

Overview of distributed query processing

Data Management for Big Data
2018-2019 (spring semester)

Dario Della Monica

These slides are a modified version of the slides provided with the book
Özsu and Valduriez, *Principles of Distributed Database Systems* (3rd Ed.), 2011
The original version of the slides is available at: extras.springer.com

Distributed DBMS

© M. T. Özsu & P. Valduriez

Ch. 6/1

Outline (distributed DB)

- Introduction (Ch. 1) *
- Distributed Database Design (Ch. 3) *
- Distributed Query Processing (Ch. 6-8) *
 - Overview (Ch. 6) *
 - Query decomposition and data localization (Ch. 7) *
 - Distributed query optimization (Ch. 8) *
- Distributed Transaction Management (Ch. 10-12) *

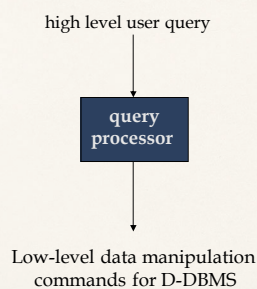
* Özsu and Valduriez, *Principles of Distributed Database Systems* (3rd Ed.), 2011

Distributed DBMS

© M. T. Özsu & P. Valduriez

Ch. 6/2

Query Processing in a D-DBMS



Distributed DBMS

© M. T. Özsu & P. Valduriez

Ch. 6/3

Selecting Alternatives

```

SELECT *
FROM EMP, ASG
WHERE EMP.ENO = ASG.ENO
AND RESP = "Manager"
  
```

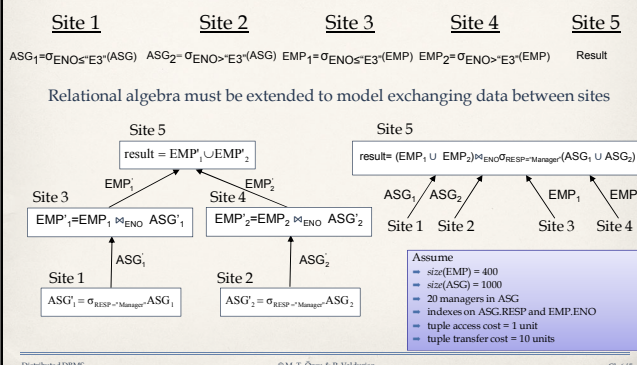
$EMP \bowtie_{ENO} (\sigma_{RESP="Manager"}(ASG))$

Distributed DBMS

© M. T. Özsu & P. Valduriez

Ch. 6/4

What are the Additional Problems?



Distributed DBMS

© M. T. Özsu & P. Valduriez

Ch. 6/5

Cost of Alternatives

- Assume
 - size(EMP) = 400, size(ASG) = 1000, 20 managers in ASG
 - cluster indexes on ASG.RESP and EMP.ENO
 - tuple access cost = 1 unit; tuple transfer cost = 10 units
- Strategy A
 - produce ASG': (10+10) * tuple access cost 20
 - transfer ASG' to the sites of EMP: (10+10) * tuple transfer cost 200
 - produce EMP': (10+10) * tuple access cost * 2 40
 - transfer EMP' to result site: (10+10) * tuple transfer cost 200
 - Total Cost 460**
- Strategy B
 - transfer EMP to site 5: 400 * tuple transfer cost 4,000
 - transfer ASG to site 5: 1000 * tuple transfer cost 10,000
 - produce ASG': 1000 * tuple access cost 1,000
 - join EMP and ASG': 400 * 20 * tuple access cost 8,000
 - Total Cost 23,000**

Distributed DBMS

© M. T. Özsu & P. Valduriez

Ch. 6/6

Query Optimization Objectives in DDBS

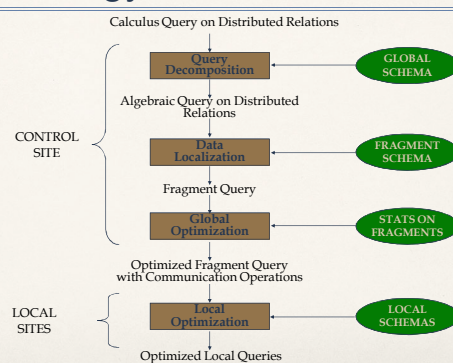
- To transform a query in a high-level language (SQL) on a distributed DB (seen as a single DB by the user) into an efficient execution strategy, expressed in a lower-level language (extension of relational algebra with communication and data transfer operators), on several local DB's
- Cost factors to minimize
 - Centralized DB: CPU and I/O cost only (actually, only I/O)
 - Distributed DB: also communication costs
 - Communication costs are the dominating ones (even though this might not be the case with increased network speed, especially within Local Area Network)

Distributed DBMS

© M. J. Özsu & P. Valduriez

Ch. 6/2

Distributed Query Processing Methodology



Distributed DBMS

© M. J. Özsu & P. Valduriez

Ch. 6/3