

Operating System Class Test-01

Question 1 : Process States & Context Switching

A Bank's Customer service system runs multiple processes to handle online transactions, chat support, and fraud detection. During peak hours, the system slows down, and some transactions get delayed.

Task :

- a) Explain the different Process States that a transaction might go through in this system.
- b) If the system switches from handling one transaction to another, which OS Component is responsible for the transition?
- c) How does a Context switch impact performance, and how can the bank reduce the overhead?

Ans To the Question No 1

a) Process States in the Bank's Customer Service System :

In an operating system, a process goes through several states during its execution. The different process states that a transaction might go through in the bank's customer service system are:

1. **New:** The transaction process is created when a customer initiates an online transaction, chat support request, or fraud detection task.
2. **Ready:** The process is waiting in the ready queue for CPU allocation after getting all the necessary resources.
3. **Running:** The process is being executed by the CPU, handling the customer's transaction, chat request, or fraud detection task.
4. **Blocked (Waiting):** If the process requires I/O operations (e.g., waiting for database verification or payment gateway response), it moves to the blocked state.
5. **Terminated:** Once the transaction is completed successfully or failed, the process is terminated and removed from the system.
6. **Suspended (Optional):** If the system is overloaded, some processes might be moved to secondary storage (suspended state) to free up resources.

b) OS Component Responsible for Process Transition :

The **CPU Scheduler** and **Dispatcher** are responsible for handling the transitions between process states:

- ❑ The **Long-Term Scheduler** decides which new processes enter the ready queue.
- ❑ The **Short-Term Scheduler** selects which ready process will be executed next.
- ❑ The **Dispatcher** performs the actual context switch between processes, moving them between Ready, Running, and Blocked states.

c) Impact of Context Switching on Performance and Reducing Overhead :

Impact of Context Switching on Performance :

- Context switching requires saving the state of the current process and loading the state of the next process. This involves saving registers, program counters, and memory mappings, which consumes CPU cycles.
- Frequent context switches increase system overhead, leading to delays in transactions.
- High context switch rates during peak hours slow down the response time for online transactions and chat support.

Ways to Reduce Context Switching Overhead :

1. **Process Prioritization:** Assign higher priority to critical banking tasks (e.g., fraud detection) to reduce unnecessary context switches.
2. **Batch Processing:** Group similar transactions to minimize frequent switching.
3. **Increase Time Quantum:** If using a round-robin scheduling algorithm, increasing the time slice can reduce the number of context switches.
4. **Multithreading:** Using threads instead of processes can reduce overhead since threads share the same process memory.
5. **Load Balancing:** Distribute the workload across multiple servers to prevent excessive switching on a single system.

Question 2 : Multiprogramming & Time-Sharing System

A gaming Company is designing a cloud-based game streaming service where multiple players share server resources in real time.

Task :

- a) Would a Multiprogramming or Time-Sharing OS be better suited for this system? justify your answer.
- b) What challenges can occur if too many players connect at the same time?
- c) How can the OS Scheduler manage fair resource allocation among all players?

Ans To the Question No 2

a) Multiprogramming vs. Time-Sharing OS for Cloud-Based Game Streaming :

A **Time-Sharing Operating System (TSOS)** would be better suited for a cloud-based game streaming service rather than a **Multiprogramming OS**.

Justification:

- **Time-Sharing OS** allows multiple users (players) to interact with the system simultaneously by rapidly switching between tasks, giving each player a small slice of CPU execution. This ensures real-time responsiveness, which is critical for gaming.
- **Multiprogramming OS**, on the other hand, focuses on maximizing CPU utilization by running multiple processes concurrently but does not guarantee immediate user interaction, which is essential for real-time gaming.

Thus, a **Time-Sharing OS** is ideal because it enables multiple players to play simultaneously with minimal delays.

b) Challenges of Too Many Players Connecting Simultaneously :

When a large number of players connect at the same time, the system may face several challenges, including:

1. **Resource Overload:** High CPU, memory, and network bandwidth consumption can cause lag, freezing, or crashes.
2. **Increased Context Switching:** Excessive process switching can degrade performance, causing delays in rendering and input processing.
3. **Network Congestion:** Too many simultaneous connections can overwhelm the network infrastructure, leading to packet loss and high latency.

