.

Each oracle configuration can be *normal* or *converse*.

A *normal* oracle configuration has the following effect. If the subtask is negative, this configuration has no children and amounts to a REJECT. Otherwise, its value (ACCEPT or REJECT) is identical to the value of its unique successor configuration.

A *converse* oracle configuration has the following effect. If the subtask is negative, this configuration has no children and amounts to an ACCEPT. Otherwise, its value (ACCEPT or REJECT) is identical to the value of its unique successor configuration.

From the definition of logspace ATM with polynomial tree-size, it follows that any polynomially tree-sized logspace ATM M with LOGCFL oracle configurations (where an oracle configuration contributes 1 to the size of an accepting subtree) is equivalent to a standard logspace ATM having polynomial tree size.

M is started with R initialized to the empty set and C_R having value $var(Q)$.

We describe the evolution of M corresponding to a procedure call k -decomposable(C_R, R).

Instruction 1 is performed by guessing an arbitrary k -set S of atoms.

The “Guess” phase of Instruction 1 is implemented by an existential configuration of the ATM. (Actually, it is implemented by a subtree of existential configurations, given that a single existential configuration can only guess one bit; note however, that each accepting computation tree will contain only one branch of this subtree.)

⁴The separate representation of R is actually slightly redundant, given that R also occurs in the description of the $[R]$ -component C_R .

Checking Step 2 is in symmetric logspace (SL). The most difficult task is to enumerate atoms of $atoms(C_R)$, which in turn – as most substantial subtask – requires to enumerate the variables of C_R . Remember that C_R is given in the form $\langle rep(R), i \rangle$ as described above. Thus, enumerating C_R amounts to cycling over all variables j and checking whether j is $[R]$ -connected to i . The latter subtask is easily seen to be in SL because it essentially amounts to a connectedness-test of two vertices in an undirected graph. It follows that the entire checking-task of Instruction 2 is in SL. Since $SL \subseteq LOGCFL$, this correspond to a LOGCFL-subtask. We can thus assume that the checking-task is performed by some normal oracle configuration. If the oracle computation fails at some branch, the branch ends in a REJECT, otherwise, the guessed k -set S corresponding to that branch satisfies all the conditions checked by Step 2 of k -decomp.

Steps 3 and 4 together intuitively correspond to a “big” universal configuration that universally quantifies over all subtrees corresponding to the procedure calls k -decomposable(C, S) for all $C \in \mathcal{C}$. This could be realized as follows. First, a subtree of universal configurations enumerates all candidates $C_i = \langle rep(S), i \rangle$ for $1 \leq i \leq |var(Q)|$, for $[var(S)]$ -components. Each branch of this subtree (of polynomial depth) computes exactly one candidate C_i . Each such branch is expanded by a converse oracle configuration checking whether C_i is effectively a $[var(S)]$ -component contained in C_R . Thus, branches that do not correspond to such a component are terminated with an ACCEPT configuration (they are of no interest), while all other branches are further expanded. Each branch C_i of the latter type is expanded by the subtree corresponding to the recursive call k -decomposable(C_i, S).

We have thus completely described a logspace ATM M with oracle configurations that implements k -decomp. It is easy to see that this machine has polynomial accepting computation trees. In fact, this is seen from the fact that there exist only a polynomial number of choices for set S in Step 1, and that no such set is chosen twice in any accepting computation tree. ■

From the lemmas above and Proposition 2.3, we obtain the following theorem.

Theorem 5.15 *Deciding whether a conjunctive query Q has k -bounded hypertree-width is in LOGCFL.*

In fact, the following proposition states that an accepting computation tree of a bounded-treesize logspace ATM can be *computed* in (the functional version of) LOGCFL.

Proposition 5.16 ([14]) *Let M be a bounded-treesize logspace ATM recognizing a language A . It is possible to construct a L^{LOGCFL} transducer T which for each input $w \in A$ outputs a single (polynomially-sized) accepting tree for M and w .*

By Lemma 5.9 and Lemma 5.12, we have a one-to-one correspondence between the NF k -width hypertree decompositions and accepting computation trees of k -decomp. Thus, by Proposition 5.16, we get that hypertree decompositions are efficiently *computable*.

