

12 Deadlock

Deadlock is one area where there is a strong theory, but much of it is ignored in practice. Reason: solutions are expensive and/or require predicting the future.

Deadlock example: semaphores. Two processes: one does $P(x)$ followed by $P(y)$, the other does the reverse.

Deadlock: a situation where each of a collection of processes is waiting for something from other processes in the collection. Since all are waiting, none can provide any of the things being waited for.

These are relatively simple-minded cases. Things may be much more complicated:

- In general, don't know in advance how many resources a process will need. If only we could predict the future....
- Deadlock can occur over separate resources, as in semaphore example, or over pieces of a single resource, as in memory, or even over totally separate classes of resources (tape drives and memory). Deadlock can occur over anything involving waiting, for example messages in a pipe system. Hard for OS to control.

In general, there are four conditions for deadlock:

- Mutually exclusive access: resources cannot be shared.
- No preemption. Once given, a resource cannot be taken away.
- Hold and wait: processes can hold resources while they are waiting for more resources.
- There is a circularity in the graph of who has what and who wants what. This graph shows processes as circles, resources as squares, arrows from process to resource waited for, from resource to owning process.

Solutions to the deadlock problem fall into two general categories:

- Detection : determine when the system is deadlocked and then take drastic action. Requires termination of one or more processes in order to release their resources. Impractical unless the system has special support for aborting processes and recovering.
- Prevention : organize the system so that it is impossible for deadlock ever to occur. May lead to less efficient resource utilization in order to guarantee no deadlocks.

Deadlock prevention: must find a way to eliminate one of the four necessary conditions for deadlock:

- Don't allow exclusive access. This is not reasonable for many applications.
- Create enough resources so that there's always plenty for all.
- Don't allow waiting (phone company solution). This punts the problem back to the user.
- Allow preemption. E.g. pre-empt students' thesis disk space?
- Make process ask for everything at once. Either get them all or wait for them all. Tricky to implement: must be able to wait on many things without locking anything. Painful for process: may be difficult to predict future needs, so must make very wasteful use of resources.
- Banker's algorithm.
 1. state maximum resource needs in advance
 2. allocate resources when needed; wait when granting request would lead to deadlock
- Require *ordered* requests. E.g. ask for resource of type S, then resources of type T, etc. If all processes allocate resource types in same order, there can be no cyclic dependencies.
 1. Resources that are never held by a process at the same time can be requested in any order relative to each other.
 2. Finding a feasible ordering of resources requires design time knowledge of the needs of all processes.

In general, deadlock prevention is difficult, expensive or inefficient. Detection is also expensive and recovery requires special system support (process may be in arbitrary state).