

Fast CQ Evaluation

Complexity of CQ

Theorem: It holds that:

- $BQE(\mathbf{CQ})$ is NP-complete (**combined complexity**)
- $BQE[\mathbf{D}](\mathbf{CQ})$ is NP-complete, for a fixed database \mathbf{D} (**query complexity**)
- $BQE[\mathbf{Q}](\mathbf{CQ})$ is in LOGSPACE, for a fixed query $\mathbf{Q} \in \mathbf{CQ}$ (**data complexity**)

Proof:

(NP-membership) Consider a database \mathbf{D} , and a Boolean CQ $\mathbf{Q} :- \text{body}$

Guess a substitution $h : \text{terms}(\text{body}) \rightarrow \text{terms}(\mathbf{D})$

Verify that h is a match of \mathbf{Q} in \mathbf{D} , i.e., $h(\text{body}) \subseteq \mathbf{D}$

(NP-hardness) Reduction from 3-colorability

(LOGSPACE-membership) Inherited from $BQE[\mathbf{Q}](\mathbf{DRC})$

Complexity of CQ

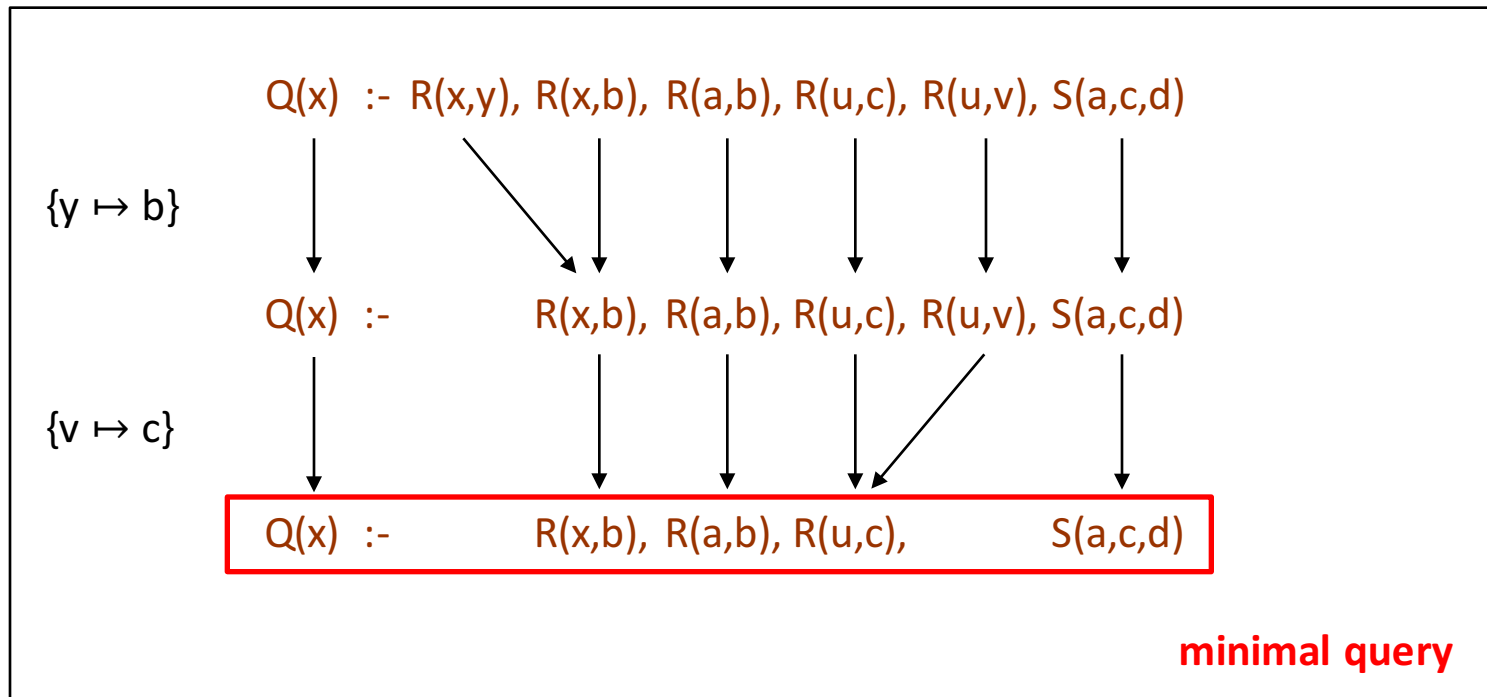
Theorem: It holds that:

- $BQE(\mathbf{CQ})$ is NP-complete (**combined complexity**)
- $BQE[\mathbf{D}](\mathbf{CQ})$ is NP-complete, for a fixed database \mathbf{D} (**query complexity**)
- $BQE[\mathbf{Q}](\mathbf{CQ})$ is in LOGSPACE, for a fixed query $\mathbf{Q} \in \mathbf{CQ}$ (**data complexity**)

Evaluating a CQ \mathbf{Q} over a database \mathbf{D} takes time $|\mathbf{D}|^{O(|\mathbf{Q}|)}$

Minimizing Conjunctive Queries

- Database theory has developed principled methods for optimizing CQs:
 - Find an equivalent CQ with minimal number of atoms (**the core**)
 - Provides a notion of “true” optimality




Minimizing Conjunctive Queries

- But, a minimal equivalent CQ might not be easier to evaluate - remains NP-hard
- “Good” classes of CQs for which query evaluation is tractable (*in combined complexity*):
 - Graph-based
 - Hypergraph-based

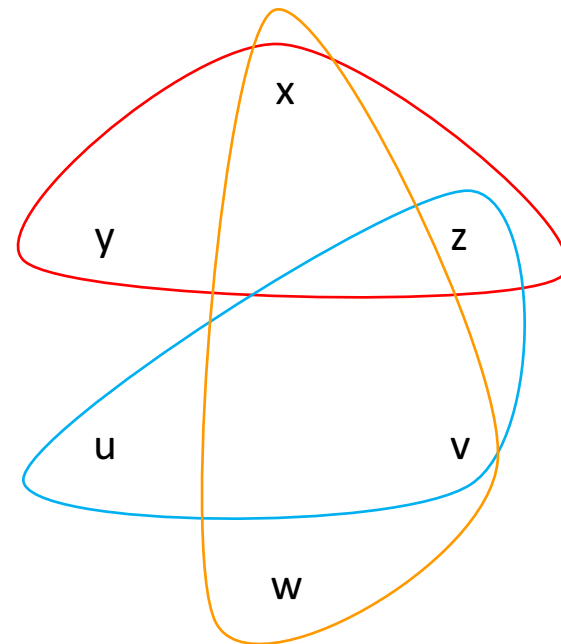
(Hyper)graph of Conjunctive Queries

$Q \text{ :- } R(x,y,z), R(z,u,v), R(v,w,x)$

graph of Q - $G(Q)$



hypergraph of Q - $H(Q)$



“Good” Classes of Conjunctive Queries

measures how close a graph is to a tree



- Graph-based
 - CQs of bounded **treewidth** - their graph has bounded treewidth

measures how close a hypergraph is to an acyclic one

- Hypergraph-based:
 - CQs of bounded **hypertree width** - their hypergraph has bounded hypertree width
 - **Acyclic CQs** - their hypergraph has **hypertree width 1**