Theorem 5.17 *Computing a k -bounded hypertree decomposition (if any) of a conjunctive query Q is in L^{LOGCFL} , i.e., in functional LOGCFL.*

Since LOGCFL is closed under L^{LOGCFL} reductions [14], the two following statements follow from the theorem above and Theorem 4.6 and Theorem 4.7, respectively.

Corollary 5.18 *Deciding whether a k -bounded hypertree-width query Q evaluates to true on a database \mathbf{DB} is LOGCFL-complete.*

Corollary 5.19 *The answer of a (non-Boolean) k -bounded hypertree-width query Q can be computed in time polynomial in the combined size of the input instance and of the output relation.*

5.3 A Datalog Program Recognizing Queries of k -bounded Hypertree-Width

In this section we show a straightforward polynomial-time implementation of the LOGCFL algorithm above. In particular, we reduce (in polynomial time) the problem of deciding whether there exists a k -bounded hypertree-width decomposition of a given conjunctive query Q to the problem of solving a Datalog program.

First, we associate an identifier (e.g., some constant number) to each k -vertex (non empty subset of Q consisting of k atoms at most) R for Q , and to each $[R]$ -component C for any k -vertex R of Q . Moreover, we have a new identifier $root$ which intuitively will be the root of any possible tree-decomposition, and a new identifier $varQ$ which encodes the set of all the variables of the query and hence is seen as a component including any subset of $var(Q)$.

Then, we compute the following relations:⁵

- k -vertex(\cdot): Contains a tuple $\langle R \rangle$ for each k -vertex R of Q .
- $component(\cdot, \cdot)$: Contains a tuple $\langle C_R, R \rangle$ for each $[R]$ -component C_R of some k -vertex R . Moreover, it contains the tuple $\langle varQ, root \rangle$.
- $meets-condition(\cdot, \cdot, \cdot)$: Contains any tuple $\langle S, R, C_R \rangle$ s.t. S and R are k -vertices, C_R is an $[R]$ -component, and the following conditions hold: $var(S) \cap C_R \neq \emptyset$, and $\forall P \in atoms(C_R) var(P) \cap var(R) \subseteq var(S)$. Moreover, it contains a tuple $\langle S, root, varQ \rangle$ for any k -vertex S .
- $subset(\cdot, \cdot)$: Encodes the standard set-inclusion relationship between $[R]$ -components of Q .

Let \mathcal{P} be the following Datalog program:

1. k -decomposable(R, C_R) \leftarrow k -vertex(S), $meets-conditions(S, R, C_R)$, \neg $undecomposable(S, C_R)$
2. $undecomposable(S, C_R)$ \leftarrow $component(C_S, S)$, $subset(C_S, C_R)$, \neg k -decomposable(S, C_S).

It is easy to see that $hw(Q) \leq k$ if and only if $\mathcal{P} \models k$ -decomposable($root, varQ$).

Note that \mathcal{P} is locally stratified on the base relations to which it is applied, and it is clearly evaluable in polynomial time.

6 Bounded Hypertree-Width vs Related Notions

Many relevant cyclic queries are – in a precise sense – close to acyclic queries because they can be decomposed via low bandwidth decompositions to acyclic queries.

The main classes of bounded-width queries considered in database theory and in artificial intelligence are the following:

- *Queries of bounded treewidth.* Treewidth is the best-known graph theoretic measure of tree-similarity. The concept of treewidth is based on the notion of *tree-decomposition* of a graph. The concept of treewidth is easily generalized to hypergraphs and thus to conjunctive queries. Conjunctive queries of bounded treewidth can be answered in polynomial time [7]. For each fixed k , deciding whether a query has treewidth k is in LOGCFL.
- *Queries of bounded degree of cyclicity.* This concept was introduced by Gyssens et al. [18, 17] and is based on the notion of *hinge-tree decomposition*. The smaller the degree of cyclicity of a hypergraph, the more the hypergraph resembles an acyclic hypergraph. Hypergraphs of bounded treewidth have also bounded degree of cyclicity, but not vice-versa. Queries of bounded degree of cyclicity can be recognized and processed in polynomial time [17].

⁵For the sake of clarity, we directly refer to objects by means of their associated identifiers (which we also use as logical terms in the program).

