



# Question & Answers

DEADLOCKS

Sercan Külcü | Operating Systems | 10.04.2023

# Contents

What is a deadlock in the context of operating systems?.....	3
What causes deadlocks to occur in multi-process/multi-threaded environments? .....	3
Why are deadlocks a challenging problem in operating system design? .....	4
What are some common techniques used to prevent deadlocks? .....	4
What is the role of resource allocation and scheduling policies in preventing and resolving deadlocks? .....	4
What happens if a non-recursive mutex is locked more than once.....	5
What are the necessary conditions which can lead to a deadlock in a system? .....	5
What is a Banker's algorithm? .....	6
When is a system in a safe state?.....	6
Why do we use precedence graphs?.....	7
Explain the resource allocation graph? .....	7
How to recover from a deadlock?.....	8
How do operating systems detect and recover from deadlocks? .....	8
What is the difference between a deadlock and a livelock, and how are they handled in operating systems? .....	9
How do operating systems handle priority-based resource allocation in the context of deadlocks?.....	9
What is the impact of deadlock prevention techniques on system performance and efficiency? .....	9
How do distributed computing and cloud computing environments handle deadlocks, and what are some strategies for preventing and resolving them?.....	10

What are some advanced techniques for deadlock detection and resolution in operating systems? ..... 10

How do real-time and safety-critical systems handle deadlocks?..... 11

What is the role of transactional memory in preventing deadlocks, and how does it differ from traditional approaches? ..... 11

How do operating systems handle deadlocks in multi-processor/multi-core environments? ..... 12

What are some emerging trends and technologies in deadlock prevention and resolution in operating systems? ..... 12

## What is a deadlock in the context of operating systems?

A deadlock is a situation in which two or more processes or threads are blocked and waiting for each other to release resources that they need to continue executing. Deadlock refers to a state in which two or more processes are waiting for each other to complete, resulting in an impasse. This situation is analogous to two trains traveling towards each other on a single track, where neither can move past the other. Similarly, in an operating system, deadlock can occur when multiple processes are holding resources and waiting for other resources held by another process. This scenario can lead to a standstill where none of the processes can proceed, causing system performance to degrade significantly. Operating systems typically implement deadlock detection and resolution algorithms to mitigate such situations and ensure the efficient execution of processes.

## What causes deadlocks to occur in multi-process/multi-threaded environments?

Deadlocks can occur when multiple processes or threads are competing for a finite set of resources, such as memory, I/O devices, or locks. If each process holds a resource that another process needs and vice versa, a circular wait can occur, leading to a deadlock.

Why are deadlocks a challenging problem in operating system design?

Deadlocks are a challenging problem because they can result in a system hang or crash, which can be very difficult to debug and resolve. Moreover, preventing and resolving deadlocks without causing starvation or performance degradation is a non-trivial task.

What are some common techniques used to prevent deadlocks?

Some common techniques used to prevent deadlocks include resource allocation and scheduling policies, deadlock detection and recovery, and avoidance of circular wait by enforcing an ordering or hierarchy among resources.

What is the role of resource allocation and scheduling policies in preventing and resolving deadlocks?

Resource allocation and scheduling policies play a critical role in preventing and resolving deadlocks by ensuring that resources are allocated and released in a safe and efficient manner. For example, a resource allocation policy may ensure that resources are granted in a way that avoids circular wait, while a scheduling policy may prioritize processes or threads that are less likely to cause or be affected by deadlocks.

What happens if a non-recursive mutex is locked more than once.

A deadlock is a situation that can arise when a thread that has already acquired a mutex attempts to acquire the same mutex again. This causes the thread to enter the waiting list for that mutex, which can result in a deadlock because no other thread can release the mutex. To prevent deadlocks, it is important for an operating system to identify the owner of the mutex and return it if it is already locked by the same thread. This careful implementation helps to prevent deadlocks and ensure the proper functioning of the system.

What are the necessary conditions which can lead to a deadlock in a system?

**Mutual Exclusion:** There is a resource that cannot be shared.

**Hold and Wait:** A process is holding at least one resource and waiting for another resource, which is with some other process.

**No Preemption:** The operating system is not allowed to take a resource back from a process until the process gives it back.

**Circular Wait:** A set of processes are waiting for each other in circular form.

## What is a Banker's algorithm?

The banker's algorithm is a resource allocation and deadlock avoidance algorithm used in operating systems. The algorithm works by simulating the allocation of predetermined maximum amounts of all resources, and then checking if the resulting system state is safe or not. The safe state is one where there is no possibility of a deadlock occurring, i.e., every process can eventually acquire the resources it needs to complete execution. If the system state is found to be unsafe, the algorithm blocks the request and waits until the requested resources become available. The banker's algorithm is used to avoid deadlocks in systems with limited resources, such as shared memory, where multiple processes compete for the same resources. By carefully managing the allocation of resources, the algorithm ensures that all processes can complete execution without getting stuck in a deadlock.

## When is a system in a safe state?

In an operating system, a set of dispatchable processes is said to be in a safe state if they can all execute in such a way that it avoids the possibility of a deadlock. A deadlock is a situation where two or more processes are waiting for each other to release resources that they need to continue execution. A safe state means that the system can allocate resources to all of the processes and still guarantee that they will eventually complete without deadlock. This is an important concept in operating system design, as it allows for the efficient use of resources and the prevention of system crashes. To ensure a safe state, the operating system must implement appropriate resource allocation policies and algorithms that can detect and prevent deadlock.

