

# Distributed Databases


---

by Farnoush Banaei-Kashani

Excerpt from “Principles of Distributed Database Systems”  
by M. Tamer Özsu and Patrick Valduriez


# File Systems

---




# Database Management System

---



# Motivation

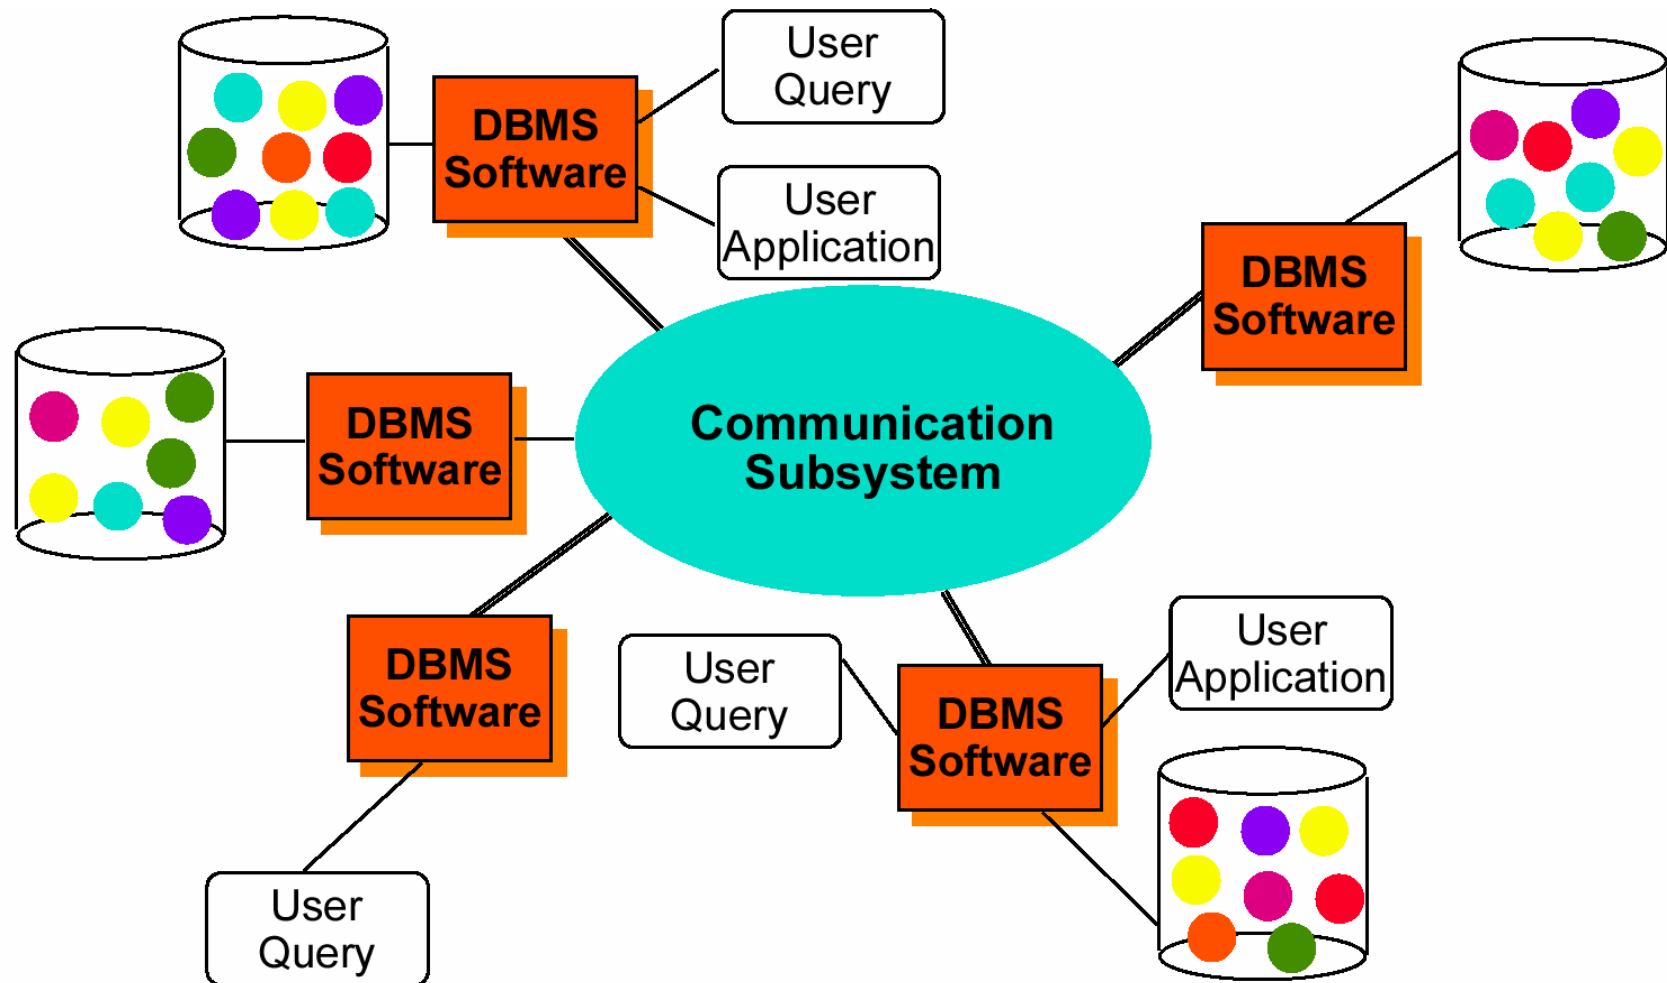
---



**integration  $\neq$  centralization**


# Distributed DBMS – Reality

---



# Distributed Database – User View

---



# Applications – First Generation

---

- Manufacturing - especially multi-plant manufacturing
- Military command and control
- EFT
- Corporate MIS
- Airlines
- Hotel chains
- Any organization which has a decentralized organization structure

