Deadlock

(Chapter 32)

CS 4410
Operating Systems
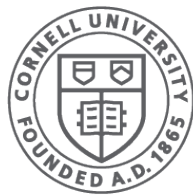
Dining Philosophers [Dijkstra 68]

Pi: **do forever**
    acquire(left(i));
    acquire(right(i));
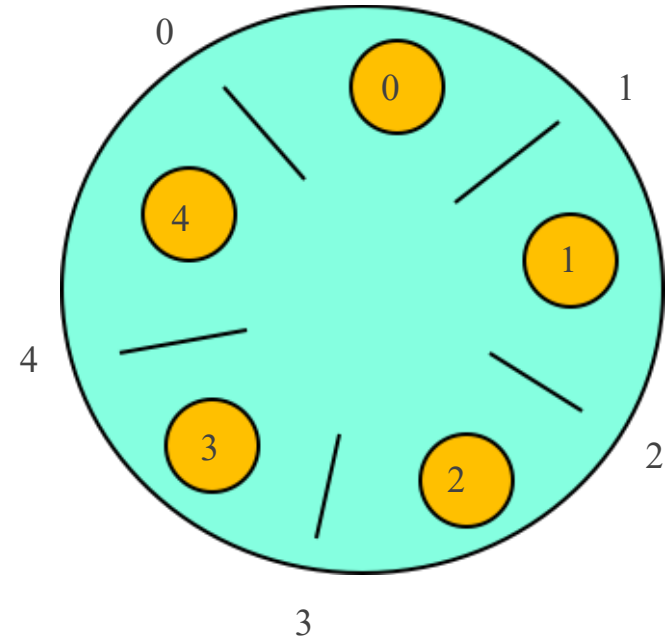   eat
    release(left(i));
    release(right(i));
end

left(i):   i+1 mod 5
right(i): i

Dining Philosophers in Harmony

```
1    import synch;
2
3    const N = 5;
4
5    def diner(which):
6        let left, right = (which, (which % N) + 1):
7            while choose({ False, True }):
8                lock(?forks[left]);
9                lock(?forks[right]);
10               # dine
11               unlock(?forks[left]);
12               unlock(?forks[right]);
13               # think
14           ;
15       ;
16   ;
17   forks = dict{ Lock() for i in {1..N} };
18   for i in {1..N}:
19       spawn diner(i);
20   ;
```

Dining Philosophers in Harmony

Issue: Non-terminating State		Shared Variable
Process		*forks*
init /()	0 865-879 584-587 880-884 869-879 :	dict{ 1:False, 2:False, 3:False, 4:False, 5:False }
diner/1	822-834 835(choose True) 836-841 58	dict{ 1:True, 2:False, 3:False, 4:False, 5:False }
diner/2	822-834 835(choose True) 836-841 58	dict{ 1:True, 2:True, 3:False, 4:False, 5:False }
diner/3	822-834 835(choose True) 836-841 58	dict{ 1:True, 2:True, 3:True, 4:False, 5:False }
diner/4	822-834 835(choose True) 836-841 58	dict{ 1:True, 2:True, 3:True, 4:True, 5:False }
diner/5	822-834 835(choose True) 836-841 58	dict{ 1:True, 2:True, 3:True, 4:True, 5:True }

14

Dining Philosophers in Harmony

Process	Status	Stack Trace	Variables
diner/1	blocked	848 diner(1): lock(?forks[right]); 593 lock(?forks[2]): await not tas(lk); 574 tas(?forks[2]): atomic:	result = () which = 1 left = 1 right = 2
diner/2	blocked	848 diner(2): lock(?forks[right]); 593 lock(?forks[3]): await not tas(lk); 574 tas(?forks[3]): atomic:	result = () which = 2 left = 2 right = 3
diner/3	blocked	848 diner(3): lock(?forks[right]); 593 lock(?forks[4]): await not tas(lk); 574 tas(?forks[4]): atomic:	result = () which = 3 left = 3 right = 4
diner/4	blocked	848 diner(4): lock(?forks[right]); 593 lock(?forks[5]): await not tas(lk); 574 tas(?forks[5]): atomic:	result = () which = 4 left = 4 right = 5
diner/5	blocked	848 diner(5): lock(?forks[right]); 593 lock(?forks[1]): await not tas(lk); 574 tas(?forks[1]): atomic:	result = () which = 5 left = 5 right = 1

Problematic Emergent Properties

Starvation: Process waits forever

Deadlock: A set of processes exists, where each is blocked and can become unblocked only by actions of another process in the set.

