

Parallel and Distributed Algorithms

Parallel and Distributed Algorithms

1. Parallel Algorithms

- Designed to run multiple computations simultaneously on multiple processors/cores within a single machine to solve a problem faster.

2. Distributed Algorithms

- Designed to run across multiple independent machines (nodes) connected by a network, often without shared memory.

3. Key Challenges

- **Synchronization**: coordinating processes correctly.
- **Communication overhead**: time spent passing data between processors/nodes.
- **Fault tolerance**: handling node/processor failures gracefully.

4. MapReduce Paradigm

- A widely used distributed computing model: "map" processes data in parallel across nodes, "reduce" aggregates the results.

5. Consensus Algorithms

- Distributed systems need protocols (e.g., Paxos, Raft) to agree on shared state even when some nodes fail or messages are delayed.

6. Why It Matters

- As data volumes grow beyond what a single machine can handle, parallel and distributed algorithms are essential for search engines, big data processing, and cloud computing.