# Applications – Second Generation

---

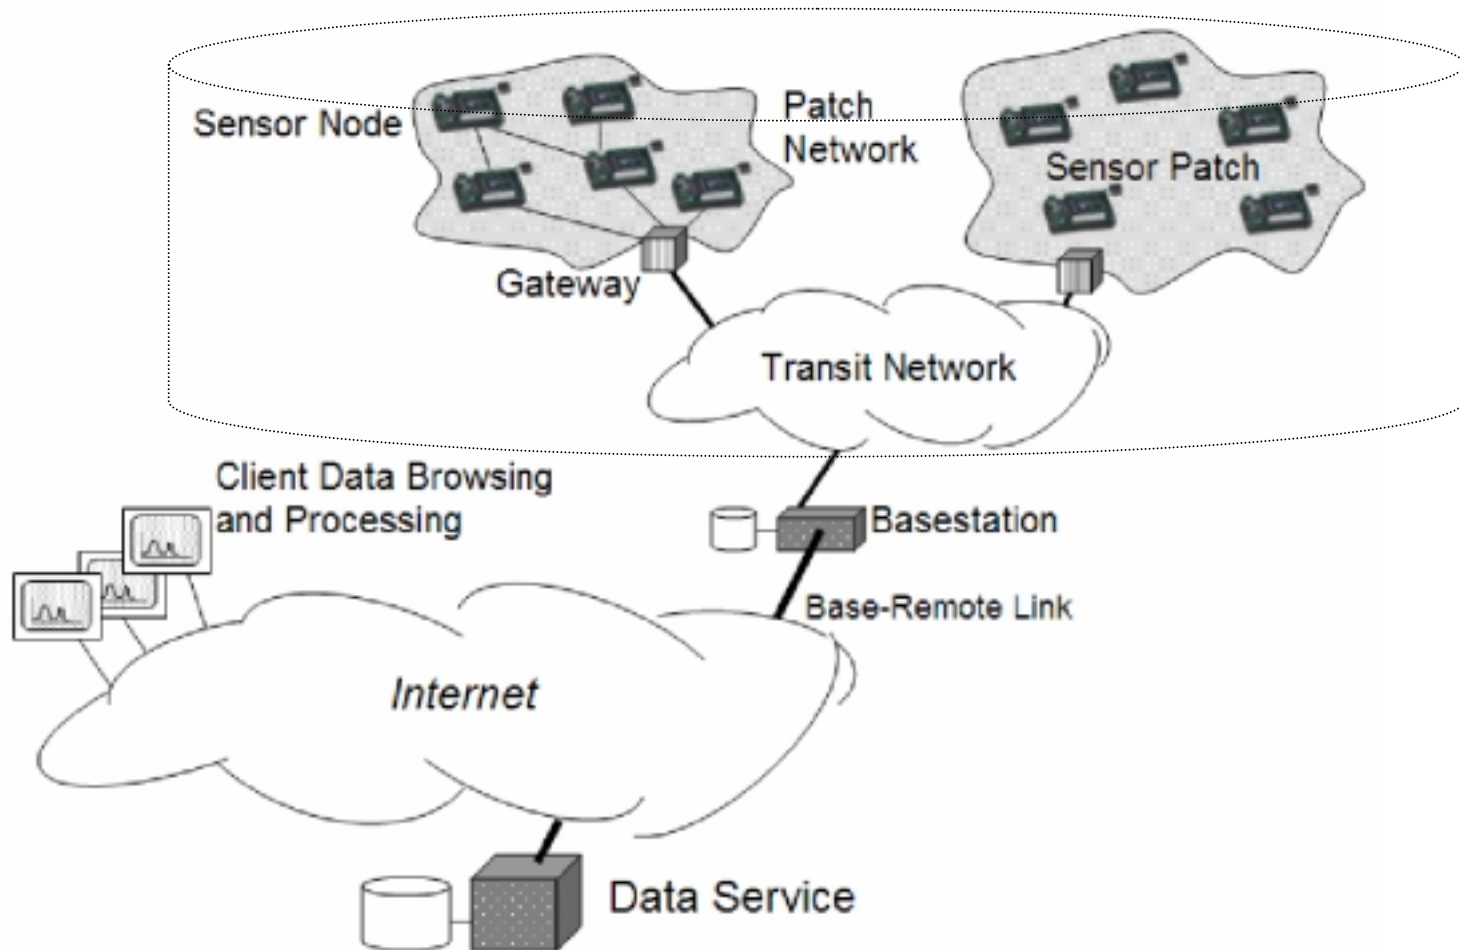
## ■ Sensor Databases





# Applications – Second Generation

---




<http://www.greatduckisland.net/>

# Applications – Second Generation


---

## ■ Peer-to-Peer Databases




# Applications – Second Generation

---



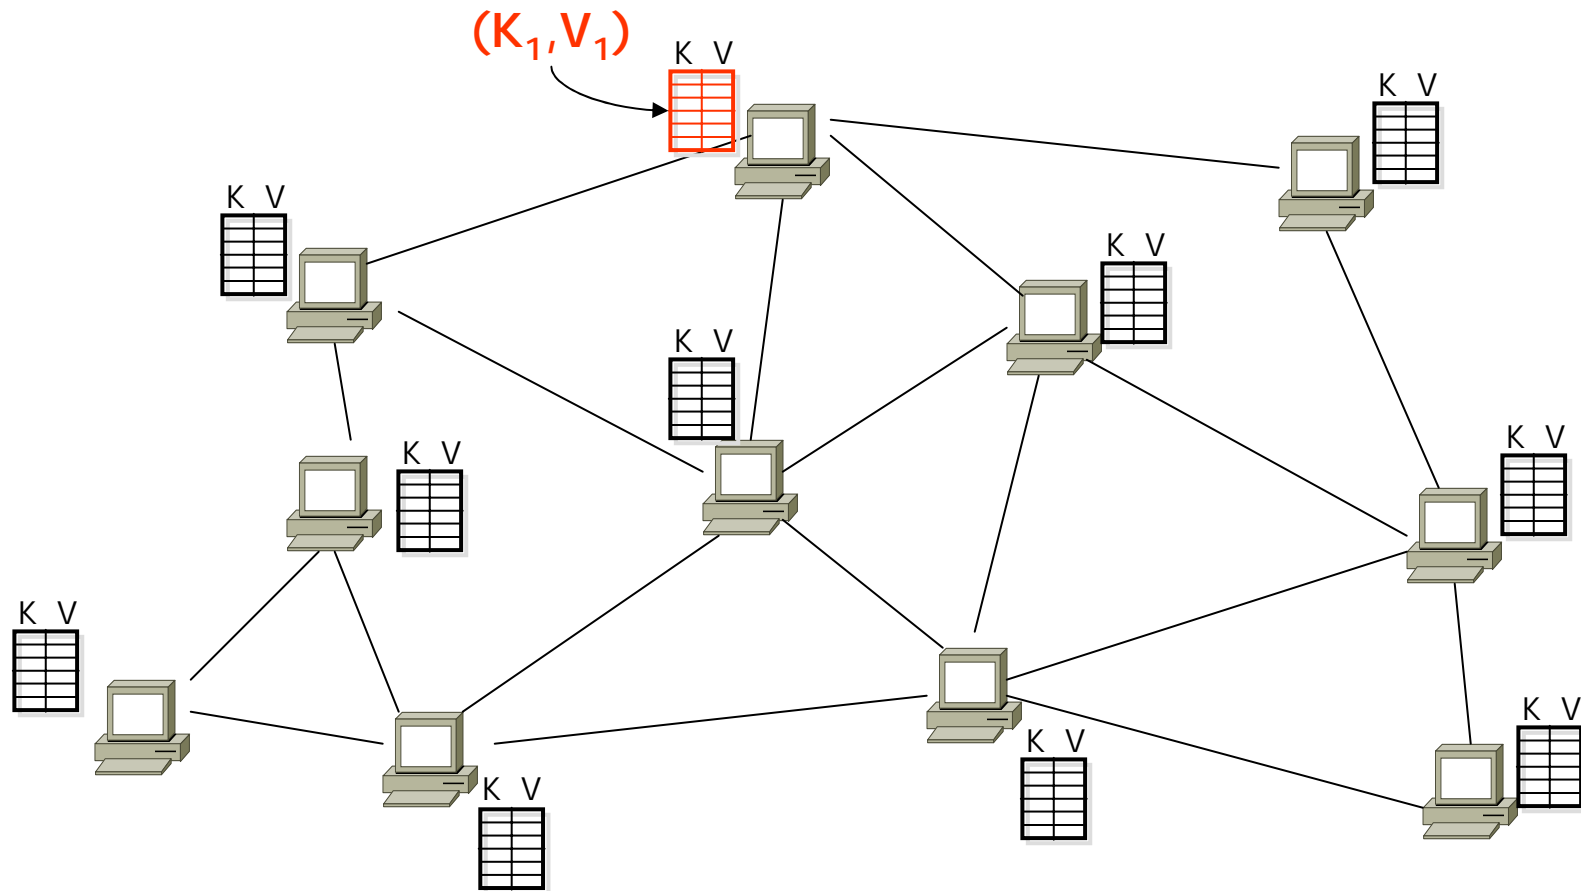
# Applications – Second Generation

---




# Applications – Second Generation

---




# Applications – Second Generation

---




# Applications – Second Generation

---



# Applications – Second Generation

---





# Topics

---

- Introduction
- Background
- Distributed DBMS Architecture
- Distributed Database Design
- Semantic Data Control
- Distributed Query Processing
- Distributed Transaction Management
- Parallel Database Systems
- Distributed Object DBMS
- Database Interoperability
- Current Issues

# Topics

---

- Introduction
- Background
- Distributed DBMS Architecture
- Distributed Database Design
- Semantic Data Control
- Distributed Query Processing
- Distributed Transaction Management
- Parallel Database Systems
- Distributed Object DBMS
- Database Interoperability
- Current Issues

# Outline

---

- ❑ Definition and Differentiation
- ❑ Promises and Challenges
- ❑ Architecture Alternatives

# Distributed Computing

---

- A concept in search of a definition and a name.
- A number of autonomous processing elements (not necessarily homogeneous) that are interconnected by a computer network and that cooperate in performing their assigned tasks.

# What is distributed ...

---

- Processing logic
- Functions
- Data
- Control

# What is Distributed Database System?

---

A distributed database (DDB) is a collection of multiple, *logically interrelated* databases distributed over a *computer network*.

A distributed database management system (D-DBMS) is the software that manages the DDB and provides an access mechanism that makes this distribution *transparent* to the users.

