Mapping the data warehouse architecture to Multiprocessor architecture

P.POOVIZHI

AP – IT

SNSCE

Parallel Relational Data Base Technology

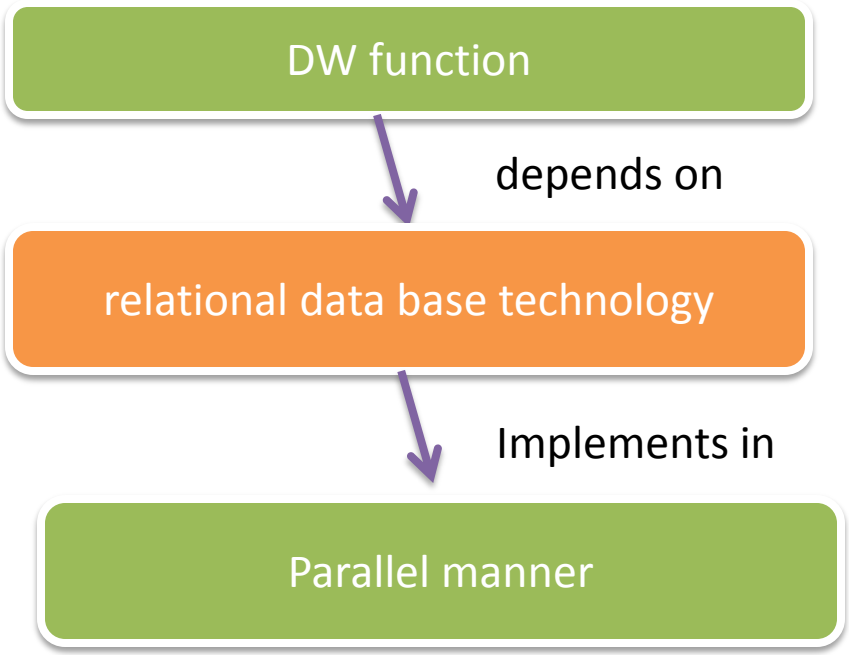

- **Two advantages:**
- **Linear speed up:** refers the ability to increase the number of processor to reduce response time
- **Linear scale up:** refers the ability to provide same performance on the same requests as the database size increases

