

## Introduction

**Big Data:** Data too large to be processed on one server

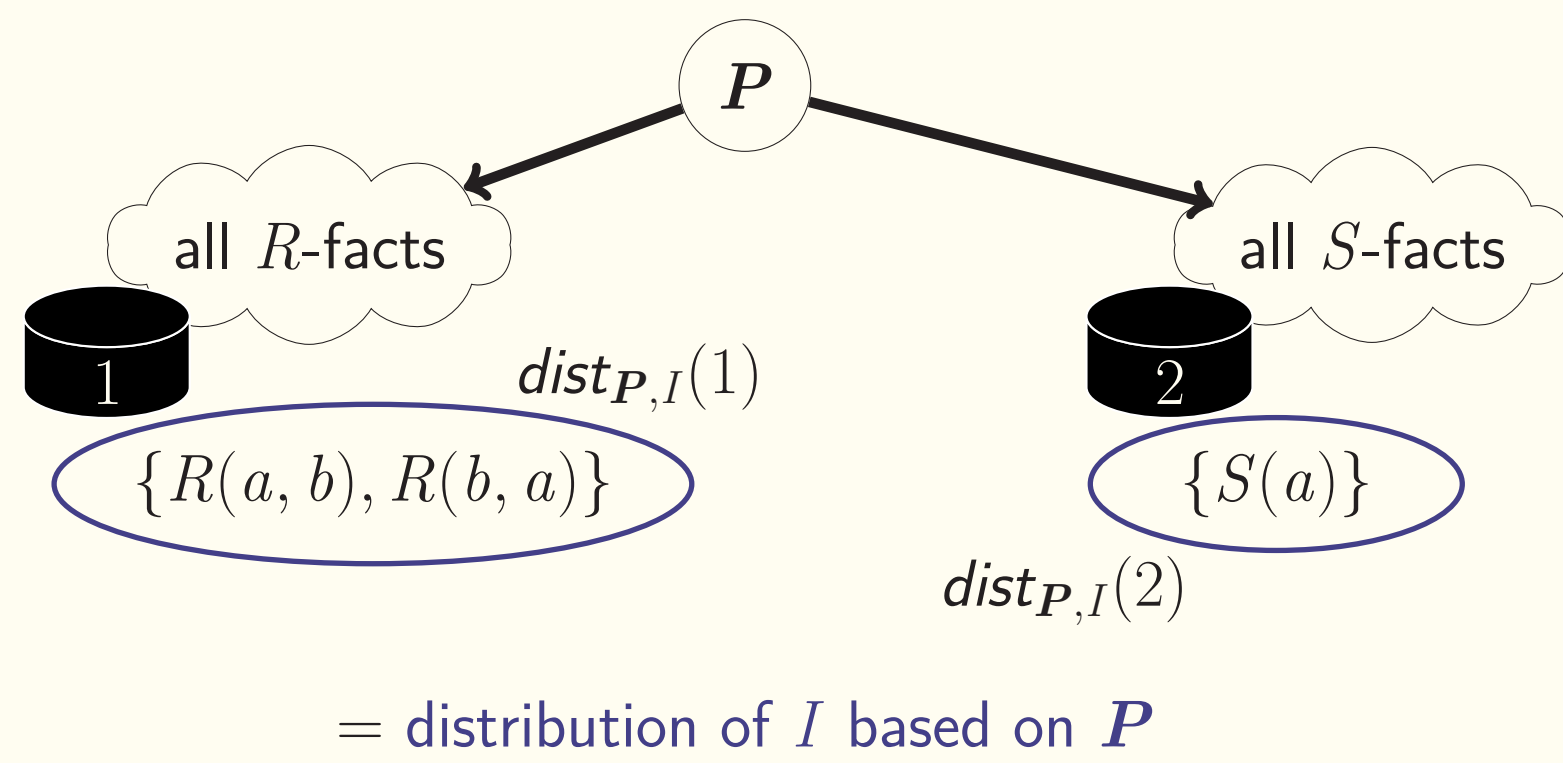
**Today systems:** Hadoop, Spark, ..., and many others

**Common Strategy:**

- ▶ Data is stored in a distributed way
- ▶ Query evaluation: Multiple rounds with reshuffling

## Distribution Policies

Network  $\mathcal{N}$  is a finite set of nodes  
 Instance  $I = \{R(a, b), R(b, a), S(a)\}$



### Definition

A distribution policy  $P$  is a total function mapping facts (over **dom**) to sets of nodes in  $\mathcal{N}$