4. **Server Bottleneck:** Limited computing resources (CPU, GPU, and RAM) may not handle all requests efficiently, leading to degraded gaming experience.
5. **Fairness Issues:** Some players may experience smooth gameplay while others face lag due to uneven resource allocation.

c) How the OS Scheduler Manages Fair Resource Allocation :

The **OS Scheduler** ensures fair distribution of CPU, memory, and network resources using the following techniques :

1. **Priority Scheduling:** Critical tasks (e.g., game physics, rendering, and input processing) are given higher priority over background tasks.
2. **Round-Robin Scheduling:** Each player's process is given a time slice, ensuring fair CPU access and preventing starvation.
3. **Dynamic Load Balancing:** The system distributes the workload across multiple servers to prevent overload and ensure smooth performance.
4. **Quality of Service (QoS) Policies:** The OS may prioritize players with stable connections to reduce lag and maintain a balanced gaming experience.
5. **Adaptive Resource Allocation:** If more players join, the OS can allocate resources dynamically based on demand, ensuring optimal usage.
6. **Virtualization & Cloud Scaling:** The system can automatically allocate additional virtual machines (VMs) or cloud resources to handle spikes in player activity.

By implementing these scheduling strategies, the OS ensures a smooth, fair, and real-time gaming experience for all players, even under heavy load.

Question 3 : Inter-Process Communication (IPC)

A ride-sharing app (like Uber) has two main processes :

1. Process A : Continuously updates the driver's location.
2. Process B : Matches Drivers with nearby passenger based on the location data.

Task :

- a) Which Inter-Process Communication (IPC) mechanism would be the best for real-time location sharing (message passing or shared memory) ? Explain.
- b) If two drivers request the same passenger at the same time, how can synchronization prevent conflicts?
- c) What problems might arise if the IPC method is not well-optimized?

Ans To the Question No 3

(a) Best IPC Mechanism for Real-Time Location Sharing :

For real-time location sharing in the ride-sharing app, **Shared Memory** would be the best IPC mechanism.

Justification:

- **Low Latency:** Shared memory allows both Process A (driver location updater) and Process B (matching passengers) to access and update location data instantly without the overhead of message passing.
- **Efficient Data Sharing:** Since location data is continuously updated, shared memory enables direct access without frequent copying, reducing communication delays.
- **Real-Time Performance:** Ride-matching requires immediate access to driver locations, which is more efficient with shared memory than message queues, where data transfer incurs additional processing time.

However, **shared memory requires synchronization mechanisms (like semaphores or mutexes) to avoid race conditions** when multiple processes try to read or update the data simultaneously.

(b) Synchronization to Prevent Conflicts When Two Drivers Request the Same Passenger :

If two drivers request the same passenger at the same time, **synchronization mechanisms** must be used to ensure only one driver is matched. Some possible solutions include:

1. **Mutex (Mutual Exclusion Lock):** Before assigning a passenger, the system locks the record to prevent multiple drivers from accessing it simultaneously. Once the passenger is assigned, the lock is released.
2. **Atomic Transactions:** The ride-matching process can be implemented as an atomic transaction, ensuring that only one driver is successfully matched at a time.
3. **First-Come, First-Serve (FCFS) Scheduling:** The system can process driver requests in order of arrival, assigning the passenger to the first valid request.
4. **Timestamp-Based Synchronization:** Each driver's request can have a timestamp, and the system assigns the passenger to the earliest valid request.
5. **Priority-Based Allocation:** If needed, the system may prioritize drivers based on factors like proximity, rating, or time since the last ride.

These synchronization techniques prevent race conditions and ensure fairness in driver-passenger matching.

(c) Problems if the IPC Method is Not Well-Optimized :

If the IPC mechanism is not optimized, several issues can arise:

1. **Data Inconsistency:** Without proper synchronization, multiple processes may read/write conflicting location data, leading to incorrect driver-passenger matching.
2. **High Latency:** Inefficient IPC (such as excessive message passing) can introduce delays, causing lag in ride matching and frustrating users.
3. **Race Conditions:** If multiple processes update shared memory simultaneously without control, it can lead to incorrect or lost updates.
4. **Deadlocks:** Improper use of locks (e.g., two processes waiting indefinitely for each other to release a resource) can freeze ride-matching operations.

5. **Resource Overhead:** Excessive IPC operations can consume CPU and memory resources, slowing down the entire system.

To prevent these issues, the system should implement **efficient shared memory access with proper synchronization techniques like mutexes, semaphores, or transactional memory operations.**