Types of parallelism

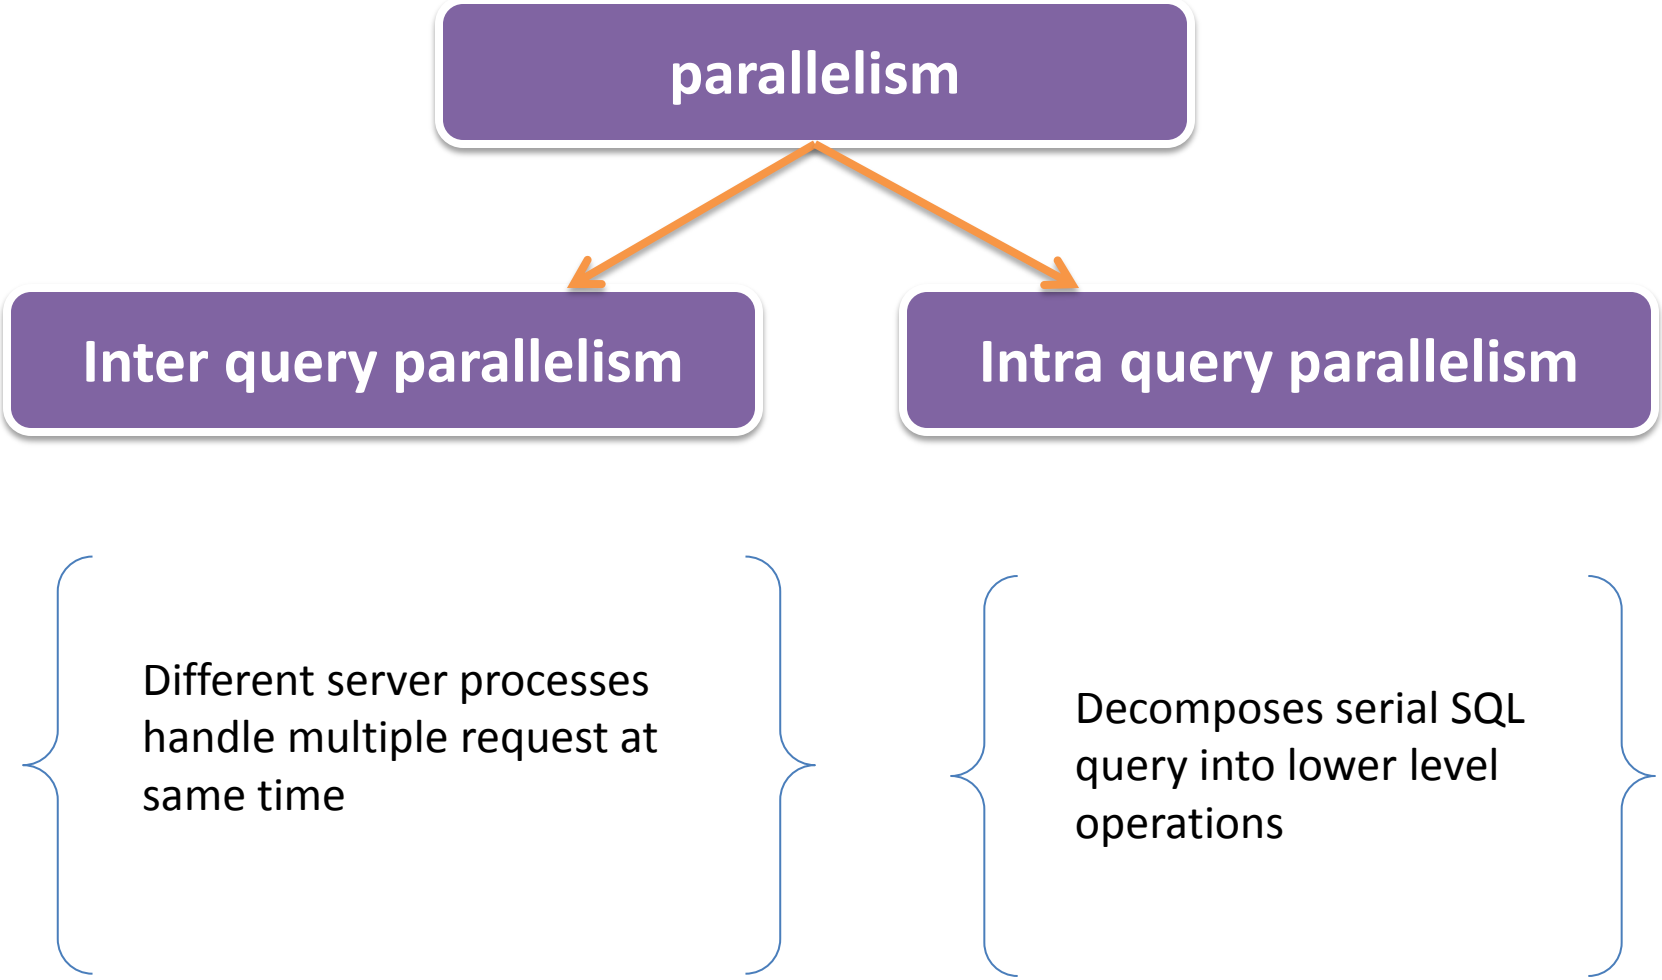

Types of parallelism

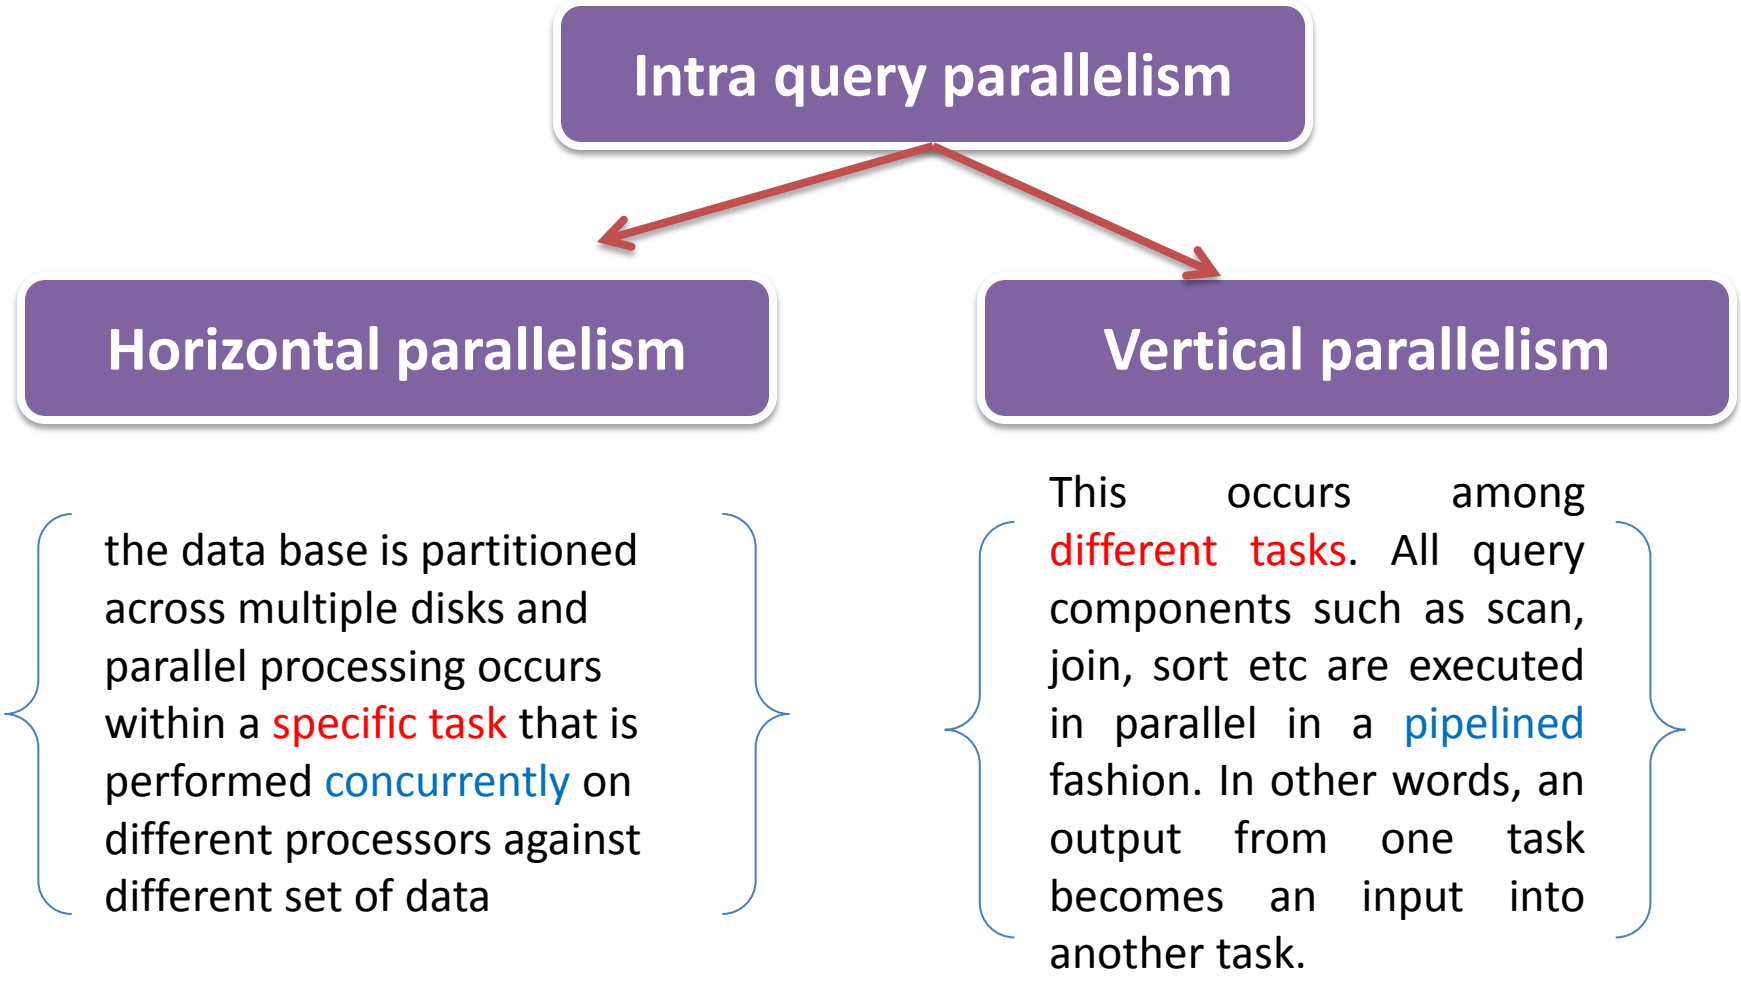

Types of parallelism

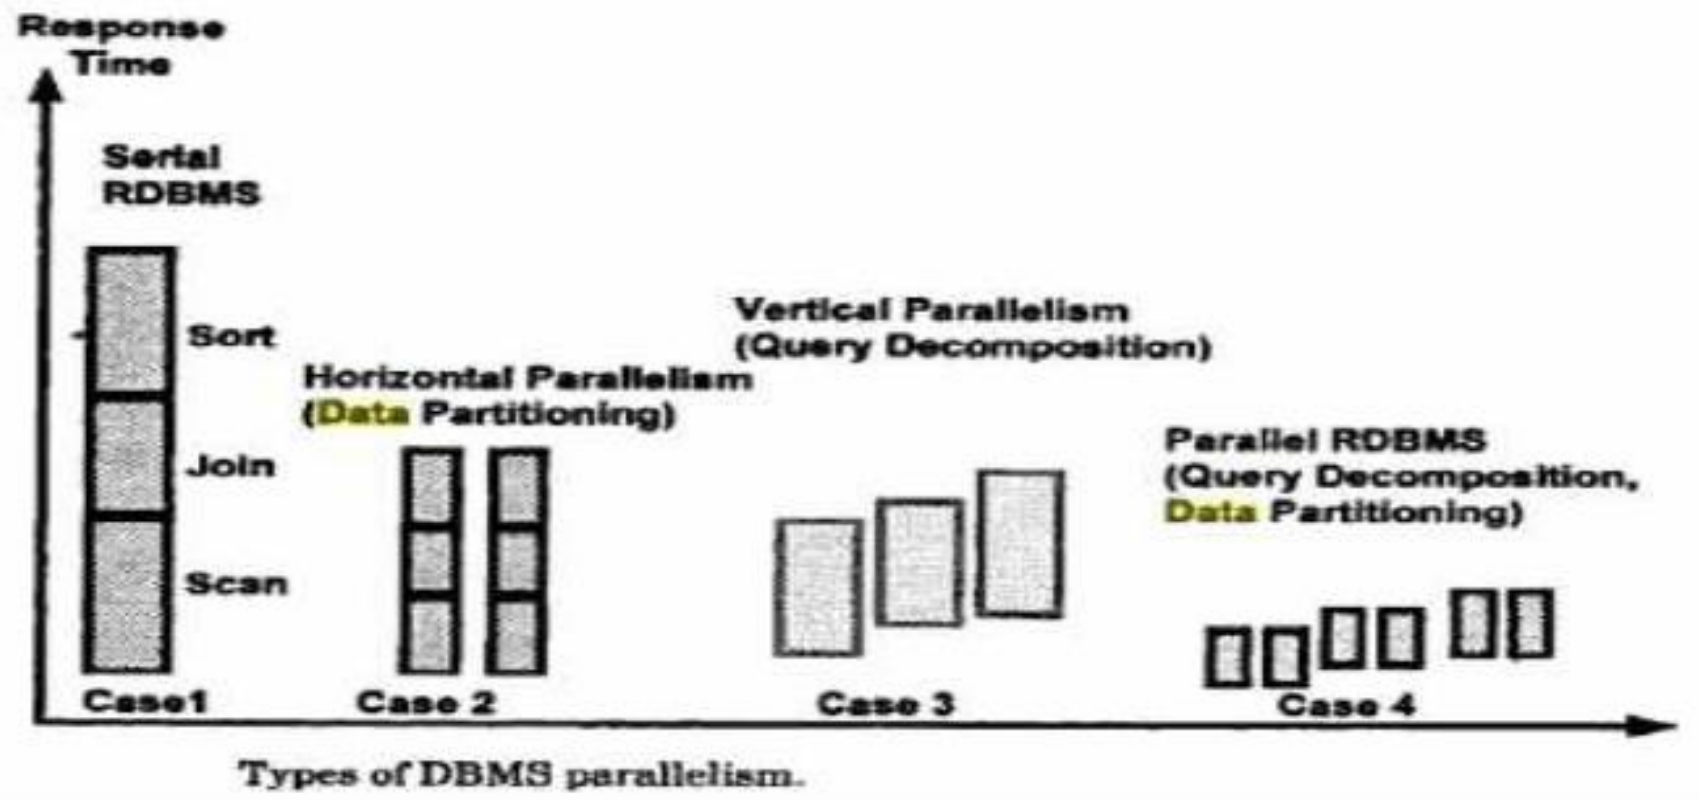

Types of DBMS parallelism.