- *Queries of bounded query-width.* This notion is based on the concept of *query decomposition* [7]. Any hinge-tree decomposition of width k is also a query decomposition of width k , but not vice-versa. Thus query decompositions are the most general (i.e., most liberal) decompositions. It follows from results in [7] that queries of bounded query-width can be answered in polynomial time, once a query decomposition is given.

Thus, the class of queries of bounded query-width is the widest of the above mentioned classes of tractable cyclic queries. We next show that this class is properly included in the class of queries of bounded hypertree-width. More precisely, we show that every k -width query-decomposition corresponds to an equivalent k -width hypertree-decomposition, but the converse is not true, in general. Recall that $hw(Q)$ and $qw(Q)$ denote the hypertree width and the query width of a conjunctive query Q .

Theorem 6.1

- For each conjunctive query Q it holds that $hw(Q) \leq qw(Q)$.
- There exist queries Q such that $hw(Q) < qw(Q)$.

Proof. (*Point a.*) Let Q be a conjunctive query and $\langle T, \lambda \rangle$ a query decomposition of Q . W.l.o.g., assume Q is pure (i.e., labels contain only atoms, see Section 3.1). Then, (T, χ, λ) is a hypertree decomposition of Q , where, for any vertex v of T , $\chi(v)$ consists of the set of variables $var(\lambda(v))$ occurring in the atoms $\lambda(v)$. Indeed, because the properties of query decompositions holds for $\langle T, \lambda \rangle$, $\langle T, \chi, \lambda \rangle$ verifies Condition 1 and 2 of Definition 4.1. Condition 3 and 4 follows immediately, as $\chi(p) = var(\lambda(p))$ by construction. Therefore, $hw(Q) \leq qw(Q)$.

(*Point b.*) The query Q_4 of Example 4.2 has no query decompositions of width 2, and it is easy to see that $qw(Q_4) = 3$. However, $hw(Q_4) = 2$, as witnessed by the hypertree decomposition shown in Figure 5. ■

Acknowledgment

We thank Cristinel Mateis for his useful comments on an earlier version of this paper.

References

- [1] S. Abiteboul, R. Hull, V. Vianu. *Foundations of Databases*, Addison-Wesley Publishing Company, 1995.
- [2] C. Beeri, R. Fagin, D. Maier, and M. Yannakakis. On the Desiderability of Acyclic Database Schemes. *Journal of ACM*, 30(3):479–513, 1983.
- [3] P.A. Bernstein, and N. Goodman. The power of natural semijoins. *SIAM J. Computing*, 10(4):751–771, 1981.
- [4] H. L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM Journal on Computing*, 25(6):1305-1317, 1996.
- [5] A.K. Chandra, D.C. Kozen, and L.J. Stockmeyer. Alternation. *Journal of ACM*, 26:114–133, 1981.
- [6] A.K. Chandra and P.M. Merlin. Optimal Implementation of Conjunctive Queries in relational Databases. In *ACM Symp. on Theory of Computing (STOC'77)*, pp.77–90, 1977.
- [7] Ch. Chekuri and A. Rajaraman. Conjunctive Query Containment Revisited. In *Proc. International Conference on Database Theory 1997 (ICDT'97), Delphi, Greece, Jan. 1997*, Springer LNCS, Vol. 1186, pp.56–70, 1997.