Acyclic Hypergraphs

- A **join tree** of a hypergraph $\mathbf{H} = (V, E)$ is a labeled tree $\mathbf{T} = (N, F, L)$, where $L : N \rightarrow E$ such that:
 1. For each hyperedge $e \in E$ of \mathbf{H} , there exists $n \in N$ such that $e = L(n)$
 2. For each node $u \in V$ of \mathbf{H} , the set $\{n \in N \mid u \in L(n)\}$ induces a connected subtree of \mathbf{T}



Acyclic Hypergraphs

- A **join tree** of a hypergraph $\mathbf{H} = (V, E)$ is a labeled tree $\mathbf{T} = (N, F, L)$, where $L : N \rightarrow E$ such that:
 1. For each hyperedge $e \in E$ of \mathbf{H} , there exists $n \in N$ such that $e = L(n)$
 2. For each node $u \in V$ of \mathbf{H} , the set $\{n \in N \mid u \in \lambda(n)\}$ induces a *connected* subtree of \mathbf{T}




condition 2 is violated

Acyclic Hypergraphs

- A **join tree** of a hypergraph $\mathbf{H} = (V, E)$ is a labeled tree $\mathbf{T} = (N, F, L)$, where $L : N \rightarrow E$ such that:
 1. For each hyperedge $e \in E$ of \mathbf{H} , there exists $n \in N$ such that $e = L(n)$
 2. For each node $u \in V$ of \mathbf{H} , the set $\{n \in N \mid u \in \lambda(n)\}$ induces a *connected* subtree of \mathbf{T}

- **Definition:** A hypergraph is **acyclic** if it has a join tree

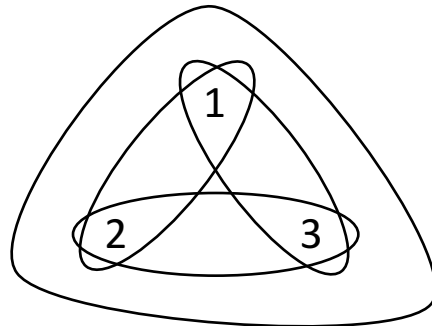


prime example of a cyclic hypergraph

Acyclic Hypergraphs

- A **join tree** of a hypergraph $\mathbf{H} = (V, E)$ is a labeled tree $\mathbf{T} = (N, F, L)$, where $L : N \rightarrow E$ such that:
 1. For each hyperedge $e \in E$ of \mathbf{H} , there exists $n \in N$ such that $e = L(n)$
 2. For each node $u \in V$ of \mathbf{H} , the set $\{n \in N \mid u \in \lambda(n)\}$ induces a *connected* subtree of \mathbf{T}

- **Definition:** A hypergraph is **acyclic** if it has a join tree



but this is acyclic

Relevant Algorithmic Tasks

ACYCLICITY

Input: a query $Q \in \mathbf{CQ}$

Question: is Q acyclic? or is $H(Q)$ acyclic?

$\{Q \in \mathbf{CQ} \mid H(Q) \text{ is acyclic}\}$

BQE(**ACQ**)


Input: a database D , a Boolean query $Q \in \mathbf{ACQ}$

Question: is $Q(D)$ non-empty?

Checking Acyclicity

Via the **GYO-reduction** (Graham, Yu and Ozsoyoglu)


1. Eliminate nodes occurring in at most one hyperedge
2. Eliminate hyperedges that are empty or contained in other hyperedges



Checking Acyclicity

Via the **GYO-reduction** (Graham, Yu and Ozsoyoglu)


1. Eliminate nodes occurring in at most one hyperedge
2. Eliminate hyperedges that are empty or contained in other hyperedges



Checking Acyclicity

Via the **GYO-reduction** (Graham, Yu and Ozsoyoglu)


1. Eliminate nodes occurring in at most one hyperedge
2. Eliminate hyperedges that are empty or contained in other hyperedges



Checking Acyclicity

Via the **GYO-reduction** (Graham, Yu and Ozsoyoglu)

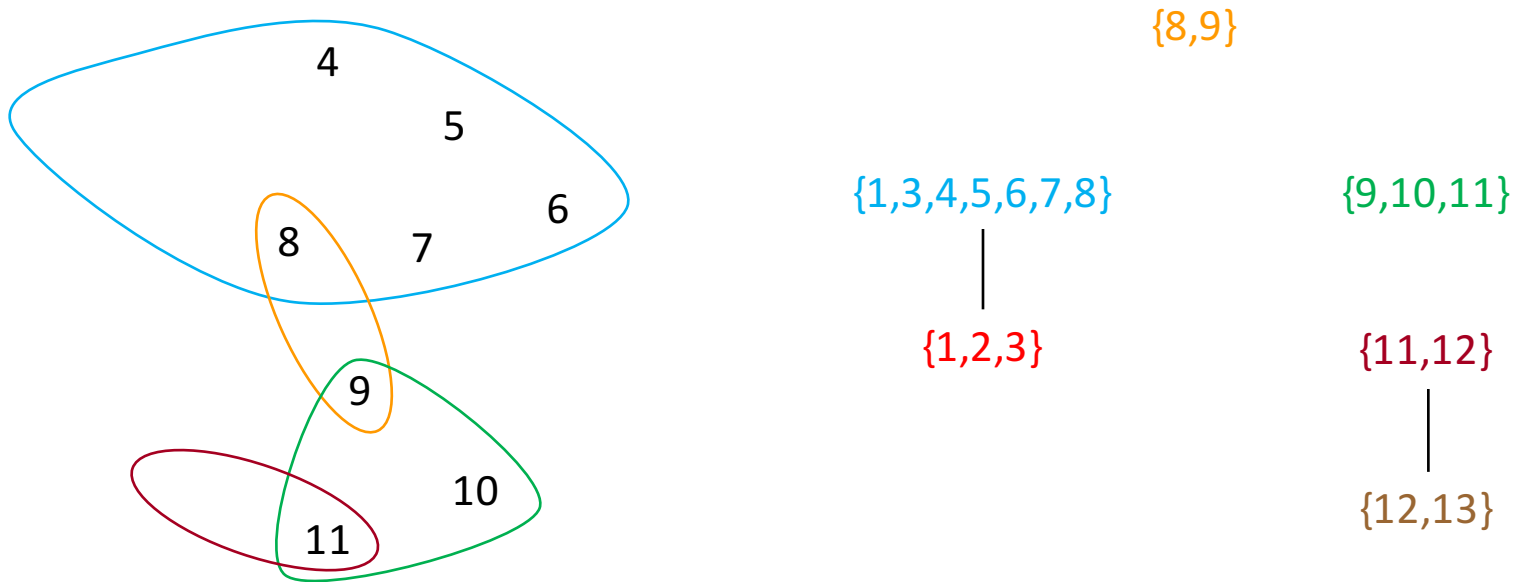
1. Eliminate nodes occurring in at most one hyperedge
2. Eliminate hyperedges that are empty or contained in other hyperedges



Checking Acyclicity

Via the **GYO-reduction** (Graham, Yu and Ozsoyoglu)