Data Partitioning

- Key component for effective parallel execution of data base operations.

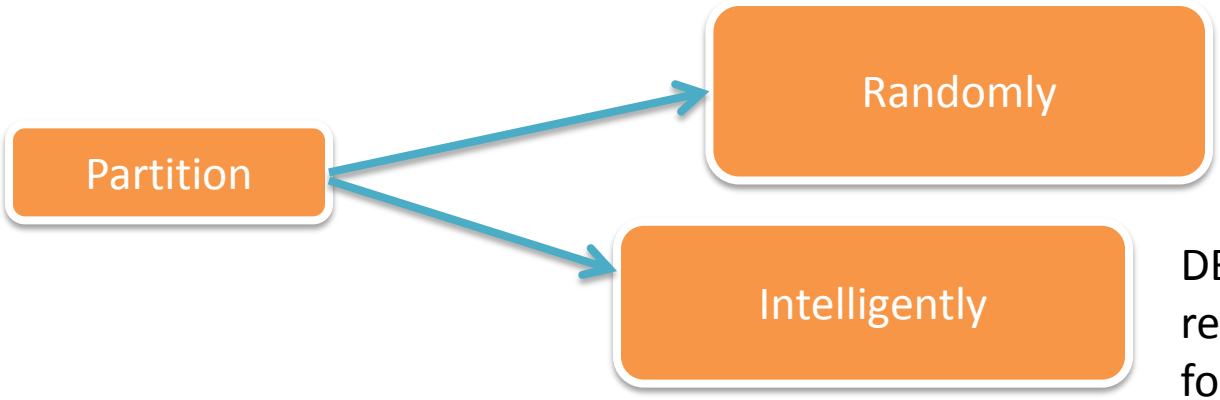

- Random data striping across multiple disks on a single server.
- Round robin

DBMS knows the location of record, does not waste time for searching

Hash partitioning: A hash algorithm is used to calculate the partition number based on the value of the partitioning key for each row

Key range partitioning: Rows are placed and located in the partitions according to the value of the partitioning key. That is all the rows with the key value from A to K are in partition 1, L to T are in partition 2 and so on.

