Distributed Computing and Systems
Chalmers university of technology

Prof Philippas Tsigas

Distributed Computing and Systems Research Group

DISTRIBUTED SYSTEMS II

RESOURCE ALLOCATION: DRINKING PHILOSOPHERS

Conflict -> Precedence Graph

Undirected GRAPH, in which edges represent shared resources between processes we call this graph CONFLICT GRAPH.

The algorithm by Chandy and Misra resolves conflicts by defining for every possible conflict a *precedence* relation:

- When two processes compete for a resource the one with higher precedence may access the resource first.

- In order to receive a solution which is fair these precedences will have to change dynamically.

The directed graph graph that changes dynamically is called *precedence graph*.

For each resource an edge of the precedence graph is directed from processes with lower precedence to processes with higher precedence.

The precedences of the graph are chosen such that it is always possible to distinguish at least one process from all other processes i.e. this process can enter its critical section. (NO DEADLOCK)

This is ensured by the existence of at least one process which has higher precedence for all its shared resources. A process with this property is called *sink*.

Its existence is guaranteed when the precedence graph is always acyclic.

By changing directions of edges it is possible to change the precedences dynamically.

This must happen in a way that the precedence graph stays acyclic, so *progress*, *fairness* and *mutual exclusion* is guaranteed.

Starting with a DAG

- The graph is initialised acyclic for example by a node-colouring algorithm.

- The graph can remain acyclic if after use of the critical section a process reverse all adjacent precedences in one step.

- Need a mechanism to keep the sense of direction:

The mechanism

Forks which have the property to be either *clean* or *dirty*.

- A fork will be cleanedbefore it is send to a neighbour process.

- A clean fork will become dirty when the holder of the resource enters the critical section.

- After use it remains DIRTY until it is sent to a neighbour process.

The dynamic DAG

- The respective precedence graph H can be defined in the following way:

- For all pairs of processes p and q which share a common resource, $<p,q>$ one of the following statements is true:

1. p holds the *fork* for the resource and the *fork* is CLEAN

2. q holds the *fork* for the resource and the *fork* is DIRTY

3. the *fork* for the resource is in transit from q to p

Requesting Forks

The request of forks is realized by *request tokens*.

For each fork there exist one request token such that only the holder of the request token can request a fork.

A *hungry* process requests a fork by sending the *request* TOKEN to the owner of the desired *fork*.

A process is not interested in accessing its resources when it holds a *request* TOKEN but not a *fork*.

The algorithm

The algorithm is initialised by an acyclic precedence graph H and all processes with lower precedence own dirty forks while processes with higher precedence own request tokens.

All processes are *thinking* i.e they are not interested in their resources.

A process which becomes *hungry* will send all its *request* TOKEN to neighbour processes and wait until it received all *forks*.

- A process which received all forks will change its state to *eating*.

- A process which leaves the CRITICAL SECTION changes the state of all its *forks* to DIRTY. Then for all held *request* TOKEN the respective *fork* is sent to neighbour processes.

The above steps assume following rules:

Receiving a *request* TOKEN for fork f:

1. If processors state is different from *eating* and f is DIRTY then f will be sent to the requesting processor.

2. If processors state was also *hungry* then the *request* TOKEN will also be sent back.

Receiving a fork f: The state of f will be set to clean.

- if diner does not have fork – it sends request to the neighbor

- process gives up dirty forks and holds clean

(R1) Requesting a fork f:
$hungry, reqf(f), \sim fork(f) \rightarrow$
send request token for fork f (to the philosopher with whom f is shared);
$reqf(f) := false$

(R2) Releasing a fork f:
$\sim eating, reqf(f), dirty(f) \rightarrow$
send fork f (to the philosopher with whom fork f is shared);
$dirty(f) := false;$
$fork(f) := false$

(R3) Receiving a request token for f:
upon receiving a request for fork $f \rightarrow$
$reqf(f) := true$