- Deadlock implies Starvation (but not *vice versa*)

- Starvation often tied to **fairness**: A process is not forever blocked awaiting a condition that (i) becomes continuously true or (ii) infinitely-often becomes true.

Testing for starvation or deadlock is difficult in practice

More Examples of Deadlock

Example (initially in1 = in2 = False):

    in1 = True; **await not** in2; in1 = False;
    //
    in2 := True; **await not** in1; in2 = False;

Example (initially semaphores m1 = m2 = 1):

    P(m1); P(m2); V(m1); V(m2);
    //
    P(m2); P(m1); V(m1); V(m2);

System Model

- Set of resources requiring "exclusive" access
    - Might be "k-exclusive access" if resource has capacity for k
    - Examples: buffers, packets, I/O devices, processors, …

- Protocol to access a resource causes blocking:
    - If resource is free then access is granted; process proceeds
    - If resource is in use then process blocks
    - Use resource
    - Release resource

When is deadlock possible?

Necessary Conditions for Deadlock

1. **Mutual Exclusion**. Acquire can block invoker

2. **Hold & wait**. A process can be blocked while holding resources

3. **No preemption**. Allocated resources cannot be reclaimed. Explicit release operation needed

4. **Circular waits** are possible

   Let p → q denote "p waits for q to release a resource". Then

   P1 → P2 → … → Pn → P1

Deadlock is Undesirable

- Deadlock <u>prevention</u>: Ensure that a necessary condition cannot hold

- Deadlock <u>avoidance</u>: System does not allocate resources that will lead to a deadlock

- Deadlock <u>detection</u>: Allow system to deadlock; detect it; recover

Deadlock Prevention: Negate 1

#1: Eliminate mutual exclusion / bounded resources:

- Make resources sharable without locks
  - Harmony book Chapter 22 has examples of non-blocking data structures
- Have sufficient resources available, so acquire never delays
 - E.g., unbounded queue, or simply make sure bounded queue is "large enough"

Deadlock Prevention: Negate 2

#2: **Eliminate hold and wait**

Don't hold some resources when waiting for others.

- Re-write code:

```
def foo():
    lock(?mutex);
    doSomeStuff();
    bar();
    doOtherStuff();
    unlock(?mutex);
```

➡

```
def foo():
    lock(?mutex);
    doSomeStuff();
    unlock(?mutex);
    bar();
    lock(?mutex);
    doOtherStuff();
    unlock(?mutex);
```

- Or request all resources before execution begins
  - Problems:
    - Processes don't know what they need ahead of time.
    - Starvation (if waiting on many popular resources).
    - Low utilization (need resource only for a bit).

Deadlock Prevention: Negate 1 of 4

#3: Allow preemption

Requires mechanism to save / restore resource state:

multiplexing   vs   undo/redo

- Examples of multiplexing:
  - processor registers
  - Regions of memory (pages)

- Examples of undo/redo
  - Database transaction processing

#4: Eliminate circular waits.

Let R = {R1, R2, ... Rn} be the set of resource types.
Let (R , <) be a non-symmetric relation:

- not r < r [irreflexive]
- if r < s and s < t then r < t [transitive]
- not r < s and s < r [non-symmetric]
- for every r and s (r ≠ s): r < s or s < r [total order]

Rule: Request resources in increasing order by <
(All resources from type Ri must be requested together)

Rule: To request the resources of type Rj, first release all
resources from type Ri where Ri < Rj.

Why < Rules Work

Thm: Total order resource allocation avoids circular waits

Proof: By contradiction. Assume a circular wait exists
$$P1 \rightarrow P2 \rightarrow P3 \rightarrow \dots \rightarrow Pn \rightarrow P1.$$
P1 requesting R1 held by P2.
P2 requesting R2 held by P3. (So R1 < R2 holds)
…

Conclude: R1 < R2, R2 < R3, …, Rn < R1
By transitivity: R1 < R1. A contradiction!

Havender's Scheme (OS/360)

Hierarchical Resource Allocation

Every resource is associated with a level.

- **Rule H1**: All resources from a given level must be acquired using a single request.
- **Rule H2**: After acquiring from level L_j must not acquire from L_i where $i < j$
- **Rule H3**: May not acquire from L_i unless already released from L_j where $j > i$.