## Simple Evaluation Algorithm

### 1-Round MPC model

Input = query  $Q$

Step 1:

Step 2:

Output = union of output at each server

## When is the simple algorithm correct on a distribution policy?

## Parallel-Correctness

### Definition

$Q$  is parallel-correct on  $I$  w.r.t.  $P$ , iff

$$Q(I) = \bigcup_{\kappa \in \mathcal{N}} Q(\text{dist}_{P,I}(\kappa))$$

⊇ by monotonicity

### Definition (w.r.t. all instances)

$Q$  is parallel-correct w.r.t.  $P$  iff  
 $Q$  is parallel-correct w.r.t.  $P$  on every  $I$

## Sufficient Condition

(C0) for every valuation  $V$  for  $Q$ ,


$$\bigcap_{f \in V(\text{body}_Q)} P(f) \neq \emptyset.$$

**Intuition:** Facts required by a valuation meet at some node

### Lemma

(C0) implies  $Q$  parallel-correct w.r.t.  $P$ .

### Distribution policy $P$



Query  $Q: T(x, z) \leftarrow R(x, y), R(y, z), R(x, x)$

$V = \{x, z \rightarrow a, y \rightarrow b\}$        $V' = \{x, y, z \rightarrow a\}$

Requires:

$R(a, b) \ R(b, a) \ R(a, a)$        $R(a, a)$

Derives:

$T(a, a)$        $T(a, a)$

Do not meet

**Notice:**  $Q$  is minimal CQ

▶ CQ is minimal iff injective valuations are minimal

### Proposition

Testing whether a valuation is minimal is coNP-complete.

## Characterization

### Lemma

$Q$  is parallel-correct w.r.t.  $P$  iff

(C1) for every minimal valuation  $V$  for  $Q$ ,

$$\bigcap_{f \in V(\text{body}_Q)} P(f) \neq \emptyset.$$

### Definition

$V$  is minimal if no  $V'$  exists, where

$V'(\text{head}_Q) = V(\text{head}_Q)$ ,  $V'(\text{body}_Q) \subsetneq V(\text{body}_Q)$ .