Schema partitioning: an entire table is placed on one disk; another table is placed on different disk etc. This is useful for small reference tables.

User defined partitioning: It allows a table to be partitioned on the basis of a user defined expression.

Data base architectures of parallel processing

- There are three DBMS software architecture styles for parallel processing:
    1. Shared memory or shared everything Architecture
    2. Shared disk architecture
    3. Shared nothing architecture

Shared Memory Architecture

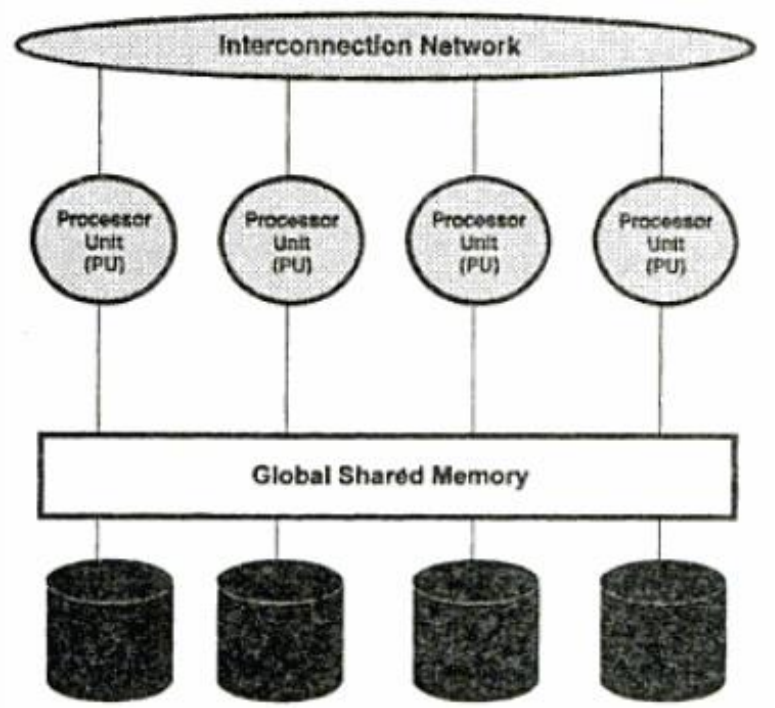

Shared-memory architecture.

- **Characteristics:**
  - Multiple PUs share memory.
  - Each PU has full access to all shared memory through a common bus.
  - Communication between nodes occurs via shared memory.
  - Performance is limited by the bandwidth of the memory bus.

Shared Memory Architecture

- Symmetric multiprocessor (SMP) machines are often nodes in a cluster.
- Multiple SMP nodes can be used with Oracle Parallel Server in a tightly coupled system, where memory is shared among the multiple PUs, and is accessible by all the PUs through a memory bus.
- Examples of tightly coupled systems include the Pyramid, Sequent, and Sun SparcServer.
- Performance is potentially limited in a tightly coupled system by a number of factors.
- These include various system components such as the memory bandwidth, PU to PU communication bandwidth, the memory available on the system, the I/O bandwidth, and the bandwidth of the common bus.

Advantages & Disadvantages – shared memory

- Advantages:
  - Memory access is cheaper than inter-node communication. This means that internal synchronization is faster than using the Lock Manager.
  - Shared memory systems are easier to administer than a cluster.
- Disadvantages:
  - Scalability is limited by bus bandwidth and latency, and by available memory.