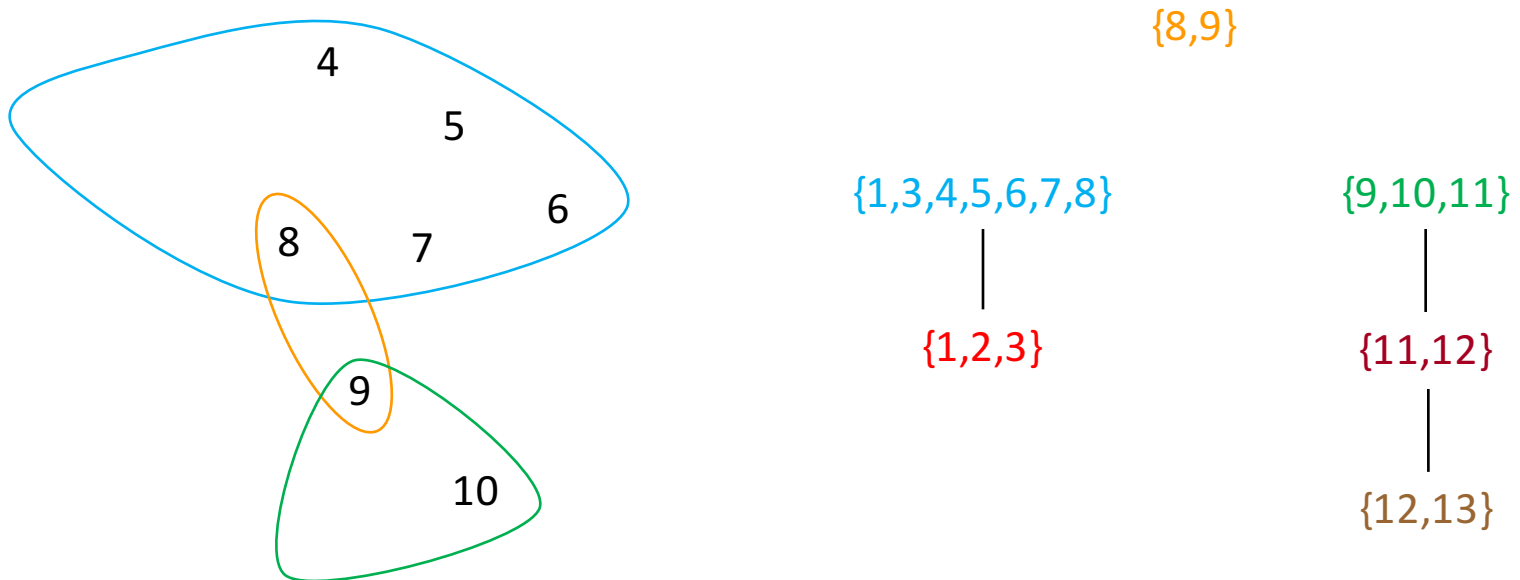
1. Eliminate nodes occurring in at most one hyperedge
2. Eliminate hyperedges that are empty or contained in other hyperedges



Checking Acyclicity

Via the **GYO-reduction** (Graham, Yu and Ozsoyoglu)


1. Eliminate nodes occurring in at most one hyperedge
2. Eliminate hyperedges that are empty or contained in other hyperedges



Checking Acyclicity

Via the **GYO-reduction** (Graham, Yu and Ozsoyoglu)



1. Eliminate nodes occurring in at most one hyperedge
2. Eliminate hyperedges that are empty or contained in other hyperedges



Checking Acyclicity

Via the **GYO-reduction** (Graham, Yu and Ozsoyoglu)



1. Eliminate nodes occurring in at most one hyperedge
2. Eliminate hyperedges that are empty or contained in other hyperedges



Checking Acyclicity

Via the **GYO-reduction** (Graham, Yu and Ozsoyoglu)



1. Eliminate nodes occurring in at most one hyperedge
2. Eliminate hyperedges that are empty or contained in other hyperedges



Checking Acyclicity

Via the **GYO-reduction** (Graham, Yu and Ozsoyoglu)


1. Eliminate nodes occurring in at most one hyperedge
2. Eliminate hyperedges that are empty or contained in other hyperedges



Checking Acyclicity

Via the **GYO-reduction** (Graham, Yu and Ozsoyoglu)

1. Eliminate nodes occurring in at most one hyperedge
2. Eliminate hyperedges that are empty or contained in other hyperedges



Checking Acyclicity

Via the **GYO-reduction** (Graham, Yu and Ozsoyoglu)

1. Eliminate nodes occurring in at most one hyperedge
2. Eliminate hyperedges that are empty or contained in other hyperedges

Theorem: A hypergraph \mathbf{H} is acyclic iff $\text{GYO}(\mathbf{H}) = \emptyset$



checking whether \mathbf{H} is acyclic is feasible in polynomial time, and if it is the case, a join tree can be found in polynomial time



Theorem: ACYCLICITY is in PTIME

Checking Acyclicity

Theorem: ACYCLICITY is in PTIME

NOTE: actually, we can check whether a CQ is acyclic in time $O(|Q|)$

linear time in the size Q

Evaluating Acyclic CQs

Theorem: BQE(ACQ) is in PTIME

NOTE: actually, if $H(Q)$ is acyclic, then Q can be evaluated in time $O(|D| \cdot |Q|)$