## Complexity

### Theorem

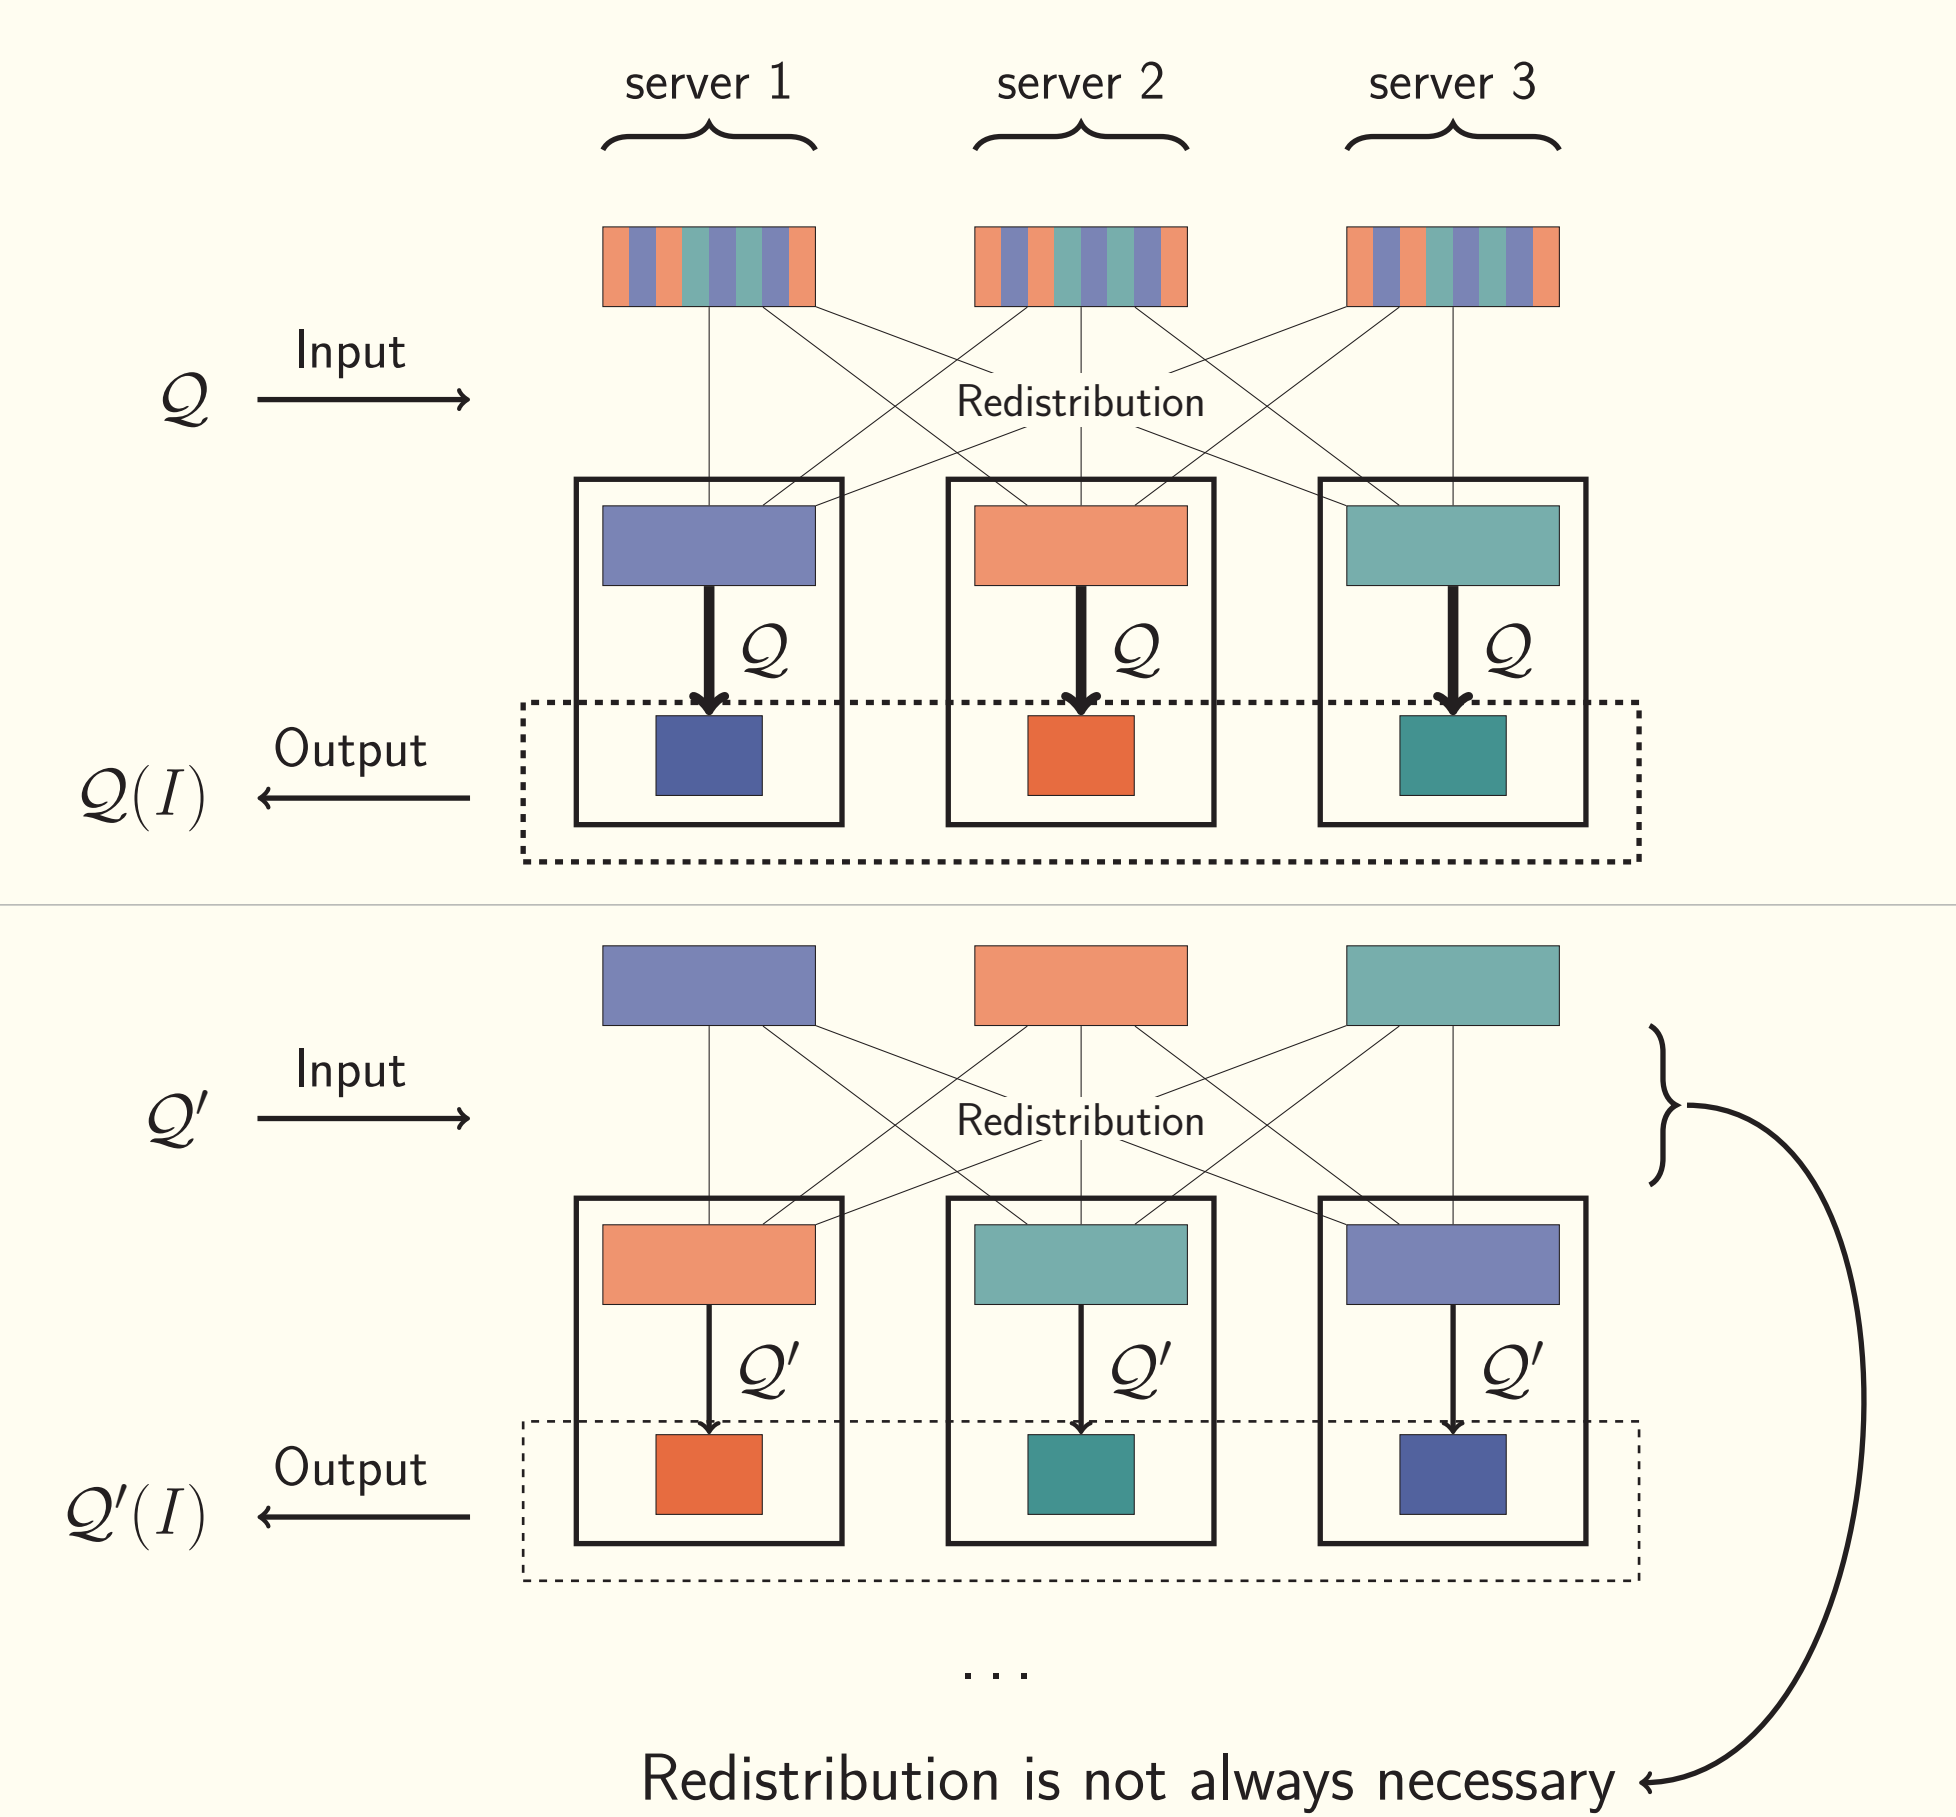
Deciding whether  $Q$  is parallel-correct w.r.t.  $P$  is  $\Pi_2^P$ -complete.

