

Lecture 4: Acyclic Conjunctive Queries

Instructor: Paris Koutris

Recall that we showed that the combined complexity of evaluating CQs is NP-complete. In this lecture, we show how to obtain polynomial time algorithms for some classes of CQs. Let us start with an example.

Example 4.1. Consider the boolean path query, which asks whether there is path of length k in the directed graph represented by the edge relation R .

$$P^k() : -R(x_1, x_2), R(x_2, x_3), \dots, R(x_k, x_{k+1})$$

Assume that $|R| = n$. Consider the following algorithm: we first project R on the attribute x_2 , then compute the semi-join with R_2 , project the result on attribute x_3 , and so on. Observe that at every point the size of the intermediate result is at most n . Hence, we can implement the algorithm with running time $O(kn)$. The running time in this case is polynomial (in combined complexity)!

Let's see now how we can generalize this intuition to compute a larger class of queries. Notice that the output of a query may be in the worst case exponential in the input size; in this case, designing an algorithm that runs in polynomial time in the size of the input is not possible. Thus, we want to design an algorithm that is polynomial with respect to the *input and the output size*. We will show that such an algorithm is possible for the class of conjunctive queries called *acyclic CQs*.

4.1 Acyclic CQs: Definition

Definition 4.2 (Query Hypergraph). The hypergraph $\mathcal{H}(q) = (V, E)$ associated to a conjunctive query q is defined as follows. The set of vertices V consists of all variables in the body of q , $\mathbf{vars}(q)$, while the set E of hyperedges contains, for each atom in the query, the set of variables that appear in this atom.

Exercise 4.3. Construct the hypergraphs for the query $q_0(x, y, z) : -R(x, y), S(y, z), T(z, x)$ and $q_1() : -R(x, y), S(y, z), T(z, w)$.

Intuitively, a CQ is acyclic if the hypergraph contains no cycles. If the hypergraph is actually a graph (which happens if every relation is binary), then query acyclicity coincides with the graph being a tree. But the situation is not as clear when the hyperedges have arity more than 2. To formally define acyclicity, we start with the following algorithm, called the GYO algorithm (from Graham-Yu-Ozsoyoglu).

Definition 4.4. An ear in a hypergraph \mathcal{H} is a hyperedge e for which we can divide its nodes into two groups: (a) those that appear exclusively in e , and (b) those that are contained in another hyperedge f of \mathcal{H} . In this case, f is called a witness of e .

For example, if a hyperedge is isolated, or if it is contained within another hyperedge, it is by definition an ear.

Example 4.5. For the query q_1 , both R and T are ears, but S is not. The query q_0 contains no ears.

The GYO algorithm works as follows:

- While \mathcal{H} has ears do:
 - Pick an ear e of H .
 - Remove all vertices that are exclusively in the ear e ; remove the ear e from E .

Example 4.6. If we apply the GYO algorithm to $\mathcal{H}(q_0)$, since no hyperedge is an ear, the output will be just $\mathcal{H}(q_0)$. For q_1 , suppose we start with T . Then, the remaining hypergraph is $R(x, y), S(y, z)$. Now both R, S are ears; if we remove S we get only $R(x, y)$. Since any isolated hyperedge is an ear, the resulting hypergraph is the empty one.

Let $\text{GYO}(\mathcal{H})$ be the resulting hypergraph if we apply the GYO algorithm to \mathcal{H} . We say that \mathcal{H} is *acyclic* if $\text{GYO}(\mathcal{H})$ is the empty hypergraph.

Definition 4.7 (CQ Acyclicity). A conjunctive query q is acyclic if $\text{GYO}(\mathcal{H}(q))$ is the empty hypergraph.

The original GYO algorithm is a slight variation of the above algorithm. Instead of removing ears one at time, it repeats the following 2 steps as long as they can be applied: (a) delete a vertex v that belongs to a single hyperedge, and (b) delete a hyperedge that is contained in another hyperedge. This variation guarantees that the remaining hypergraph will always be the same, independent of the sequence with which we remove hyperedges or vertices.

This above definition of acyclicity is formally called α -acyclicity. There are other definitions of acyclicity for hypergraphs (Berge-acyclicity, β -acyclicity, γ -acyclicity), but they are more restrictive. Next, we will see some equivalent definitions of acyclic CQs.