Shared Disk Architecture

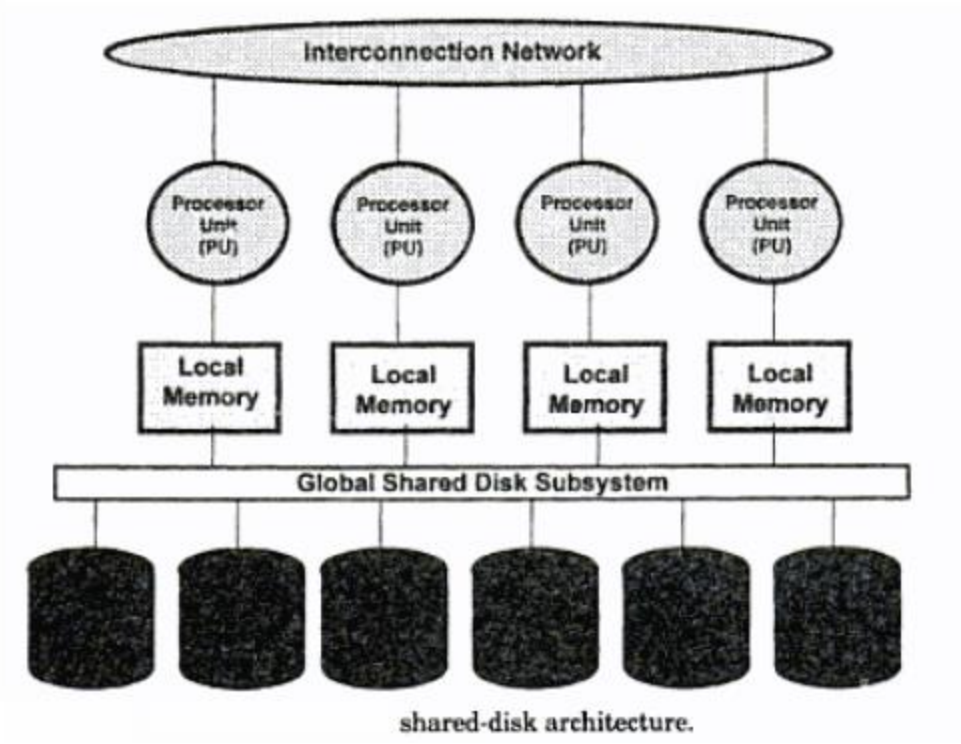

shared-disk architecture.

characteristics:

- Each node consists of one or more PUs and associated memory.
- Memory is not shared between nodes.
- Communication occurs over a common high-speed bus.
- Each node has access to the same disks and other resources.
- A node can be an SMP if the hardware supports it.
- Bandwidth of the high-speed bus limits the number of nodes (scalability) of the system.

Shared Disk Architecture

- The cluster illustrated in figure is composed of multiple tightly coupled nodes. The Distributed Lock Manager (DLM) is required. Examples of loosely coupled systems are VAX clusters or Sun clusters.
- Since the memory is not shared among the nodes, each node has its own data cache. Cache consistency must be maintained across the nodes and a lock manager is needed to maintain the consistency. Additionally, instance locks using the DLM on the Oracle level must be maintained to ensure that all nodes in the cluster see identical data.
- There is additional overhead in maintaining the locks and ensuring that the data caches are consistent. The performance impact is dependent on the hardware and software components, such as the bandwidth of the high-speed bus through which the nodes communicate, and DLM performance.

Advantages & Disadvantages – shared disk system

- **Advantages:**
  - Shared disk systems permit high availability. All data is accessible even if one node dies.
  - These systems have the concept of one database, which is an advantage over shared nothing systems.
  - Shared disk systems provide for incremental growth.
- Disadvantages:
  - If the workload is not partitioned well there may be high synchronization overhead.
  - There is operating system overhead of running shared disk software.

Shared Nothing Architecture

- In shared nothing systems only one CPU is connected to a given disk.
- If a table or database is located on that disk, access depends entirely on the PU which owns it.