### Proof:

- ▶ Lower bound: Reduction from  $\Pi_2$ -QBF
- ▶ Upper bound: Characterization  
 but, requires proper formalization of  $P$

## Multi-Query Optimization

Computing a set of queries:  $\{Q, Q', \dots\}$



## Which queries allow to reuse the distribution obtained for another query?

## Transferability



### Definition

$Q \rightarrow_T Q'$  iff  $Q'$  is parallel-correct on every  $P$  where  $Q$  is parallel-correct on


### Example

$Q: T() \leftarrow R(x, y), R(y, z), R(z, w)$        $Q': N() \leftarrow R(x, y), R(y, x)$

Minimal valuations for  $Q$ :



Minimal valuations for  $Q'$ :



Facts required by minimal valuations for  $Q'$  are also required by minimal valuations for  $Q$

$Q \rightarrow_T Q'$

## Characterization

### Lemma

$Q \rightarrow_T Q'$  iff

(C2) for every minimal valuation  $V'$  for  $Q'$  there is a minimal valuation  $V$  for  $Q$ , s.t.  $V'(\text{body}_{Q'}) \subseteq V(\text{body}_Q)$ .

Based on query structure alone, not on distribution policies

## Complexity

### Theorem

Deciding  $Q \rightarrow_T Q'$  is  $\Pi_3^P$ -complete.

### Proof:

- ▶ Lower bound: Reduction from  $\Pi_3$ -QBF
- ▶ Upper bound: Characterization

## Strongly Minimal CQs

### Definition

A CQ is strongly minimal if all its valuations are minimal

### Full-CQs

$T(x, y) \leftarrow R(x, y), R(x, x)$

### CQs without self-joins

$T() \leftarrow R(x, y), S(x, x)$

### Hybrids

$T(y) \leftarrow R(x, y), R(x, x), R(z, x), S(z)$

A minimal CQ is not always strongly minimal

### Lemma

Deciding whether  $Q$  is strongly minimal is coNP-complete

### Theorem

Deciding  $Q \rightarrow_T Q'$  is NP-complete for strongly minimal  $Q$

## Hypercube

▶ Invented in the context of Datalog evaluation

[Ganguli, Silberschatz & Tsur 1990]

▶ Described in Map-Reduce context


[Afrati & Ullman 2010]

▶ Intensively studied by many people

[Beame, Koutris & Suciu 2014]

### Algorithm sketch:

▶ Reshuffling based on structure of  $Q$



Partitioning of complete valuations over servers in instance independent way through hashing of domain values

Let  $\mathcal{H}(Q)$  be the family of Hypercube distributions for  $Q$ .

### Definition

$Q \rightarrow_H Q'$  iff

$Q'$  is parallel-correct w.r.t. every  $P \in \mathcal{H}(Q)$ .

### Two properties:

- ▶  $Q$ -generous: for every valuation facts meet on some node ( $\forall P \in \mathcal{H}(Q)$ )
- ▶  $Q$ -scattered: there is a policy scattering facts in such a way that no facts meet by coincidence ( $\forall I$ )

### Theorem

Deciding whether  $Q \rightarrow_H Q'$  is NP-complete

(also when  $Q$  or  $Q'$  is acyclic)

## Conclusion & Future Work

### Summary:

Formal framework for reasoning about correctness of query evaluation and optimization in a distributed setting

### Based on two concepts:

- ▶ Parallel-correctness
- ▶ Transferability

Independent of expression mechanism

## Related Concepts

### Containment

$Q \subseteq Q'$

### Lemma

Containment and transferability are incomparable

### Determinacy

(Data-Integration)

$Q'(I) = Q'(J)$  implies  $Q(I) = Q(J)$ , for every  $I, J$

### Lemma

Determinacy and transferability are incomparable

### Transferability

$Q \rightarrow_T Q'$

(Distributed)

## Future Work

### Expression Formalism for distribution policies

▶ Other than Hypercube?

### Distribution policy for set of queries

▶ Given CQ: which distribution policy? **Hypercube**

▶ Given set of CQs: which distribution policy? **Open question**

### Tractable Results

- ▶ Other classes of queries?
- ▶ Other families of distribution policies?

### More expressive classes of queries

- ▶ This work: CQs
- ▶ FO: undecidable
- ▶ initial results: UCQs, CQs with negation