## Why do we use precedence graphs?

A precedence graph is a useful tool in analyzing the execution order of several processes in an operating system. It is a directed acyclic graph (DAG), where the nodes of the graph represent individual statements of program code, and the edges between the nodes indicate the execution order. A directed edge from node A to node B shows that statement A executes first, and then statement B executes. This graph is often used in concurrent programming to identify dependencies between processes and to detect potential issues such as race conditions, deadlocks, and starvation. By constructing a precedence graph, developers can optimize the execution order of processes and ensure the correct and efficient execution of the program.

## Explain the resource allocation graph?

The resource allocation graph (RAG) is a visual representation of the state of the system in terms of processes and resources. It is used to represent how resources are being used and allocated among different processes in the system. In a RAG, each process is represented by a node, and each resource is represented by a resource type that is connected by edges to the nodes that are currently using it. A circle around a set of nodes and resources represents a potential deadlock. One of the advantages of having a diagram like a RAG is that sometimes it is possible to identify a deadlock directly, by examining whether the diagram contains any circular dependencies or loops. The RAG can also be used to detect other potential problems, such as resource starvation or inefficient resource allocation.



## How to recover from a deadlock?

Deadlocks are a serious problem in operating systems, and recovering from them is crucial to prevent system failure. There are several methods to recover from a deadlock. One method is **process termination**, which involves aborting all the deadlock processes or aborting one process at a time until the deadlock is eliminated. Another method is resource preemption, which involves releasing some resources from one or more processes to allow the others to continue execution. **Rollback** is another method, which involves reversing the actions of one or more processes to a previously saved state. Finally, selecting a victim is a method in which one of the processes involved in the deadlock is selected to be terminated or have its resources preempted to allow the others to proceed. It is important to note that recovery from a deadlock can be complex and may not always be successful. Therefore, it is best to prevent deadlocks from occurring in the first place through careful system design and resource allocation.

## How do operating systems detect and recover from deadlocks?

Operating systems use various algorithms to detect deadlocks, such as the deadlock detection algorithm that searches for cycles in the resource allocation graph. Once a deadlock is detected, the operating system can use several methods to recover from it, such as aborting one or more processes, preempting resources, or rolling back processes to a safe state.

What is the difference between a deadlock and a livelock, and how are they handled in operating systems?

A deadlock occurs when two or more processes are blocked waiting for each other to release resources. In contrast, a livelock occurs when two or more processes keep changing their state in response to each other's state changes, without making progress. Livelocks are generally harder to detect and recover from than deadlocks.

How do operating systems handle priority-based resource allocation in the context of deadlocks?

Priority-based resource allocation can increase the likelihood of deadlocks, as higher-priority processes may block lower-priority ones from acquiring resources. Operating systems can use various techniques, such as priority inheritance or priority ceiling, to prevent priority inversions and reduce the chances of deadlocks.

What is the impact of deadlock prevention techniques on system performance and efficiency?

Deadlock prevention techniques can add overhead to the system and reduce its performance and efficiency. For example, priority inheritance can increase the amount of time spent in kernel mode, while using larger timeouts in deadlock detection algorithms can delay the

detection of deadlocks. Therefore, a balance must be struck between the benefits of deadlock prevention and the costs of its implementation.

How do distributed computing and cloud computing environments handle deadlocks, and what are some strategies for preventing and resolving them?

Deadlocks can be a significant problem in distributed and cloud computing environments, where resources are distributed across multiple nodes. Techniques such as distributed deadlock detection and resolution algorithms can be used to detect and recover from deadlocks in these environments. Additionally, resource allocation and scheduling policies can be designed to prevent deadlocks from occurring in the first place.

What are some advanced techniques for deadlock detection and resolution in operating systems?

Advanced techniques for deadlock detection and resolution in operating systems include:

- **State prevention:** This involves modifying the resource allocation policy to avoid deadlocks from occurring in the first place.
- **Dynamic ordering:** This involves dynamically ordering the acquisition and release of resources to prevent deadlocks.

- **Wait-for graph algorithms:** These algorithms analyze the wait-for graph to detect and resolve deadlocks.
- **Resource preemption:** This involves preempting resources from low-priority processes to prevent deadlocks.
- **Timeouts:** This involves setting a timeout for processes waiting on a resource, and if the timeout expires, the process is aborted to prevent deadlocks.

How do real-time and safety-critical systems handle deadlocks?

Real-time and safety-critical systems typically use deterministic scheduling and resource allocation policies to avoid deadlocks. These systems also often have specialized hardware and software to detect and recover from deadlocks quickly.

What is the role of transactional memory in preventing deadlocks, and how does it differ from traditional approaches?

Transactional memory can be used to prevent deadlocks by providing a mechanism for multiple threads to execute a series of operations atomically. If an operation in the series fails, the transaction is aborted, and any resources acquired during the transaction are released.

## How do operating systems handle deadlocks in multi-processor/multi-core environments?

Deadlocks in multi-processor/multi-core environments can be challenging to handle because the resource allocation and scheduling policies must be coordinated across all processors/cores. This coordination can lead to increased overhead and decreased performance.

## What are some emerging trends and technologies in deadlock prevention and resolution in operating systems?

Emerging trends and technologies in deadlock prevention and resolution in operating systems include:

- Machine learning and AI-based approaches for predicting and preventing deadlocks.
- Hardware support for deadlock prevention and resolution.
- Distributed algorithms for deadlock detection and resolution.
- More sophisticated resource allocation and scheduling policies to prevent and resolve deadlocks in high-performance computing environments.