Example of allowed sequence:
1. acquire(W@L1, X@L1)
2. acquire(Y@L3)
3. release(Y@L3)
4. acquire(Z@L2)

acquire

release

L_1

L_2

L_n
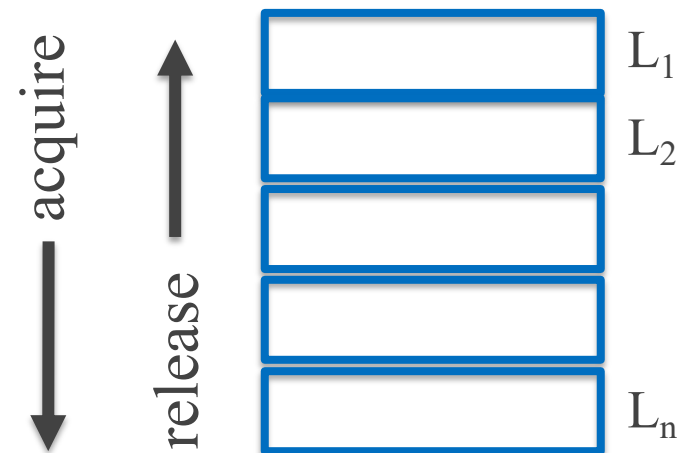
Dining Philosophers (Again)

Pi: **do forever**
    acquire(F(i));
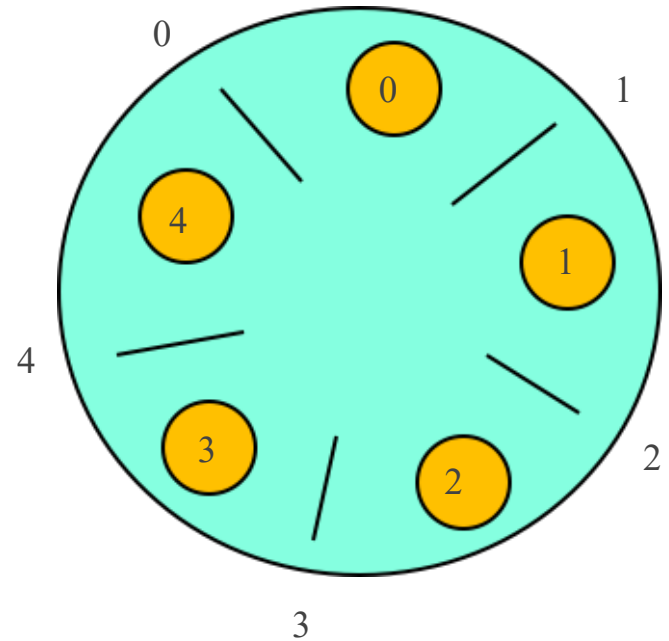   acquire(G(i));
    eat
    release(F(i));
    release(G(i));
    end

F(i):   min(i, i+1 mod 5)
G(i):   max(i, i+1 mod 5)

Ordering Resources in Harmony

```
def diner(which):
    let left, right = which, (which % N) + 1:
        if left < right:
            lock(?forks[left]);
            lock(?forks[right]);
        else:
            lock(?forks[right]);
            lock(?forks[left]);
        ;
        # dine
        unlock(?forks[left]);
        unlock(?forks[right]);
        # think
    ;
;
```

Simultaneous Acquisition in Harmony

```
5    def diner(which):
6      let left, right = (which, (which % N) + 1):
7        while choose({ False, True }):
8          lock(?mutex);
9          while forks[left] or forks[right]:
10           if forks[left]:
11             wait(?conds[left]);
12             ;
13           if forks[right]:
14             wait(?conds[right]);
15             ;
16           ;
17         assert not (forks[left] or forks[right]);
18         forks[left] = True;
19         forks[right] = True;
20         unlock(?mutex);
21
22         # dine
23
24         lock(?mutex);
25         forks[left] = False;
26         forks[right] = False;
27         notify(?conds[left]);
28         notify(?conds[right]);
29         unlock(?mutex);
30
31         # think
32         ;
33       ;
34     ;
```

```
35   mutex = Lock();
36   forks = dict{ False for i in {1..N} };
37   conds = dict{ Condition(?mutex) for i in {1..N} };
```