Distributed database system (DDBS) = DDB + D-DBMS

# What is not a DDBS?

---

- Distributed database is a **database, not a collection of files**    data logically related as exhibited in the users' access patterns
  - ▶ relational data model
- D-DBMS is a **full-fledged DBMS**
  - ▶ not remote file system, not a TP system

# What else is not a DDBS?


---

- A database system which resides at one of the nodes of a network of computers - this is a centralized database on a network node




# Centralized DBMS on Network

---



# Distributed DBMS Environment

---




# What else is not a DDBS?

---

- Processors at different sites are interconnected by a computer network    no multiprocessors
  - ▶ parallel database systems


# Shared-Memory Architecture

---




# Shared-Disk Architecture

---



# Shared-Nothing Architecture

---



# Outline

---

- Definition and Differentiation
- Promises and Challenges
- Architecture Alternatives

# Distributed Database Promises

---

- ❶ Transparent management of distributed, fragmented, and replicated data
- ❷ Improved reliability/availability through distributed transactions
- ❸ Improved performance
- ❹ Easier and more economical system expansion



# Example

---

EMP

ENO	ENAME	TITLE
E1	J. Doe	Elect. Eng.
E2	M. Smith	Syst. Anal.
E3	A. Lee	Mech. Eng.
E4	J. Miller	Programmer
E5	B. Casey	Syst. Anal.
E6	L. Chu	Elect. Eng.
E7	R. Davis	Mech. Eng.
E8	J. Jones	Syst. Anal.

ASG

ENO	PNO	RESP	DUR
E1	P1	Manager	12
E2	P1	Analyst	24
E2	P2	Analyst	6
E3	P3	Consultant	10
E3	P4	Engineer	48
E4	P2	Programmer	18
E5	P2	Manager	24
E6	P4	Manager	48
E7	P3	Engineer	36
E7	P5	Engineer	23
E8	P3	Manager	40

PROJ

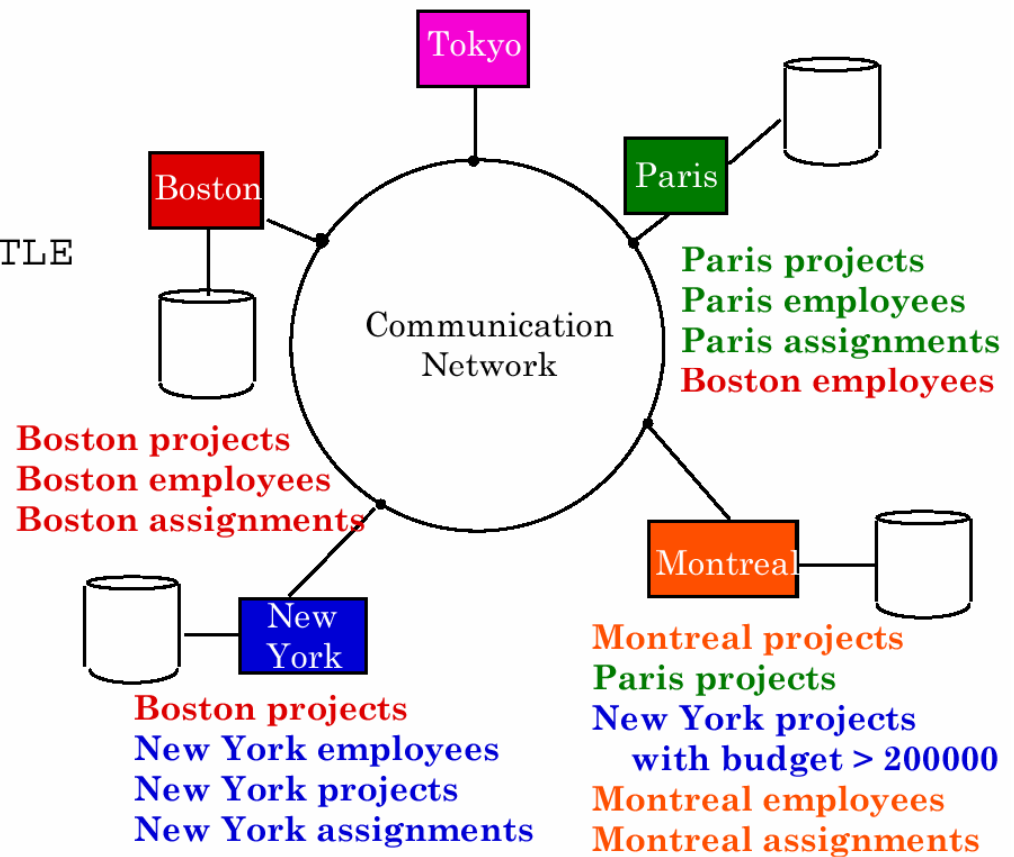
PNO	PNAME	BUDGET
P1	Instrumentation	150000
P2	Database Develop	135000
P3	CAD/CAM	250000
P4	Maintenance	310000

PAY

TITLE	SAL
Elect. Eng.	40000
Syst. Anal.	34000
Mech. Eng.	27000
Programmer	24000

# Example

```
SELECT ENAME, SAL
FROM   EMP, ASG, PAY
WHERE  DUR > 12
AND    EMP.ENO = ASG.ENO
AND    PAY.TITLE = EMP.TITLE
```




# Transparency

---

- Transparency is the separation of the higher level semantics of a system from the lower level implementation issues.
- Fundamental issue is to provide  
**data independence**  
in the distributed environment
  - Network (distribution) transparency
  - Replication transparency
  - Fragmentation transparency
    - ◆ horizontal fragmentation: selection
    - ◆ vertical fragmentation: projection
    - ◆ hybrid


# Distributed DBMS – Reality

---



# Distributed Database – User View

---



# Distributed Database Promises

---

- ✓ ❶ Transparent management of distributed, fragmented, and replicated data
- ❷ Improved reliability/availability through distributed transactions
- ❸ Improved performance
- ❹ Easier and more economical system expansion

# Distributed Database Promises

---

- ✓ ❶ Transparent management of distributed, fragmented, and replicated data
- ✓ ❷ Improved reliability/availability through distributed transactions
- ❸ Improved performance
- ❹ Easier and more economical system expansion

# Performance Improvement

---

- Proximity of data to its points of use

- ▢ Requires some support for fragmentation and replication

- Parallelism in execution

- ▢ Inter-query parallelism

- ▢ Intra-query parallelism



# Parallelism Requirements

---

- Have as much of the data required by *each* application at the site where the application executes

- Full replication

- How about updates?