- [8] A. D’Atri, and M. Moscarini. Recognition algorithms and design methodologies for acyclic database schemes. In *Advances in Computing Research*, 3:43–68, P. Kannellakis, JAI press, London, 1986.
- [9] R. Fagin. Degrees of acyclicity for hypergraphs and relational database schemes. *J. ACM*, 30(3):514–550, 1983.
- [10] R. Fagin, A.O. Mendelzon, and J.D. Ullman. A simplified universal relation assumption and its properties. *ACM Trans. on Database Systems*, 7(3):343–360, 1982.
- [11] M.R. Garey and D.S. Johnson. *Computers and Intractability. A Guide to the Theory of NP-completeness*. Freeman and Comp., NY, USA, 1979.
- [12] N. Goodman, and O. Shmueli. Tree queries: a simple class of relational queries. *ACM Trans. on Database Systems*, 7(4):653–677, 1982.
- [13] G. Gottlob, N. Leone, and F. Scarcello. The Complexity of Acyclic Conjunctive Queries. Technical Report DBAI-TR-98/17, available on the web as: <http://www.dbai.tuwien.ac.at/staff/gottlob/acyclic.ps>, or by email from the authors. An extended abstract concerning part of this work has been published in *Proc. of the IEEE Symposium on Foundations of Computer Science (FOCS’98)*, pp.706–715, Palo Alto, CA, 1998.
- [14] G. Gottlob, N. Leone, and F. Scarcello. Computing LOGCFL Certificates. Technical Report DBAI-TR-98/19, available on the web as: <http://www.dbai.tuwien.ac.at/staff/gottlob/certificates.ps>, or by email from the authors.
- [15] G. Gottlob, N. Leone, and F. Scarcello. Hypertree Decompositions and Tractable Queries. Technical Report DBAI-TR-98/21, available on the web as: <http://www.dbai.tuwien.ac.at/staff/gottlob/hypertrees.ps>, or by email from the authors.
- [16] S.H. Greibach. The Hardest Context-Free Language. *SIAM Journal on Comput.*, 2(4):304–310, 1973.
- [17] M. Gyssens, P.G. Jeavons, and D.A. Cohen. Decomposing constraint satisfaction problems using database techniques. *Artificial Intelligence*, 66:57–89, 1994.
- [18] M. Gyssens, and J. Paradaens. A Decomposition Methodology for Cyclic Databases. In *Advances in Database Theory*, volume 2, pp. 85-122. Plenum Press New York, NY, 1984.
- [19] D. S. Johnson. A Catalog of Complexity Classes. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume A, chapter 2, pp.67–161. Elsevier Science Publishers B.V. (North-Holland), 1990.
- [20] Ph. G. Kolaitis and M. Y. Vardi. Conjunctive-Query Containment and Constraint Satisfaction. To appear in: *Proc. of Symp. on Principles of Database Systems (PODS’98)*, 1998.
- [21] D. Maier. *The Theory of Relational Databases*, Rochville, Md, Computer Science Press, 1986.
- [22] C.H. Papadimitriou and M. Yannakakis. On the Complexity of Database Queries. In *Proc. of Symp. on Principles of Database Systems (PODS’97)*, pp.12–19, Tucson, Arizona, 1997.
- [23] X. Qian. Query Folding. In *Proc. Internatl. Conf. on Data Engineering, New Orleans, Louisiana, Feb.26-March 1, 1996 (ICDE’96)*, pp.48–55, 1996.
- [24] N. Robertson and P.D. Seymour. Graph Minors II. Algorithmic Aspects of Tree-Width. *J. Algorithms*, 7:309-322, 1986.

- [25] W.L. Ruzzo. Tree-Size Bounded Alternation. *Journal of Computer and System Sciences*, 21:218–235, 1980.
- [26] S. Skyum and L.G. Valiant. A complexity theory based on Boolean algebra. *J. Assoc. Comput. Mach.*, 32:484–502, 1985.
- [27] I.H. Sudborough. Time and Tape Bounded Auxiliary Pushdown Automata. In *Mathematical Foundations of Computer Science 1977*, LNCS 53, Springer-Verlag, pp.493–503, 1977.
- [28] R.E. Tarjan, and M. Yannakakis. Simple linear-time algorithms to test chordality of graphs, test acyclicity of hypergraphs, and selectively reduce acyclic hypergraphs. *SIAM J. Computing*, 13(3):566-579, 1984.
- [29] J.D. Ullman. *Principles of Database and Knowledge Base Systems, Vol II*, Computer Science Press, Rockville, MD, 1989.
- [30] J.D. Ullman. Information integration using logical views. In *Database Theory - ICDT'97*, LNCS 1186, pp. 19–40, Springer-Verlag, 1997.
- [31] M. Vardi. Complexity of Relational Query Languages. In *Proc. of 14th ACM STOC*, pp. 137–146, 1982.
- [32] M. Yannakakis. Algorithms for Acyclic Database Schemes. In *Proc. of Int. Conf. on Very Large Data Bases (VLDB'81)*, pp. 82–94, C. Zaniolo and C. Delobel Eds., Cannes, France, 1981.
- [33] C.T. Yu, and M.Z. Ozsoyoglu. On determining tree-query membership of a distributed query. *Infor*, 22(3):261–282, 1984.