Simultaneous Acquisition in Harmony

```
5      def diner(which):
6         let left, right = (which, (which % N) + 1):
7            while choose({ False, True }):
8               lock(?mutex);
9               while forks[left] or forks[right]:
10                 if forks[left]:
11                    wait(?conds[left]);
12                 ;
13                 if forks[right]:
14                    wait(?conds[right]);
15                 ;
16              ;
17              assert not (forks[left] or forks[right]);
18              forks[left] = True;
19              forks[right] = True;
20              unlock(?mutex);
21
22              # dine
23
24              lock(?mutex);
25              forks[left] = False;
26              forks[right] = False;
27              notify(?conds[left]);
28              notify(?conds[right]);
29              unlock(?mutex);
30
31              # think
32           ;
33        ;
34     ;
```

```
35     mutex = Lock();
36     forks = dict{ False for i in {1..N} };
37     conds = dict{ Condition(?mutex) for i in {1..N} };
```

30

Simultaneous Acquisition in Harmony

```
5    def diner(which):
6        let left, right = (which, (which % N) + 1):
7            while choose({ False, True }):
8                lock(?mutex);
9                while forks[left] or forks[right]:
10                   if forks[left]:
11                       wait(?conds[left]);
12                   ;
13                   if forks[right]:
14                       wait(?conds[right]);
15                   ;
16               ;
17               assert not (forks[left] or forks[right]);
18               forks[left] = True;
19               forks[right] = True;
20               unlock(?mutex);
21
22               # dine
23
24               lock(?mutex);
25               forks[left] = False;
26               forks[right] = False;
27               notify(?conds[left]);
28               notify(?conds[right]);
29               unlock(?mutex);
30
31               # think
32           ;
33       ;
34   ;
```
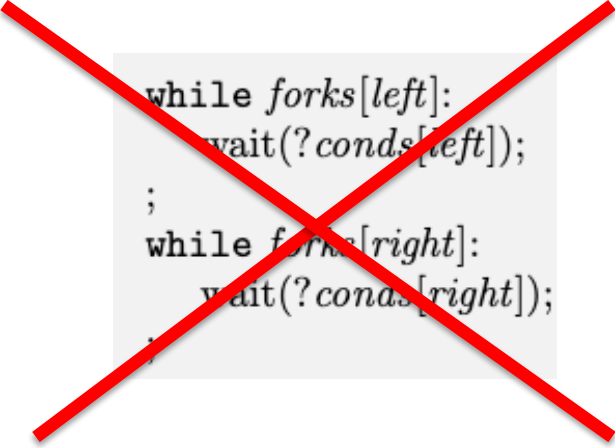
```
35   mutex = Lock();
36   forks = dict{ False for i in {1..N} };
37   conds = dict{ Condition(?mutex) for i in {1..N} };
```

```
while forks[left]:
    wait(?conds[left]);
;
while forks[right]:
    wait(?conds[right]);
;
```

Simultaneous Acquisition in Harmony

```
5    def diner(which):
6        let left, right = (which, (which % N) + 1):
7            while choose({ False, True }):
8                lock(?mutex);
9                while forks[left] or forks[right]:
10                   if forks[left]:
11                       wait(?conds[left]);
12                   ;
13                   if forks[right]:
14                       wait(?conds[right]);
15                   ;
16               ;
17               assert not (forks[left] or forks[right]);
18               forks[left] = True;
19               forks[right] = True;
20               unlock(?mutex);
21
22               # dine
23
24               lock(?mutex);
25               forks[left] = False;
26               forks[right] = False;
27               notify(?conds[left]);
28               notify(?conds[right]);
29               unlock(?mutex);
30
31               # think
32           ;
33       ;
34   ;
```

```
35   mutex = Lock();
36   forks = dict{ False for i in {1..N} };
37   conds = dict{ Condition(?mutex) for i in {1..N} };
```
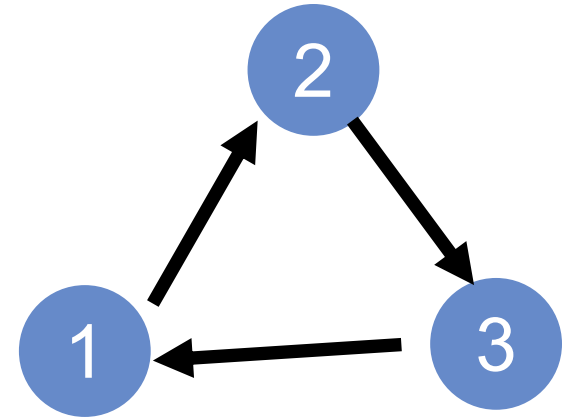
```
while forks[left]:
    wait(?conds[left]);
;
while forks[right]:
    wait(?conds[right]);
;
```