- Updates to replicated data requires implementation of distributed concurrency control and commit protocols

# Distributed Database Promises

---

- ✓ ❶ Transparent management of distributed, fragmented, and replicated data
- ✓ ❷ Improved reliability/availability through distributed transactions
- ✓ ❸ Improved performance
- ❹ Easier and more economical system expansion

# Distributed Database Promises

---

- ✓ ❶ Transparent management of distributed, fragmented, and replicated data
- ✓ ❷ Improved reliability/availability through distributed transactions
- ✓ ❸ Improved performance
- ✓ ❹ Easier and more economical system expansion


# Outline

---

- Definition and Differentiation
- Promises and Challenges
- Architecture Alternatives

# Standard ANSI Database Architecture

---



# Dimensions of the Problem

---

## ■ Distribution

- ▢ Whether the components of the system are located on the same machine or not

## ■ Heterogeneity


- ▢ Various levels (hardware, communications, operating system)
- ▢ DBMS important one
  - ◆ data model, query language, transaction management algorithms

## ■ Autonomy

- ▢ Not well understood and most troublesome
- ▢ Various versions
  - ◆ **Design autonomy**: Ability of a component DBMS to decide on issues related to its own design.
  - ◆ **Communication autonomy**: Ability of a component DBMS to decide whether and how to communicate with other DBMSs.
  - ◆ **Execution autonomy**: Ability of a component DBMS to execute local operations in any manner it wants to.