linear time in the size of D and Q

Yannakaki's Algorithm


Dynamic programming algorithm over the join tree

Given a database D , and an acyclic Boolean CQ Q

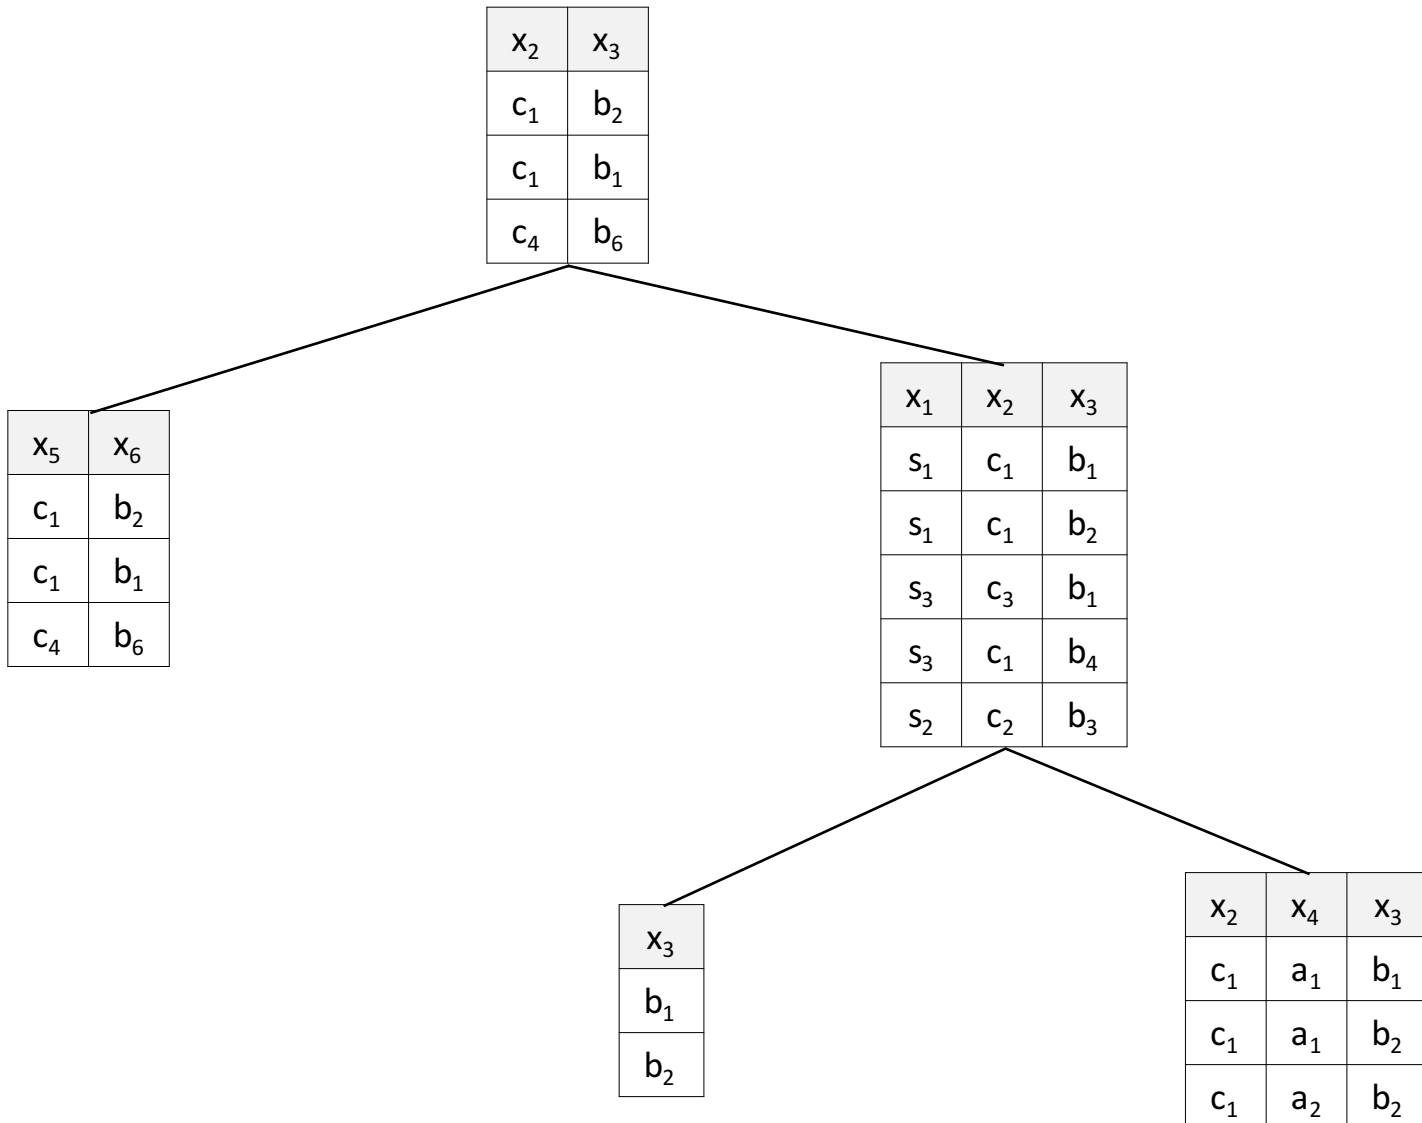
1. Compute the join tree T of $H(Q)$
2. Assign to each node of T the corresponding relation of D
3. Compute semi-joins in a bottom up traversal of T
4. Return YES if the resulting relation at the root of T is non-empty;
otherwise, return NO

Yannakaki's Algorithm: Step 1

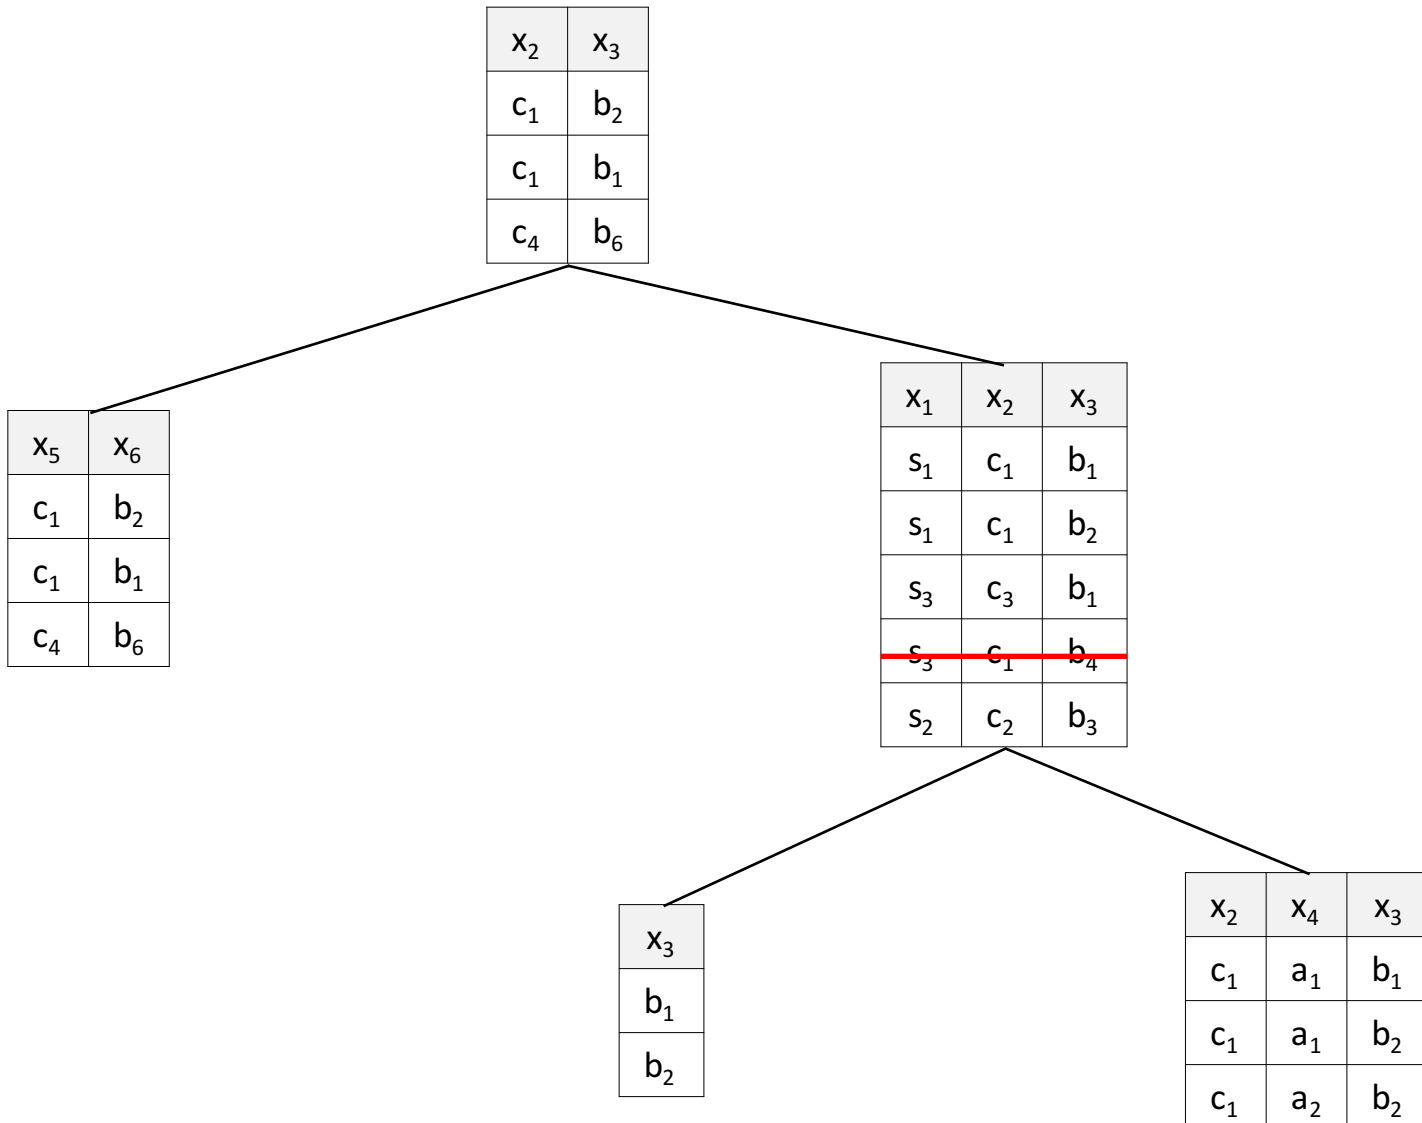
$Q := R_1(x_1, x_2, x_3), R_2(x_2, x_3), R_2(x_5, x_6), R_3(x_3), R_4(x_2, x_4, x_3)$