Deadlock Detection

Create a [Wait-For Graph](#)

- 1 Node per Process
- 1 Edge per Waiting Process, P
(from P to the process it's waiting for)

Note: graph holds for a single instant in time

Cycle in graph indicates deadlock

Testing for cycles (= deadlock)

Reduction Algorithm:
Find a node with no outgoing edges
- Erase node
- Erase any edges coming into it

Intuition: Deleted node is for process that is not waiting. It will eventually finish and release its resources, so any process waiting for those resources will longer be waiting.

Erase whole graph ↔ graph has no cycles
Graph remains ↔ deadlock

Graph Reduction: Example 1

Find node w/o outgoing edges
Erase node
Erase edges coming into it
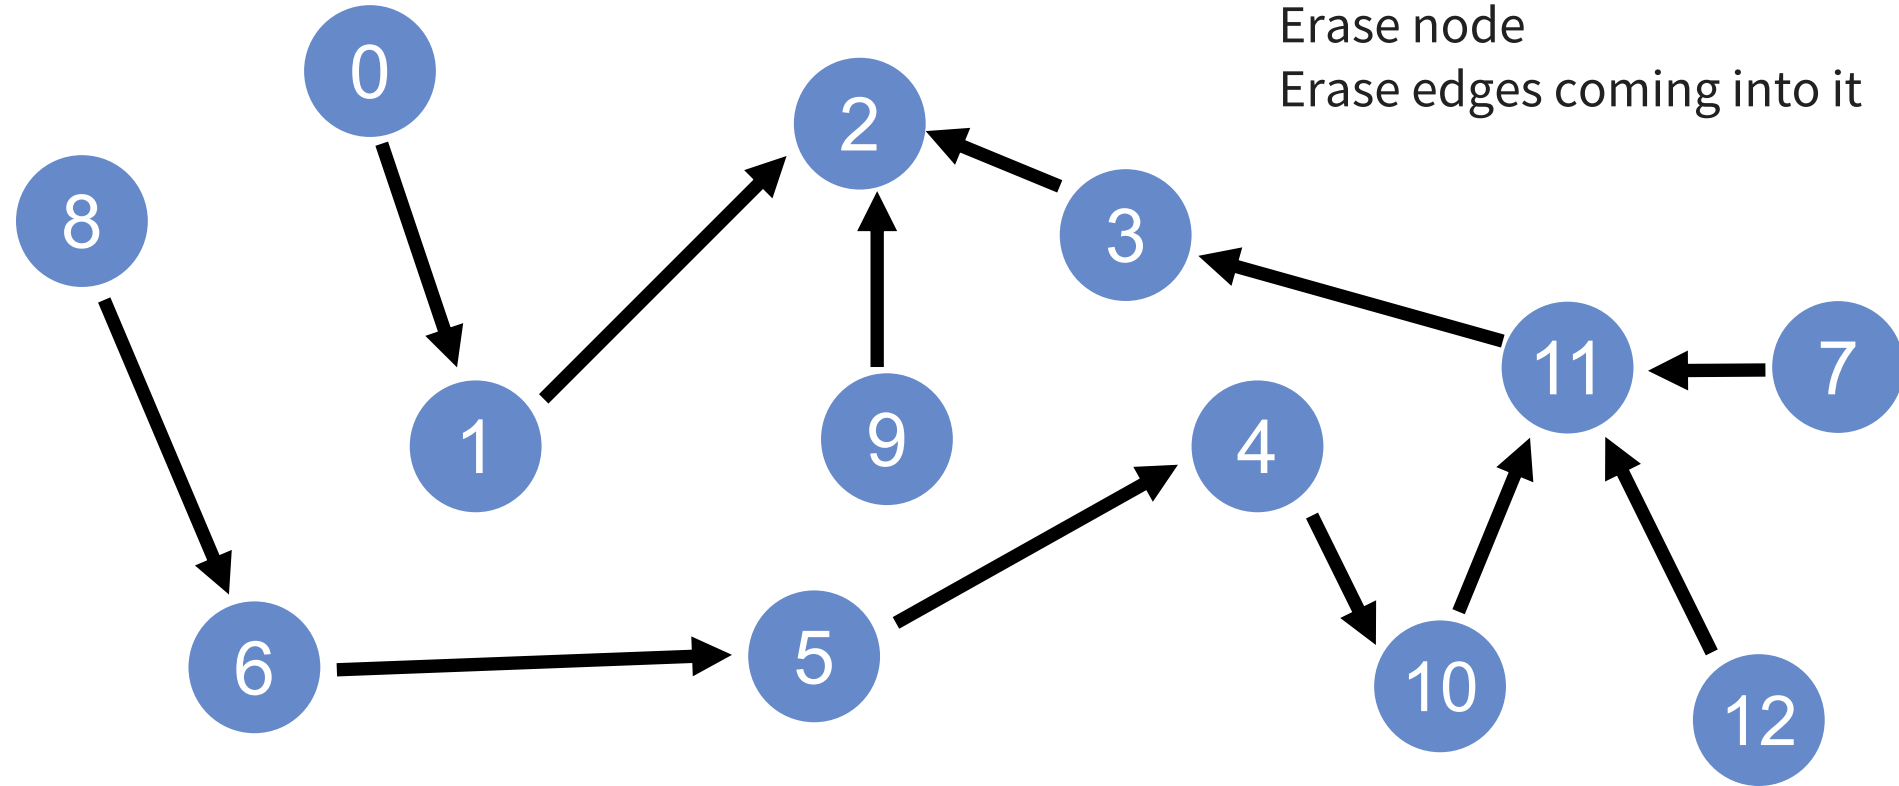
Graph can be fully reduced, hence there was no deadlock at the time the graph was drawn. (Obviously, things could change later!)

