

Deadlock Detection and Recovery

Deadlock is a complex and potentially detrimental situation that can arise in computer systems where multiple processes are competing for the same shared resources. When two or more processes become deadlocked, it means that each process is holding on to resources that are necessary for the other process(es) to complete their tasks. This can lead to a complete standstill, as none of the processes can move forward without the release of the required resources. Deadlocks can cause severe performance and stability issues in a system, which can ultimately result in system downtime or even failure. Therefore, it is essential to promptly detect and recover from deadlocks to ensure the smooth functioning of the system. Failure to detect and recover from deadlocks can result in long wait times for processes and unresponsive systems, leading to user frustration and business losses.

In order to detect deadlocks, the system must monitor the state of all processes and resources in the system. This involves examining the current resource allocation state and predicting the future resource allocation state. Various methods can be used to detect deadlocks, such as the wait-for graph (WFG) and resource allocation graph (RAG) methods. Once a deadlock is detected, recovery techniques must be implemented to break the deadlock and restore the system to a functional state.


Deadlock Detection

Deadlock detection is a crucial task in computer systems that utilize shared resources. It involves identifying and flagging situations where two or more processes are blocked, unable to proceed because they are waiting for resources that are being held by other processes. When a deadlock is detected, the system can initiate recovery procedures to break the deadlock and restore the system to a functional state. It is essential to monitor and detect deadlocks as early as possible to prevent any negative impact on the system's performance and stability. A delay in detecting deadlocks can result in significant wait times for processes and unresponsive systems, leading to user frustration and potential business losses. There are two primary methods for detecting deadlocks: resource allocation graph (RAG) and wait-for graph (WFG).

Resource Allocation Graph (RAG)


The resource allocation graph (RAG) is a widely used method for deadlock detection in computer systems. The RAG is a graphical representation of the current allocation state of resources and the processes that are holding them. The nodes of the graph represent the resources and the processes, and the edges represent the allocation relationship between them. In the RAG method, a cycle in the graph indicates the presence of a deadlock. When a cycle is detected, it means that each process in the cycle is holding on to at least one resource that is

required by another process in the cycle, resulting in a deadlock. The RAG method is highly efficient and can quickly detect deadlocks, making it an essential technique in modern operating systems.



Wait-For Graph (WFG)

The wait-for graph (WFG) is a common method used in deadlock detection in computer systems. The WFG is a graphical representation of the dependencies between the processes and the resources that they are waiting for. In the WFG, the nodes represent processes, and the resources are represented as edges. Each edge points from the process that is waiting for a resource to the process that currently holds that resource. The WFG method can efficiently detect deadlocks by analysing the graph for cycles. If a cycle is found, it indicates that a set of processes is waiting for resources that are being held by other processes in the same set, resulting in a deadlock. The system can then take appropriate actions to break the deadlock, such as rolling back the transactions or aborting some of the processes.



Deadlock Recovery

Deadlock recovery is a critical process that is initiated after a deadlock has been detected in a computer system. This complex process involves a set of actions and procedures that are undertaken to resolve the deadlock by breaking the cycle of resource dependency between the processes involved in the deadlock. The recovery process usually entails the identification of

the process or processes that are causing the deadlock and releasing one or more of the resources that they are holding. This involves a careful analysis of the system's state to determine which resources can be safely released without compromising the system's integrity. There are four primary methods of deadlock recovery: process termination, resource pre-emption, priority inversion, and rollback.

Process Termination

Process termination is a simple method for resolving deadlocks. In this method, the operating system identifies the processes involved in the deadlock and terminates one or more processes. This frees up the resources held by the terminated processes, which can be used by the remaining processes to continue their execution. However, this method has some drawbacks, such as loss of data, abrupt termination of processes, and inconsistency in the system.

Resource Pre-emption

Resource pre-emption is a more complex method for resolving deadlocks. In this method, the operating system identifies the resources involved in the deadlock and selects one or more resources to be pre-empted. The resources are then taken away from the process holding them and allocated to the waiting processes. The pre-empted process is suspended until the required resources become available again. This method can cause delays in the execution of the pre-empted process and can result in a suboptimal allocation of resources.

Priority Inversion

Priority inversion is a method for resolving deadlocks in real-time systems. In this method, the priority of the processes is changed to avoid deadlock situations. The process holding the required resources is given a higher priority, and the process waiting for the resources is given a lower priority. This method can lead to the inversion of priorities, which can cause performance issues and degrade the performance of the system. Additionally, this method can also lead to starvation of lower priority processes, as higher priority processes can keep preempting the resources.

Rollback

Rollback is a method for resolving deadlocks that is commonly used in database systems. In this method, the system rolls back the transactions of the involved processes to a previous state where they were not deadlocked. This method requires that the system maintains a log of all the transactions and the state of the system at different points in time. The system can then roll

back the transactions to the previous state and re-execute them. This method can cause significant delays in the execution of the transactions and can result in a loss of data.

Conclusion

Deadlocks can cause significant performance and stability issues in a system, and it is crucial to detect and recover from deadlocks as soon as possible. Deadlock detection involves checking for the presence of deadlocks in the system using methods such as RAG and WFG. Deadlock recovery involves breaking the deadlock by releasing the resources held by the processes in the deadlock using methods such as process termination, resource pre-emption, priority inversion, and rollback. Each method has its advantages and disadvantages, and the choice of method depends on the nature of the system and the requirements of the application. Proper management of deadlocks can help ensure the smooth functioning of the system and prevent system failure.