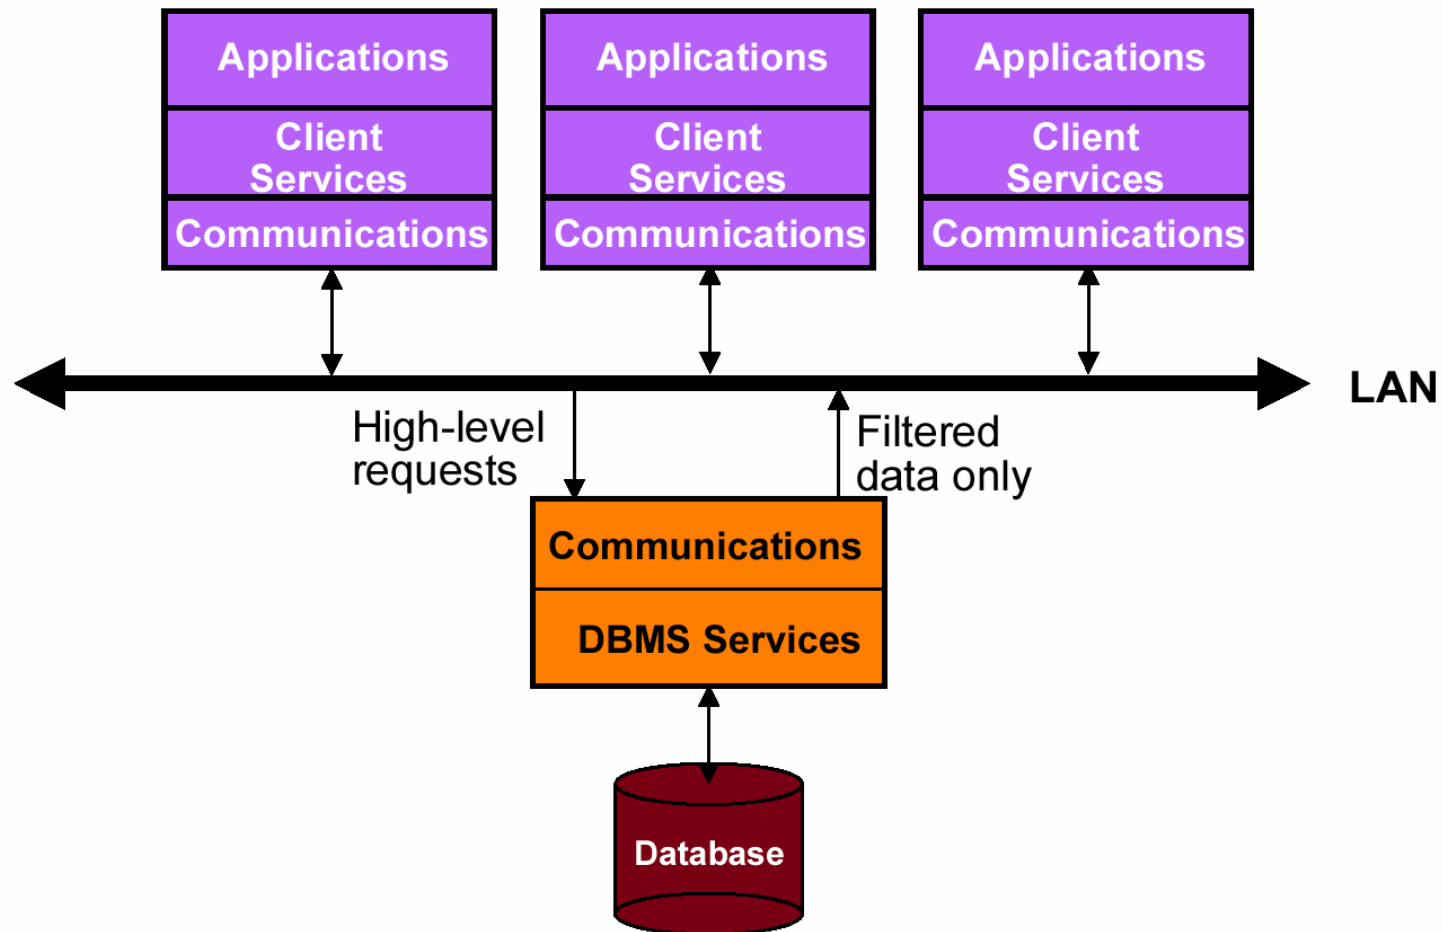
# DBMS Architecture Alternatives

---



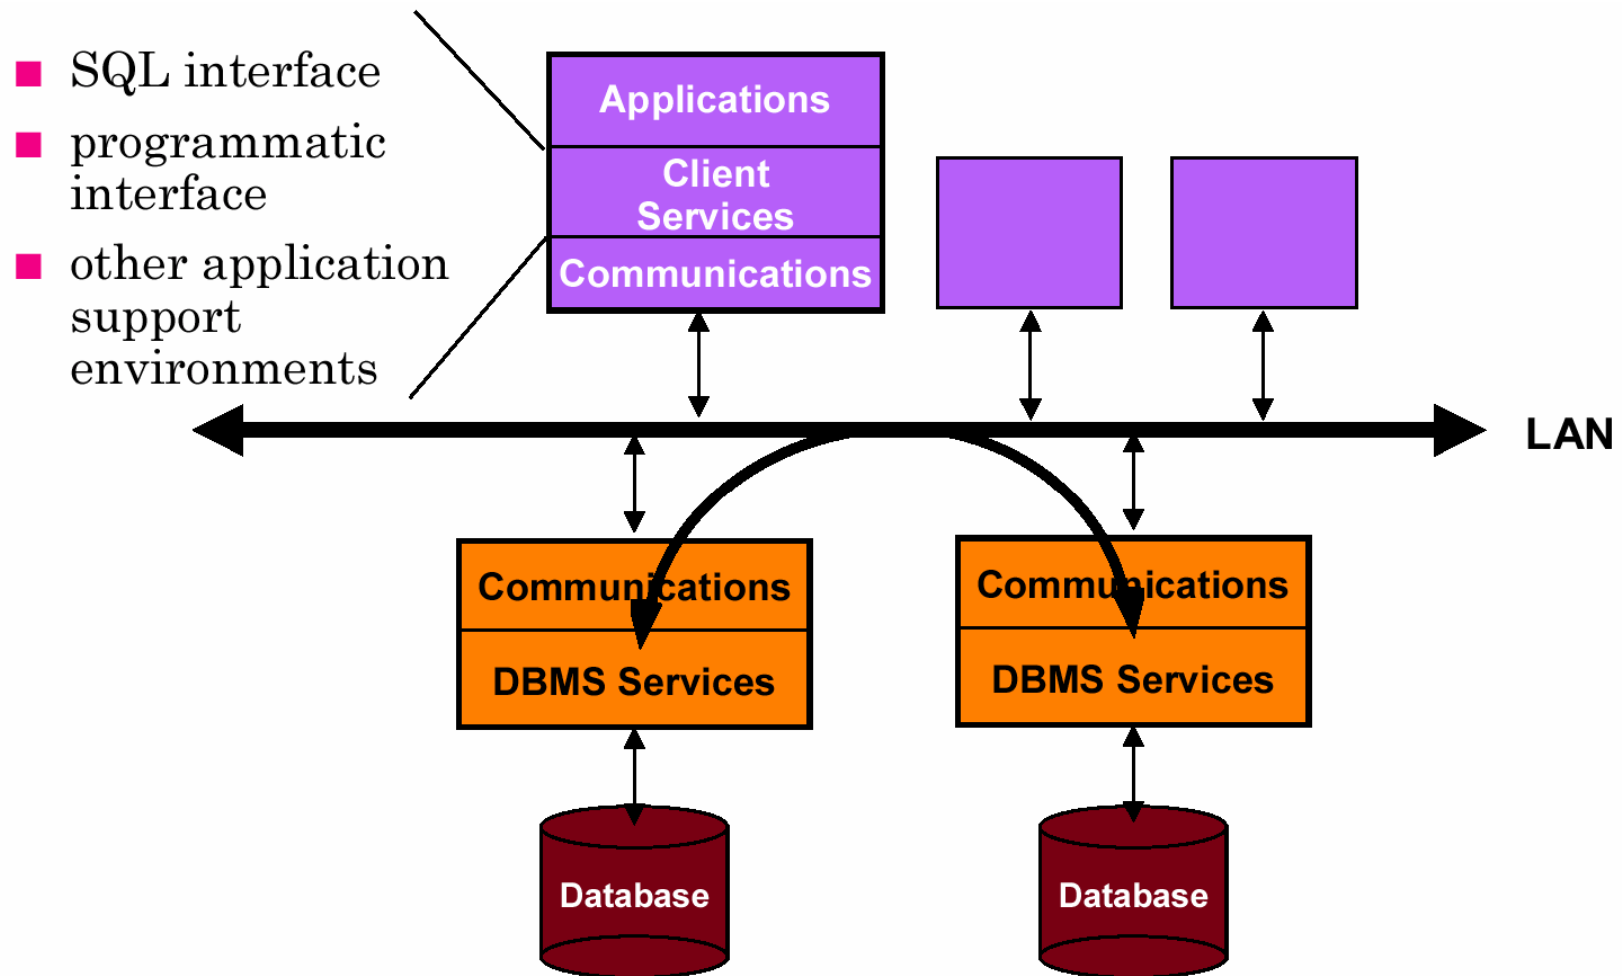
# 1. Multiple Clients / Single Server

---






# 1'. Multiple Clients / Multiple Server




## 2. Datalogical Distributed Database Architecture



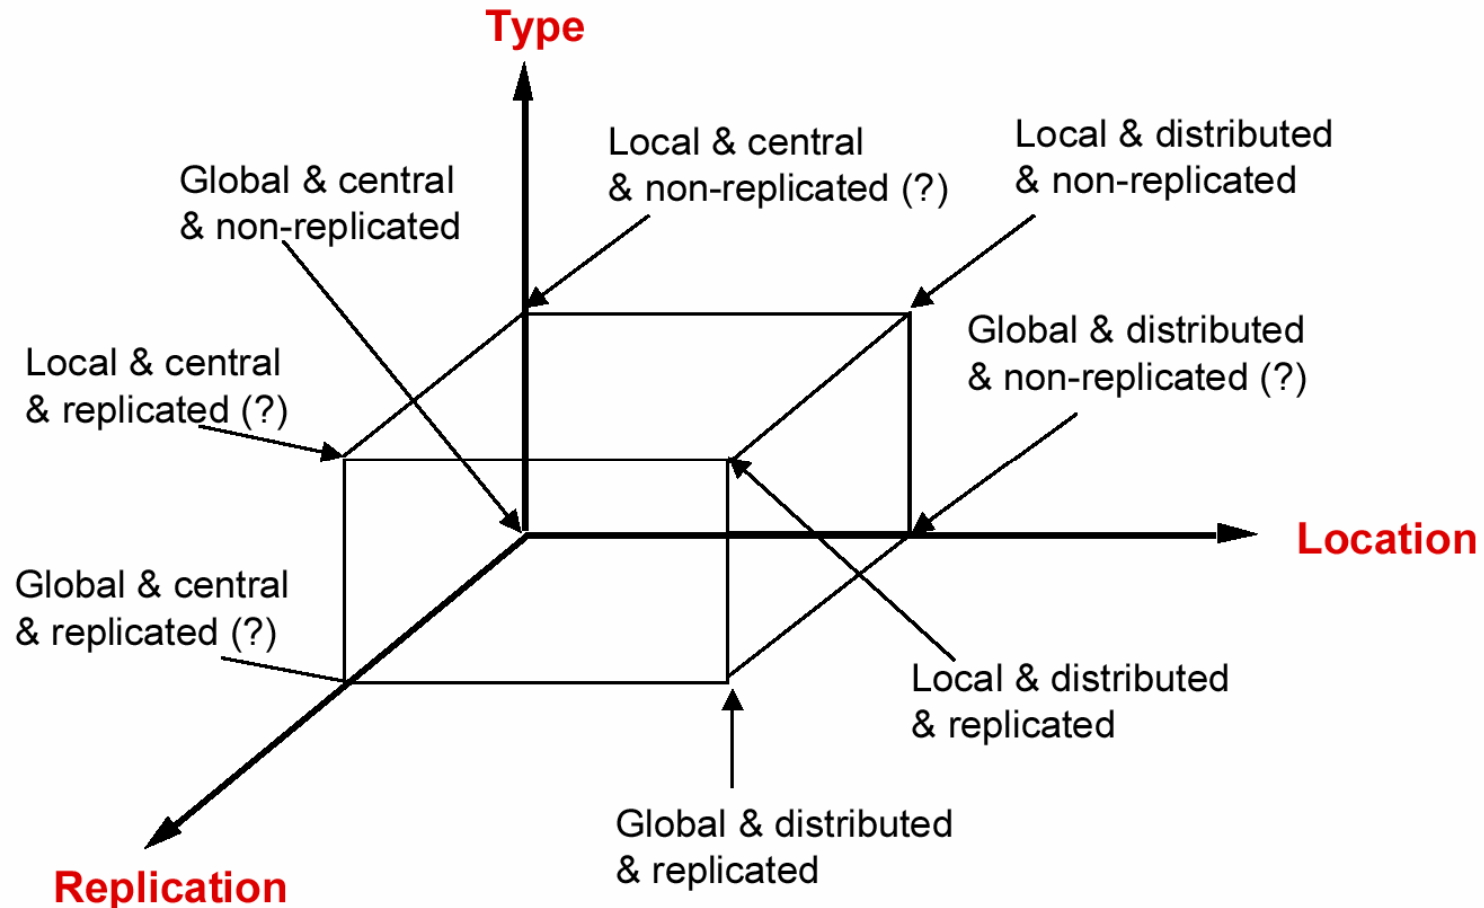
### 3. Datalogical Multi-DBMS Architecture

---



# Global Directory - Design Dimensions

---



# Next: Focus Issues

---

## ■ Distributed Database Design

- ▶ how to distribute the database
- ▶ replicated & non-replicated database distribution
- ▶ a related problem in directory management

## ■ Query Processing

- ▶ convert user transactions to data manipulation instructions
- ▶ optimization problem
- ▶  $\min\{\text{cost} = \text{data transmission} + \text{local processing}\}$
- ▶ general formulation is NP-hard

# References

---

- M.T. Özsu and P. Valduriez. *Principles of Distributed Database Systems*, 2<sup>nd</sup> edition. Prentice-Hall, 1999.
- M.T. Özsu and P. Valduriez. “Distributed and Parallel Database Systems”, In *Handbook of Computer Science and Engineering*, A. Tucker (ed.), CRC Press, 1997, pages 1093–1111 (Chapter 48).
- M.T. Özsu and P. Valduriez. “Distributed Database Systems: Where Are We Now?” *Computer*, August 1991, 24(8).
- M. T. Özsu. “The Push/Pull Effect - Can Distributed Database Technology Meet The Challenges of New Applications?” *Database Programming & Design*, April 1997.