Graph Reduction: Example 2
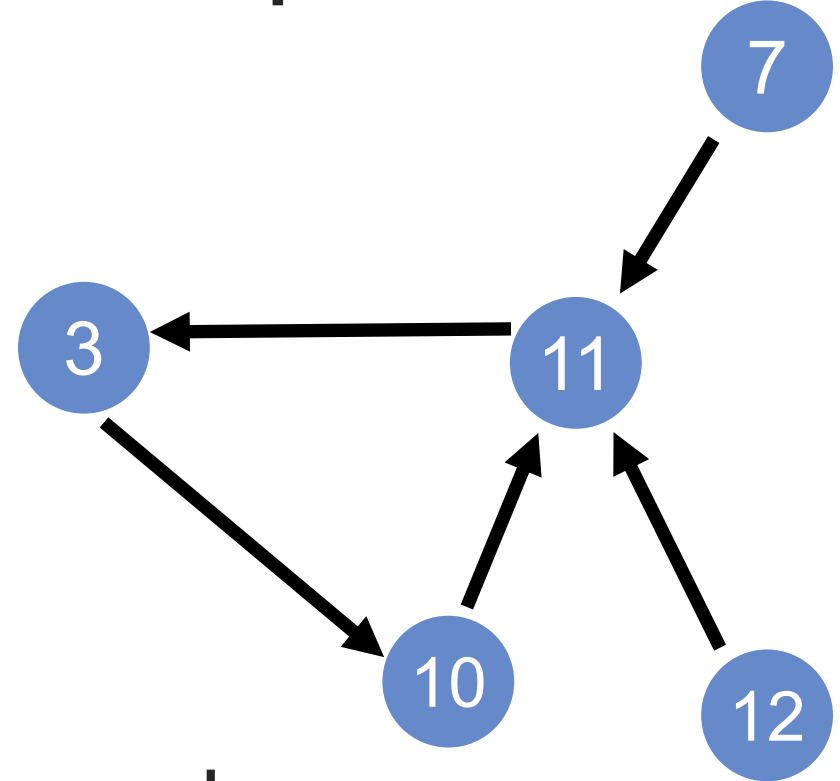
No node with no outgoing edges…
Irreducible graph, contains a cycle
    (only some processes are in the cycle)
➤ deadlock

Question:

Does choice of node for reduction matter?

Answer: No.
Explanation: an unchosen candidate at one step remains a candidate for later steps. Eventually—regardless of order—every node will be reduced.

Question:

Suppose no deadlock detected at time T.
Can we infer about a later time T+x?

Answer: Nothing.
Explanation: The very next step could be to run some process that will request a resource…
	… establishing a cyclic wait
	… and causing deadlock

Implementing Deadlock Detection

- Track resource allocation (who has what)
- Track pending requests (who's waiting for what)

Maintain a wait-for graph.