A *forest* is a graph that is a collection of disjoint trees.

Definition 4.8 (Join Forest). A join forest for a conjunctive query q is a forest $F = (V, E)$ whose vertices are the atoms in q and such that for each pair of atoms R, S having variables in common the following conditions hold:

1. R, S are in the same connected component of F , and
2. all variables common to R and S occur on the unique path from R to S .

In the case where F is a tree, it is called a *join tree* of the query q . It turns out that we can equivalently characterize acyclicity by looking at whether a CQ admits a join forest.

Lemma 4.9. A conjunctive query q is acyclic if and only if it has a join forest.

We next show a third equivalent characterization of acyclic CQs. Say that, given a CQ q , we want to remove from a relation R the tuples that we are certain that they do not participate in the final result of q , called *dangling tuples*. We can attempt to do this using the *semijoin* operator. A semijoin between relations R, S is defined as:

$$R \ltimes S = \pi_{\text{att}(R)}(R \bowtie S)$$

In other words, the result is the set of tuples from R that we are certain that they join with a tuple from S . Semijoins are often used in distributed databases to reduce the communication cost of joining two relations. An instance without any dangling tuples is called *globally consistent*. If we can remove all dangling tuples by using only a (finite) sequence semijoin operations, such a sequence is called a *full reducer*.

Example 4.10. Here is one full reducer for the query P :

$$\begin{aligned} S &:= S \ltimes R \\ T &:= T \ltimes S \\ S &:= S \ltimes T \\ R &:= R \ltimes S \end{aligned}$$

The triangle query T does not admit any full reducer. For example, consider the following instance: $\{R(a, a), R(b, b), S(a, a), S(b, b), T(a, b), T(b, a)\}$.

Lemma 4.11. A conjunctive query q is acyclic if and only if it admits a full reducer.

[Alice] contains a detailed proof of this equivalence result. The idea is to compute semi-joins in the order in which we remove the ears in the GYO algorithm (and in the reverse order as well).

4.2 Acyclic CQs: PTIME Algorithm

We will now use join forests and full reducers to compute acyclic CQs. Without loss of generality, we can assume that the query q is connected (since we can compute each connected component separately and then take the cartesian product). The algorithm works as follows:

1. Apply a full reducer to obtain a globally consistent instance.
2. Construct a join tree T for q and choose a root.
3. Choose a join order in which each node is joined with its parent in any bottom-up order.
4. After each join into a relation R , project the result onto the set of attributes that are in $\text{var}(R) \cup \text{head}(q)$.

Theorem 4.12. We can evaluate an acyclic conjunctive query q in time polynomial in the size of the query, the input and the output.

Proof. Why is this algorithm correct? The key property is that if we have a variable x that appears in a node and not its parent in the join tree, the variable will not appear at any ancestor (so it will not be needed further in the evaluation).

To show that we obtain the desired complexity, we will prove that no intermediate result can grow more than a polynomial of the input and output. Consider some node with atom R and children S_1, \dots, S_n . Then the algorithm will replace this node with a new node $T = \pi_{Y \cup Z}(T)$, where $T = R \bowtie S_1 \bowtie \dots \bowtie S_n$, Y are the variables in R , and Z are the variables in $\text{head}(q) \setminus Y$. Since Y, Z are disjoint, $\pi_{Y \cup Z}(T) \subseteq \pi_Y(T) \times \pi_Z(T)$. Now notice that $\pi_Y(T)$ is at most the size of T . Also, $\pi_Z(T)$ is bounded by the size of the output (but only because we took care of dangling tuples by applying the full reducer). \square

Example 4.13. *How does the acyclic algorithm work for the following query?*

$$q(x, t) : -R(x, y, z), S(y, v), T(y, z, u), U(z, u, w), V(u, w, t).$$

Run the GYO algorithm, find the join tree and explain what is the query plan produced by the polynomial-time algorithm.

References

[Alice] S. ABITEBOUL, R. HULL and V. VIANU, “Foundations of Databases.”

[Y81] M. YANNAKAKIS, “Algorithms for acyclic database schemes,” *VLDB 1981*.