(R4) Receiving a fork f:
upon receiving fork $f \rightarrow$
$fork(f) := true$
$\{\sim dirty(f)\}$

- diner can only wait for higher priority neighbor

- notice after eating, diner is a source (all his forks are dirty)
- graph remains acyclic throughout operation
- acyclic graph has a sink (lowest priority diner), that diner is guaranteed to eat

Correctness

Mutual Exclusion:

Proof. The precedence graph H is acyclic. □

No Starvation

Proof. Let the depth in H of any process p be defined as the maximum number of edges along a path from p to another process without predecessor. The proof will show by induction that a process of depth k will eventually eat if predecessors at depth k-1 can EAT. □

Complexities

Communication Complexity: *O(degree)*

Proof. A process sends at most one *request* TOKEN to each neighbour and receives from each neighbour at most one *fork*. □
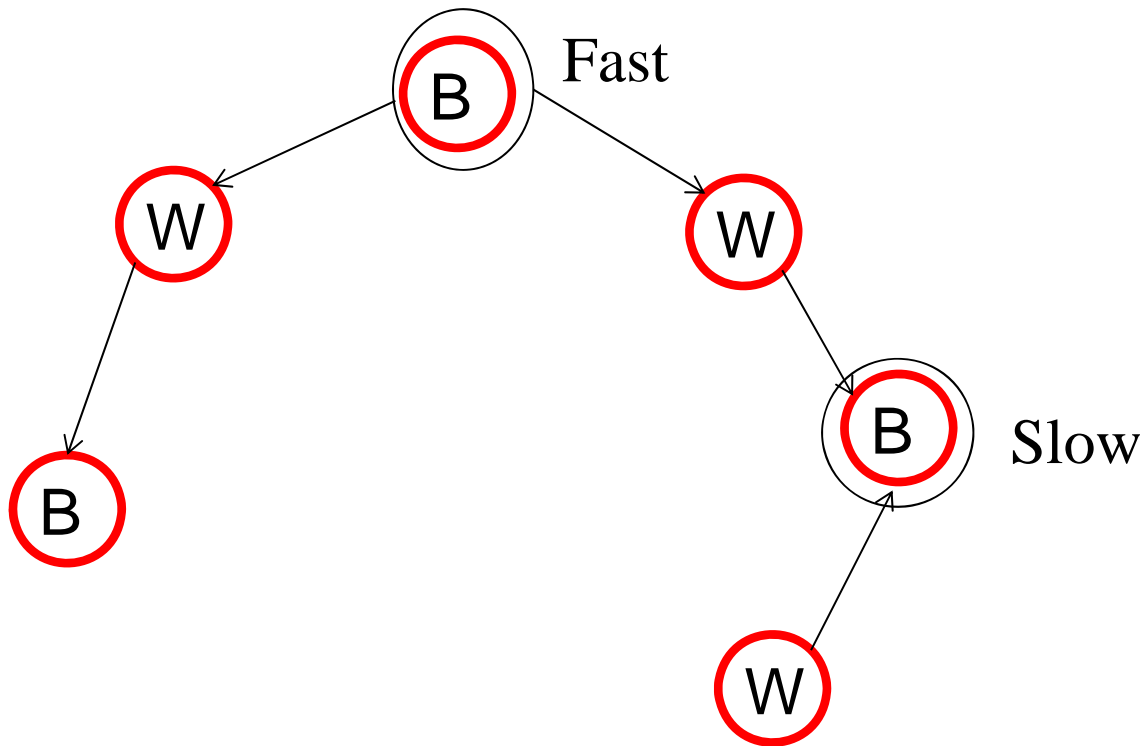
**Tine** Complexity: *O(n)*

Time Complexity

Max distance from sink = 1

Max distance from sink = 2

Max distance from sink = 2

Fast

Slow

Max distance from sink = 2

Max distance from sink = 3