When to run graph reduction?

- Whenever a request is blocked?
- Periodically?
- Once CPU utilization drops below a threshold?

Deadlock Recovery

Blue screen & reboot?

Kill one/all deadlocked processes
- Pick a victim
- Terminate
- Repeat if needed

Preempt resource/processes till deadlock broken
- Pick a victim (# resources held, execution time)
- Rollback (partial or total, not always possible)

Deadlock Avoidance

How do cars do it?
- Try not to block an intersection
- Don't drive into the intersection if you can see that you'll be stuck there.

Why does this work?
- Prevents a wait-for relationship
- Cars won't take up a resource if they see they won't be able to acquire the next one…

Deadlock Avoidance

state: allocation to each process
safe state: a state from which some execution is possible that does not cause deadlock.

- Requires knowing max allocation for each process.
- Check that
  - Exists sequence P1 P2 … Pn of processes where:
      Forall i where $1 \leq i \leq n$:
          Pi can be satisfied by Avail + resources held by P1 … Pi-1.

Assumes no synchronization between processes, except for resource requests.

Safe State Example

Suppose: 12 tape drives and 3 processes: p0, p1, and p2

	max need	current usage	could still ask for
p0	10	5	5
p1	4	2	2
p2	9	2	7

3 drives remain

Current state is *safe* because a safe sequence exists: [p1, p0, p2]
- p1 can complete with remaining resources
- p0 can complete with remaining+p1
- p2 can complete with remaining+p1+p0

What if p2 requests 1 drive? Grant or not?

Safe State Example

Suppose: 12 tape drives and 3 processes: p0, p1, and p2

	max need	current usage	could still ask for
p0	10	5	5
p1	4	2	2
p2	9	3	6

2 drives remain

Is this state safe? (Is there a sequence of requests that works?)

Safe State Example

Suppose: 12 tape drives and 3 processes: p0, p1, and p2

	max need	current usage	could still ask for
p0	10	5	5
~~p1~~	~~0~~	~~0~~	~~0~~
p2	9	3	6

4 drives remain

Is this state safe? (Is there a sequence of requests that works?)

Safe State Example

Suppose: 12 tape drives and 3 processes: p0, p1, and p2

	max need	current usage	could still ask for
p0	10	5	5
~~p1~~	~~0~~	~~0~~	~~0~~
p2	9	3	6

`4 drives remain`

Is this state safe? (Is there a sequence of requests that works?)

STUCK…
(non-terminating state)

Safe State Example

Suppose: 12 tape drives and 3 processes: p0, p1, and p2

	max need	current usage	could still ask for
p0	10	5	5
p1	4	2	2
p2	9	2	7

3 drives remain

Current state is *safe* because a safe sequence exists: [p1, p0, p2]
- p1 can complete with remaining resources
- p0 can complete with remaining+p1
- p2 can complete with remaining+p1+p0

What if p2 requests 1 drive? Grant or not?

Safe State Example

Suppose: 12 tape drives and 3 processes: p0, p1, and p2

	max need	current usage	could still ask for
p0	10	5	5
p1	4	2	2
p2	9	2	7

3 drives remain

Current state is *safe* because a safe sequence exists: [p1, p0, p2]
- p1 can complete with remaining resources
- p0 can complete with remaining+p1
- p2 can complete with remaining+p1+p0

What if p2 requests 1 drive? Grant or not?   NO

Banker's Algorithm

- from 10,000 feet:
  - Process declares its worst-case needs, asks for what it "really" needs, a little at a time
  - Algorithm decides when to grant requests
    - Build a graph assuming request granted
    - Reducible? yes: grant request, no: wait

Problems:
  - Fixed number of processes
  - Need worst-case needs ahead of time
  - Expensive

Avoiding Deadlock in Harmony

```
5     def diner(which):
6         let left, right = (which, (which % N) + 1):
7             while choose({ False, True }):
8                 P(?sema);                          ⬅
9                 lock(?forks[left]);
10                lock(?forks[right]);
11                # dine
12                unlock(?forks[left]);
13                unlock(?forks[right]);
14                V(?sema);
15                # think
16            ;
17          ;
18        ;
19     forks = dict{ Lock() for i in {1..N} };
20     sema = Semaphore(N − 1);                      ⬅
```