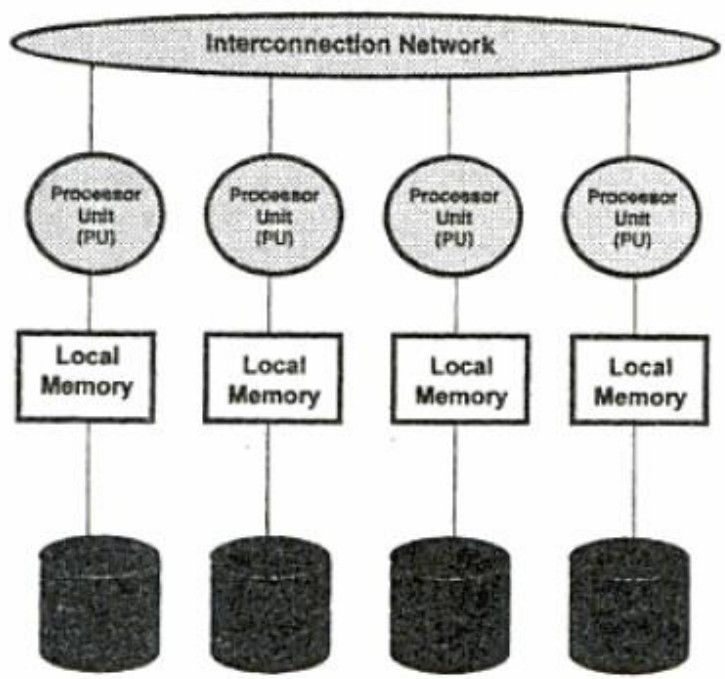

Distributed memory architecture.

advantages and disadvantages for parallel processing – shared nothing

- **Advantages**
  - Shared nothing systems provide for incremental growth.
  - System growth is practically unlimited.
  - MPPs are good for read-only databases and decision support applications.
  - Failure is local: if one node fails, the others stay up.
- **Disadvantages**
  - More coordination is required.
  - More overhead is required for a process working on a disk belonging to another node.
  - If there is a heavy workload of updates or inserts, as in an online transaction processing system, it may be worthwhile to consider data-dependent routing to alleviate contention.

Shared disk vs shared nothing

Shared disk architecture vs. shared nothing architecture	
Shared disk Architecture	Shared nothing architecture
Requires special hardware	Does not require special hardware
Non linear scalability	Provides near linear scalability
Balanced CPU or node fail-over	Balanced/Unbalanced CPU or node fail-over
Requires CPU level communication at disk access	Minimal communication
Non disruptive maintenance	Non disruptive maintenance

Parallel DBMS features

- Scope and techniques of parallel DBMS operations
- Optimizer implementation
- Application transparency
- Parallel environment which allows the DBMS server to take full advantage of the existing facilities on a very low level
- DBMS management tools help to configure, tune, admin and monitor a parallel RDBMS as effectively as if it were a serial RDBMS
- Price / Performance: The parallel RDBMS can demonstrate a non linear speed up and scale up at reasonable costs.

Parallel DBMS vendors

- 1. Oracle: Parallel Query Option (PQO)
- Architecture: shared disk arch
- Data partition: Key range, hash, round robin
- Parallel operations: hash joins, scan and sort
- 2. Informix: eXtended Parallel Server (XPS)
- Architecture: Shared memory, shared disk and shared nothing models
- Data partition: round robin, hash, schema, key range and user defined
- Parallel operations: INSERT, UPDATE, DELELTE
- 3. IBM: DB2 Parallel Edition (DB2 PE)
- Architecture: Shared nothing models
- Data partition: hash
- Parallel operations: INSERT, UPDATE, DELELTE, load, recovery, index creation, backup, table reorganization
- 4. SYBASE: SYBASE MPP
- Architecture: Shared nothing models
- Data partition: hash, key range, Schema
- Parallel operations: Horizontal and vertical parallelism