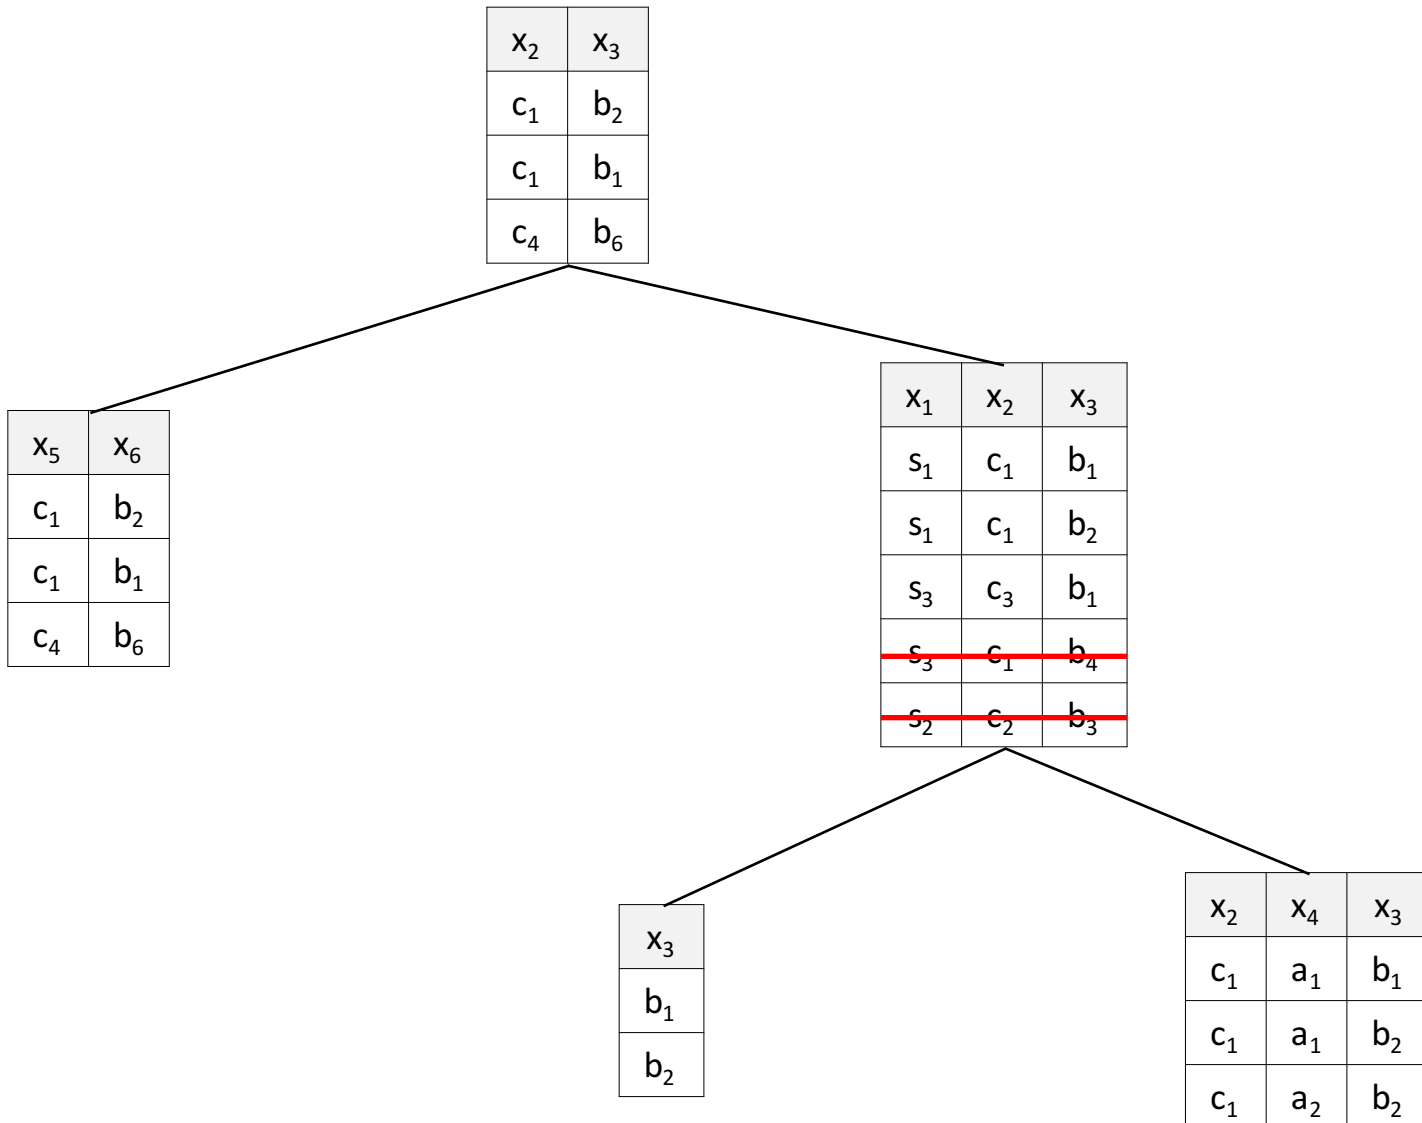
Yannakaki's Algorithm: Step 2



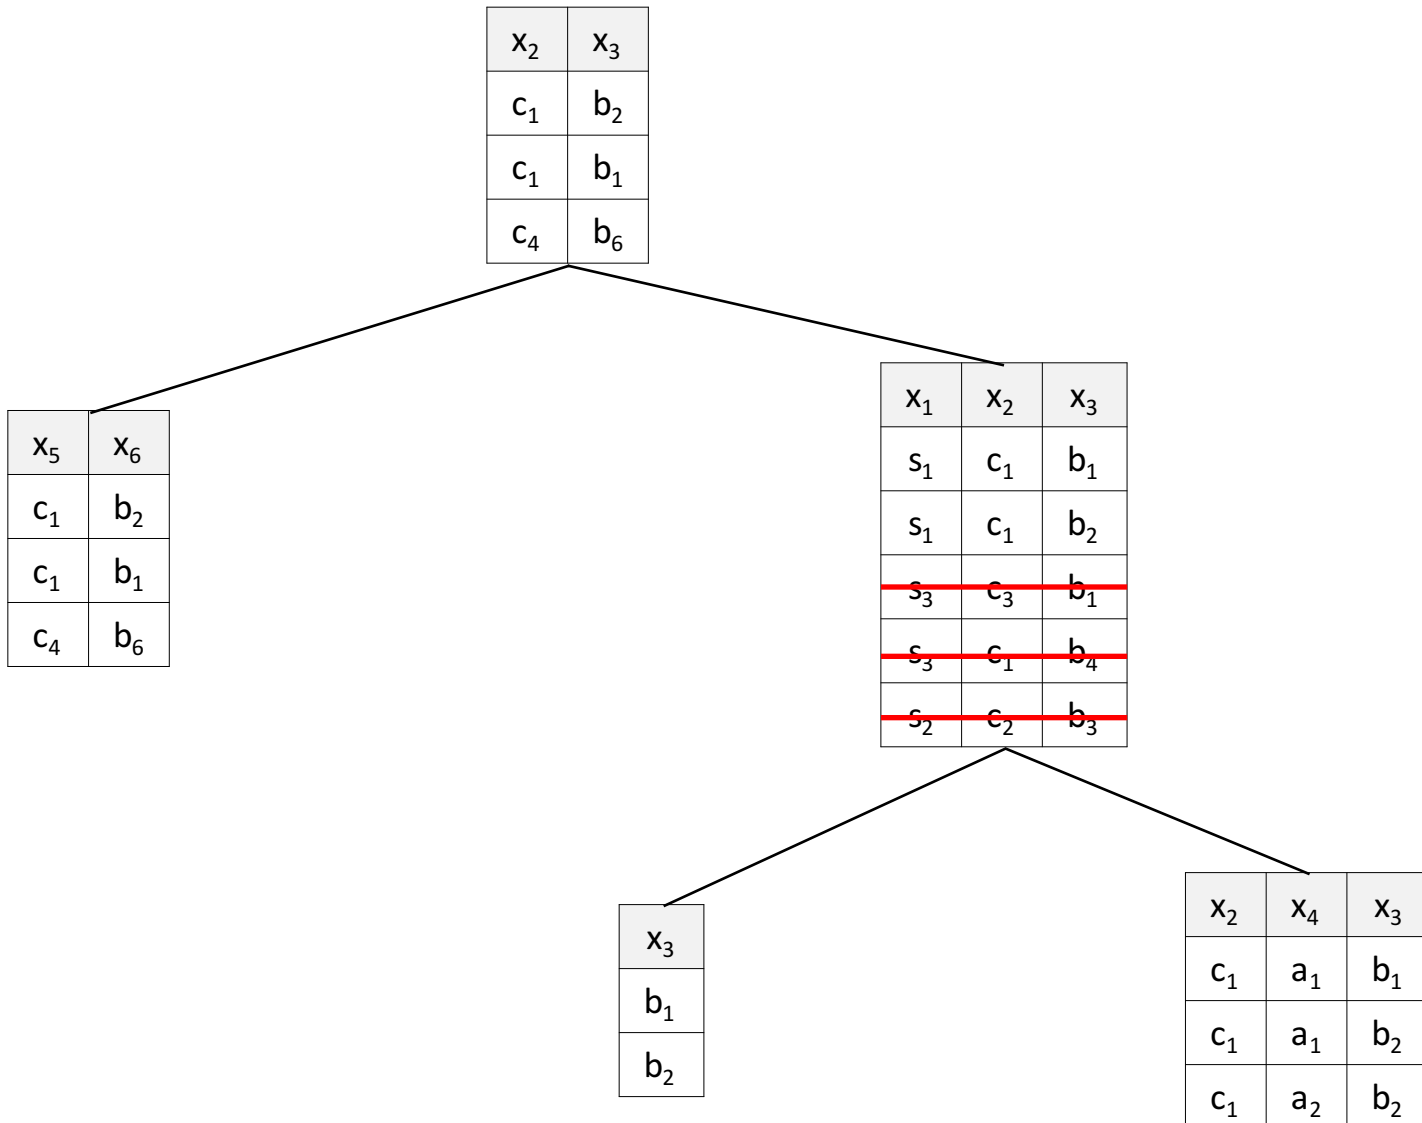
Yannakaki's Algorithm: Step 3



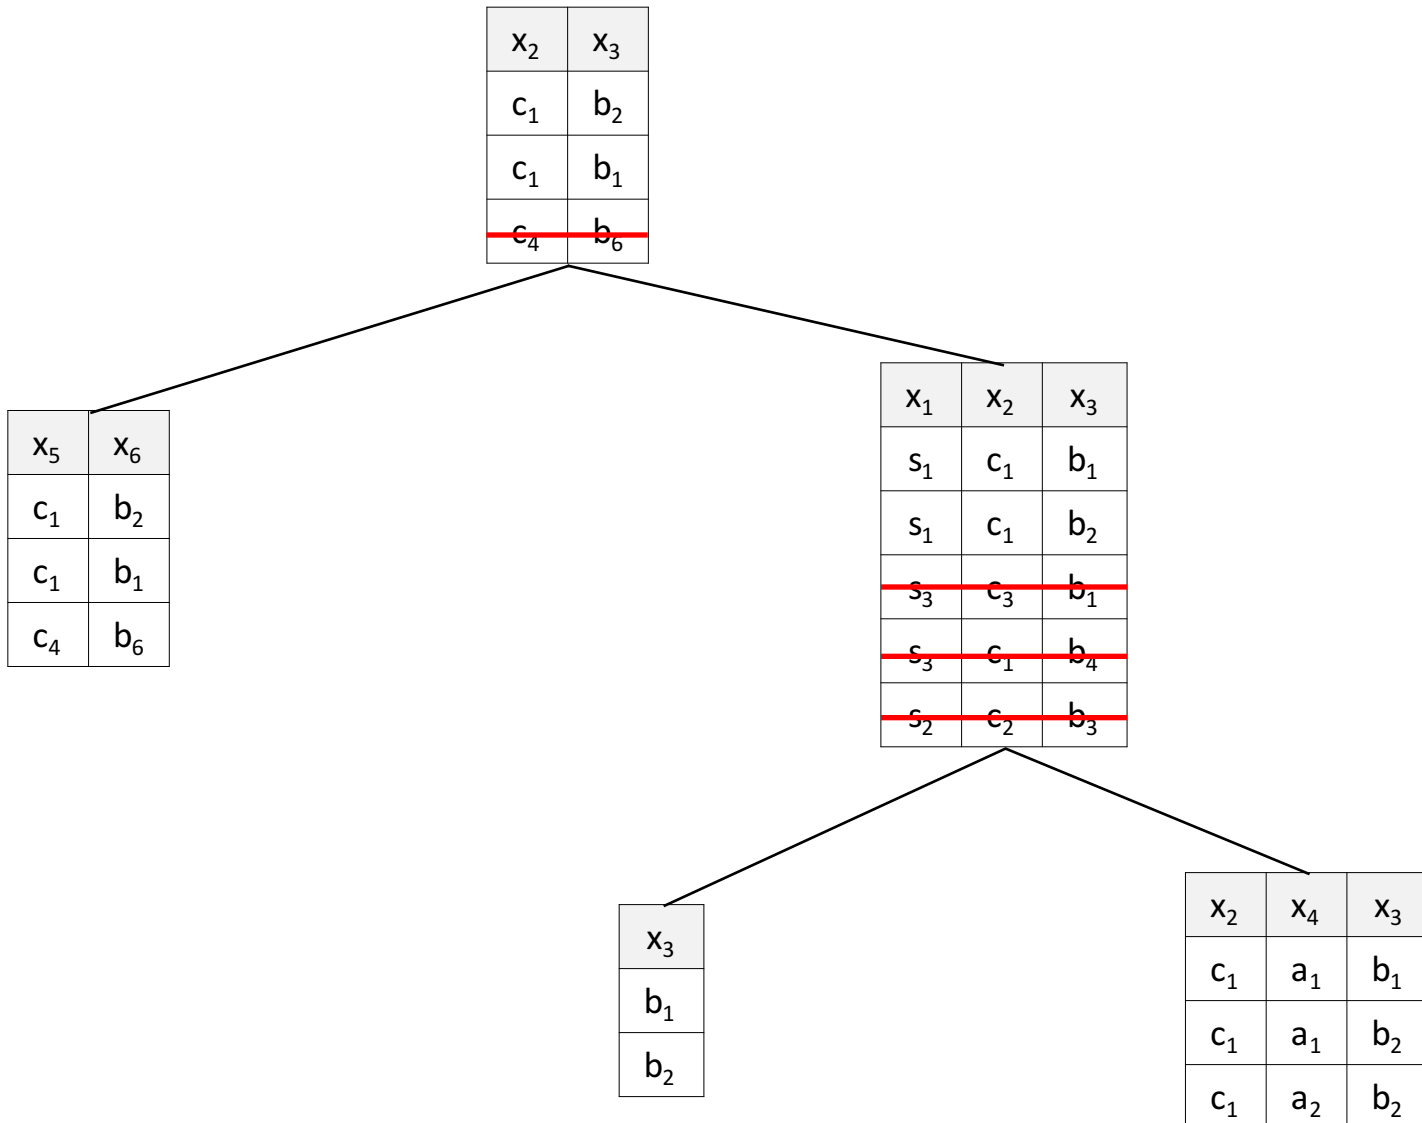
Yannakaki's Algorithm: Step 3



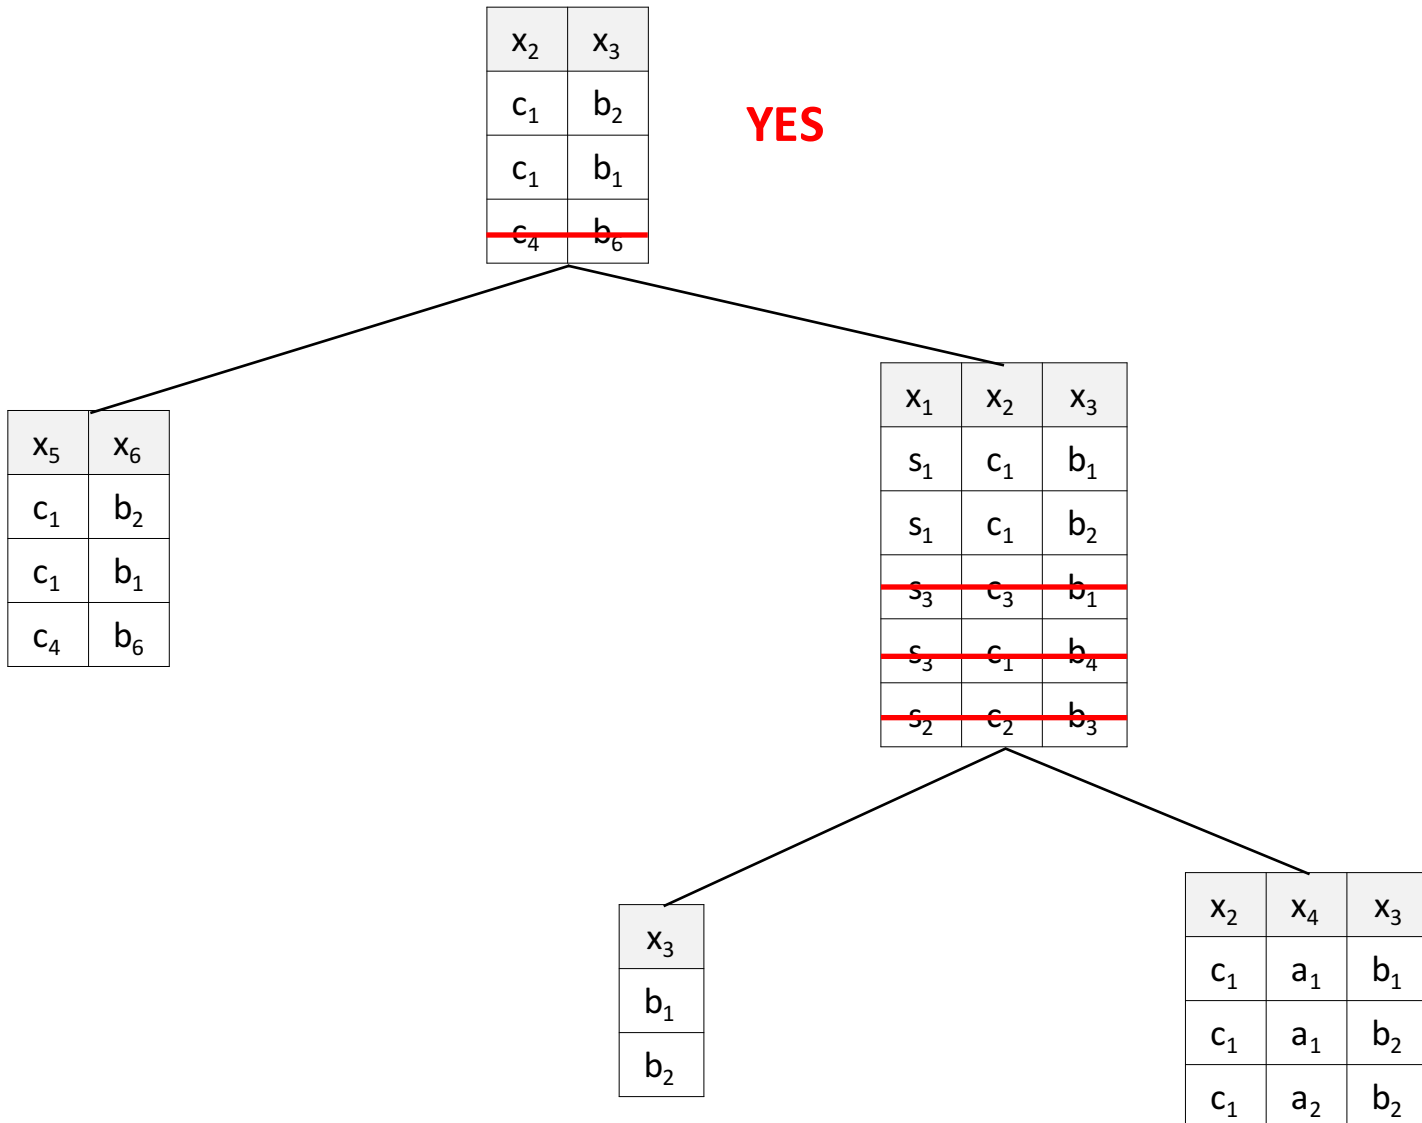
Yannakaki's Algorithm: Step 3



Yannakaki's Algorithm: Step 3



Yannakaki's Algorithm: Step 4



Acyclic CQs: Recap

ACYCLICITY

Input: a query $Q \in \mathbf{CQ}$

Question: is Q acyclic? or is $H(Q)$ acyclic?

BQE(ACQ)

Input: a database D , a Boolean query $Q \in \mathbf{ACQ}$

Question: is $Q(D)$ non-empty?

both problems are feasible in linear time

Query Optimization

Replace a given CQ with one that is much faster to execute

or

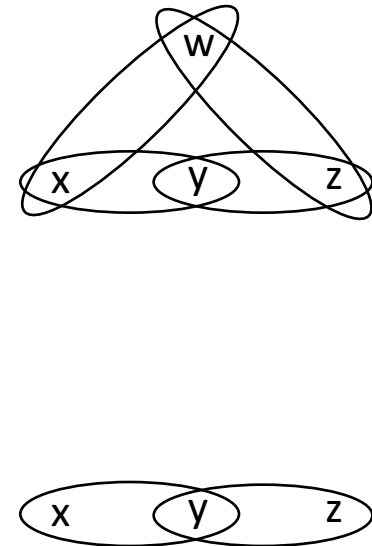
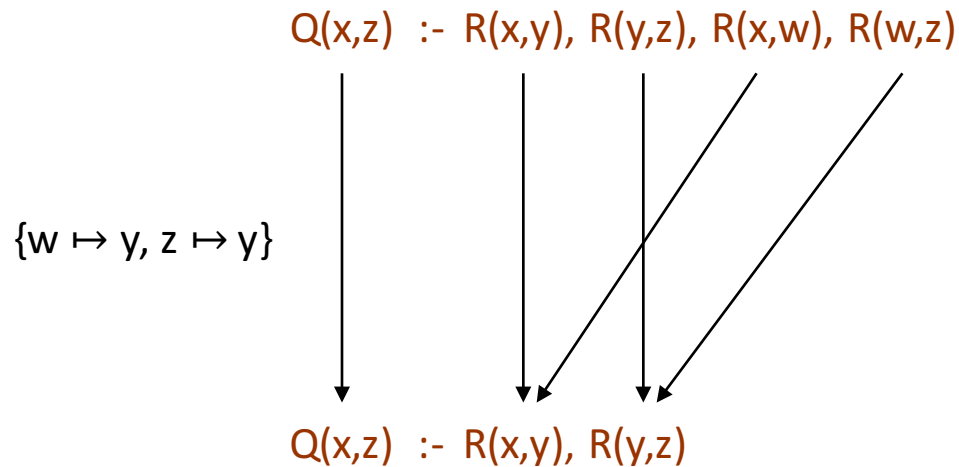