# References

---

## ■ General:

- ▶ S. Ceri and G. Pelagatti. *Distributed Databases - Principles and Systems*. McGraw Hill, 1984.
- ▶ D.A. Bell and J.B. Grimson. *Distributed Database Systems*. Addison-Wesley, 1992.

## ■ Distributed Query Processing:

- ▶ C.T. Yu and W. Meng. *Principles of Database Query Processing for Advanced Applications*. Morgan-Kaufmann, 1998.
- ▶ J.C. Freytag, D. Maier, and G. Vossen. *Query Processing for Advanced Database Systems*. Morgan-Kaufmann, 1994.
- ▶ W. Kim, D.S. Reiner and D.S. Batory. *Query Processing in Database Systems*. Springer-Verlag, 1985.
- ▶ G. Graefe. “Query Evaluation Techniques for Large Databases”, *ACM Computing Surveys*, 25(2): 73-170, June 1993.
- ▶ P. Mishra and M. H. Eich. “Join Processing in Relational Databases”, *ACM Computing Surveys*, 24(1): 63-113, March 1992.
- ▶ M. Jarke and J. Koch. “Query Optimization in Database Systems,” *Computing Surveys*, June 1984, 16(2): 227-269.

# References

---

## ■ Transaction Management:

- ▶ P.A. Bernstein; V. Hadzilacos and N. Goodman. *Concurrency Control and Recovery in Database Systems*. Addison-Wesley, 1987.
- ▶ J. Gray and A. Reuter. *Transaction Processing - Concepts and Techniques*. Morgan Kaufmann, 1993.
- ▶ A.K. Elmagarmid. *Transaction Models for Advanced Database Applications*. Morgan Kaufmann, 1992.
- ▶ P.A. Bernstein. “Transaction Processing Monitors,” *Communications of ACM*, November 1990, 33(11): 75 – 86.
- ▶ B.K. Bhargava (ed.). *Concurrency Control and Reliability in Distributed Systems*. Van Nostrand Reinhold; 1987.
- ▶ C.H. Papadimitriou. *The Theory of Concurrency Control*. Computer Science Press, 1986.
- ▶ S.K. Shrivastava (ed.) *Reliable Computer Systems*. Springer-Verlag, 1985.



# References

---

## ■ Open Systems and Interoperability:

- ▶▶▶ A. Umar. *Distributed Computing – A Practical Synthesis*. Prentice-Hall, 1993.
- ▶▶▶ S. Khoshafian, A. Chan, A. Wong and H.K.T. Wong. *A Guide to Developing Client/Server SQL Applications*. Morgan Kaufmann, 1992.
- ▶▶▶ G. Nutt. *Open Systems*. Prentice-Hall, 1992.
- ▶▶▶ A.K. Elmagarmid and C. Pu (guest eds.). *ACM Computing Surveys*, Special Issue on Heterogeneous Databases, September 1990, 22(3).
- ▶▶▶ S. Ram (guest editor). *Computer*, Special Issue on Heterogeneous Distributed Database Systems, December 1991, 24(12).

# References

---

## ■ Parallel Database Servers:

- ▶▶▶ P. Valduriez (ed). *Data Management and Parallel Processing*. Chapman and Hall, 1992.
- ▶▶▶ M. Abdelguerfi and K-F. Wong. *Parallel Database Techniques*. IEEE Computer Society Press, 1988.
- ▶▶▶ P. Valduriez. “Parallel Database Systems: Open Problems and New Issues,” *Parallel and Distributed Databases*, April 1993, 1(2): 137–165.
- ▶▶▶ D. DeWitt and J. Gray. Parallel Database Systems: The Future of High-Performance Database Systems. *Communications of ACM*, June 1992, 35(6), 1992.

# References

---

## ■ Distributed Object Management:

- ⇒ E. Bertino and L. Martino. *Object-Oriented Database Systems*. Addison-Wesley, 1993.
- ⇒ A. Kemper and G. Moerkotte. *Object-Oriented Database Management*. Prentice-Hall, 1994.
- ⇒ A. Dogac, M.T. Özsu, A. Biliris, T. Sellis (ed.). *Advances in Object-Oriented Database Systems*. Springer-Verlag, 1994.
- ⇒ M.T. Özsu, U. Dayal and P. Valduriez (eds.). *Distributed Object Management*. Morgan Kaufman, 1994.
- ⇒ W. Kim (ed). *Modern Database Management – Object-Oriented and Multidatabase Technologies*. Addison-Wesley/ACM Press, 1994.
- ⇒ S. Zdonik and D. Maier (eds.). *Readings in Object-Oriented Database Systems*. Morgan Kaufmann, 1990.
- ⇒ R.G.G. Cattell. *Object Data Management*, 2nd edition. Addison-Wesley, 1994.

# References

---

## ■ Mobile Databases

- ▶ T. Imielinski and H. Korth. *Mobile Computing*. Kluwer Publishers, 1996.
- ▶ E. Pitoura and G. Samaras. *Data Management for Mobile Computing*. Kluwer Publishers, 1998.
- ▶ T. Imielinski and B.R. Badrinath. Data Management Issues in Mobile Computing. *Communications of ACM*, October 1994, 37(10):18-28.
- ▶ M. H. Dunham and A. Helal. Mobile Computing and Databases: Anything New? *ACM SIGMOD Record*, December 1995, 24(4): 5-9.
- ▶ G. H. Forman and J. Zahorjan. The Challenges of Mobile Computing, *Computer*, April 1994, 27(4):38-47.