Replace a given CQ with one that falls in a “good” class of CQs



preferably, with an acyclic CQ
since evaluation is in linear time

Semantic Acyclicity

Definition: A CQ Q is **semantically acyclic** if there exists an acyclic CQ Q' such that $Q \equiv Q'$



Relevant Algorithmic Tasks

SemACYCLICITY

Input: a query $Q \in \mathbf{CQ}$

Question: is there an acyclic CQ Q' such that $Q \equiv Q'$?

$\{Q \in \mathbf{CQ} \mid Q \text{ semantically acyclic}\}$



BQE(**SACQ**)

Input: a database D , a Boolean query $Q \in \mathbf{SACQ}$

Question: is $Q(D)$ non-empty?

Checking Semantic Acyclicity


Theorem: A CQ Q is semantically acyclic iff its core is acyclic

Theorem: SemACYCLICITY is NP-complete

Proof idea (upper bound):

- If Q is semantically acyclic, then there exists an acyclic CQ Q' such that $|Q'| \leq |Q|$ and $Q \equiv Q'$ (why?)
- Then, we can guess in polynomial time:
 - An acyclic CQ Q' such that $|Q'| \leq |Q|$
 - A mapping $h_1 : \text{terms}(Q) \rightarrow \text{terms}(Q')$
 - A mapping $h_2 : \text{terms}(Q') \rightarrow \text{terms}(Q)$
- And verify in polynomial time that h_1 is a query homomorphism from Q to Q' (i.e., $Q' \subseteq Q$), and h_2 is a query homomorphism from Q' to Q (i.e., $Q \subseteq Q'$)

Evaluating Semantically Acyclic CQs



an improvement compare to $|D|^{o(|Q|)}$ for evaluating arbitrary CQs

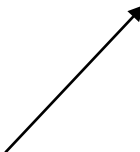
Theorem: BQE(SACQ) is fixed-parameter tractable

Evaluating Semantically Acyclic CQs

Theorem: $\text{BQE}(\text{SACQ})$ is in PTIME

assuming Q belongs to **SACQ**: $Q(D)$ is non-empty $\Leftrightarrow Q \rightarrow_{\exists 1C} D$

the duplicator has a winning strategy
for the existential 1-cover game,
which can be checked in polynomial time



Semantically Acyclic CQs: Recap

SemACYCLICITY

Input: a query $Q \in \mathbf{CQ}$

Question: is there an acyclic CQ Q' such that $Q \equiv Q'$?

NP-complete - but no database is involved

BQE(**SACQ**)

Input: a database D , a Boolean query $Q \in \mathbf{SACQ}$

Question: is $Q(D)$ non-empty?

in PTIME (combined complexity)

Recap

- “Good” classes of CQs for which query evaluation is tractable - conditions based on the graph or hypergraph of the CQ
- Acyclic CQs - their hypergraph is acyclic, can be checked in linear time
- Evaluating acyclic CQs is feasible in linear time (Yannakaki’s algorithm)
- Semantic acyclicity - difficult to check, but ensures tractable evaluation