Distributed Deadlock-Avoidance

César Sánchez

IMDEA Software Institute, Spain



DRV Workshop, Bertinoro

19-May, 2016



IMDEA Software Institute, Spain



DRV Workshop, Bertinoro

19-May, 2016

Introduction

Goal: Formalization of middleware services



Deadlocks

Deadlock is one of the classical problems in CS

One (common) approach is the ostrich approach

The other approaches are: detection, prevention and avoidance.

| | centralized | distributed |
|------------|-------------|-------------|
| detection | OK | OK |
| prevention | OK | OK |
| avoidance | Banker's | impractical |

Efficient dynamic resource allocation can have a big practical impact.



A deadlock state



Detection:



Prevention:



Avoidance:



Distributed Avoidance:





























Sequence of calls:



Distributed Real-Time Embedded Systems

Distributed Real-Time Embedded Systems:

- Asynchronous distributed system
- Limited Resources
- ► Wait-on-connection
- Arbitrary number of processes spawned
- All processes terminate

Problem: deadlocks are possible if no controller is used

Two sites, with two resources each:





Two sites, with two resources each:







Two sites, with two resources each:







Two sites, with two resources each:







Two sites, with two resources each:





Two sites, with two resources each:





Two sites, with two resources each:







Summary of Contributions

Contribution #1

Efficient deadlock Avoidance can is possible provided call-graphs are know statically

Summary of Contributions

Contribution #1

Efficient deadlock Avoidance can is possible provided call-graphs are know statically

Contribution #2 Optimal annotations can be efficiently computed. If annotations are not followed *anomalies* can occurr.

Summary of Contributions

Contribution #1 Efficient deadlock Avoidance can is possible provided call-graphs are know statically

Contribution #2 Optimal annotations can be efficiently computed. If annotations are not followed *anomalies* can occurr.

Contribution #3 Distributed Deadlock Avoidance with (individual) liveness guarantees can be efficiently achieved.

Model of Computation

- Remote procedure call (with Wait-On-Connection)
- Asynchronous messages
- All to all communication
- Finite resources: T_A total number of threads



Model of Computation

- Remote procedure call (with Wait-On-Connection)
- Asynchronous messages
- All to all communication
- Finite resources: T_A total number of threads



We seek a *deadlock avoidance* solution with no extra communication

Distributed Deadlock Avoidance Solution

Two parts:

1. Static:

2. Dynamic:

Distributed Deadlock Avoidance Solution

Two parts:

1. Static:



2. Dynamic:
Distributed Deadlock Avoidance Solution

Two parts:

1. Static:



2. Dynamic:



Annotations are computed statically



Annotations are computed statically



Dependency edges $n \rightarrow m$ whenever $\alpha(n) \leq m$ for two calls in the same node.

Annotations are computed statically



Dependency edges $n \rightarrow m$ whenever $\alpha(n) \leq m$ for two calls in the same node.



Annotations are computed statically



Dependency edges $n \rightarrow m$ whenever $\alpha(n) \leq m$ for two calls in the same node.



 $\blacktriangleright \ n$ depends on m if there is a path from n to m containing a \rightarrow

Annotations are computed statically



Dependency edges $n \rightarrow m$ whenever $\alpha(n) \leq m$ for two calls in the same node.



 $\blacktriangleright\ n$ depends on m if there is a path from n to m containing a \rightarrow

Annotations are computed statically



Dependency edges $n \rightarrow m$ whenever $\alpha(n) \leq m$ for two calls in the same node.



 $\blacktriangleright n$ depends on m if there is a path from n to m containing a \rightarrow

Annotations are computed statically



Dependency edges $n \rightarrow m$ whenever $\alpha(n) \leq m$ for two calls in the same node.



 $\blacktriangleright n$ depends on m if there is a path from n to m containing a \rightarrow

Basic Solution Deadlock Avoidance

Protocol BASIC-P: $\begin{array}{c}
 \alpha \\
 \hline
 n_1 & A
\end{array}$ $\begin{bmatrix}
 \mathbf{when} \ \alpha < t_A \ \mathbf{do} \\
 t_A - \end{bmatrix}$ $\begin{bmatrix}
 n_1() \\
 t_A + +
\end{bmatrix}$

Basic Solution Deadlock Avoidance



Theorem: If α has no cyclic dependencies, then BASIC-P guarantees absence of deadlock.

Basic Solution Deadlock Avoidance



Theorem: If α has no cyclic dependencies, then BASIC-P guarantees absence of deadlock.

Lemma: The following is an **invariant**:

The number of processes running methods with annotation ior higher is at most $T_A - i$.

The Annotation Theorem

Theorem: If α has no cyclic dependencies, then BASIC-P guarantees absence of deadlock.

Lemma: The following is an **invariant**:

The number of processes running methods with annotation ior higher is at most $T_A - i$.

| annotation $lpha$ | Max num of procs |
|-------------------|------------------|
| 0 | T_A |
| 1 | $T_A - 1$ |
| | |
| $T_A - 1$ | 1 |

The Annotation Theorem

Theorem: If α has no cyclic dependencies, then BASIC-P guarantees absence of deadlock.

Lemma: The following is an **invariant**:

The number of processes running methods with annotation i or higher is at most $T_A - i$.

Lemma: If a request $\boxed{n_1 A}^{\alpha}$ is disabled, then there is an active process running $\boxed{n_2 A}^{\alpha_2}$ with $\alpha_2 \leq \alpha$.

Two immediate questions:

1. How to compute acyclic annotations

Two immediate questions:

- 1. How to compute acyclic annotations
- Visit nodes following some reverse topological order.

- When visiting n, compute the set of nodes S previously visited and reachable following $(\rightarrow \cup \cdots)^*$.

- Set $\alpha(n)$ to 1 plus the largest node in S that resides in the same site.

Two immediate questions:

- 1. How to compute acyclic annotations
- Visit nodes following some reverse topological order.

- When visiting n, compute the set of nodes S previously visited and reachable following $(\rightarrow \cup \cdots)^*$.

- Set $\alpha(n)$ to 1 plus the largest node in S that resides in the same site.



Two immediate questions:

- 1. How to compute acyclic annotations
- Visit nodes following some reverse topological order.

- When visiting n, compute the set of nodes S previously visited and reachable following $(\rightarrow \cup \cdots)^*$.

- Set $\alpha(n)$ to 1 plus the largest node in S that resides in the same site.



Consider two nodes, with two resource each $(T_A = T_B = 2)$:





Consider two nodes, with two resource each $(T_A = T_B = 2)$:





Consider two nodes, with two resource each $(T_A = T_B = 2)$:





Consider two nodes, with two resource each $(T_A = T_B = 2)$:





Consider two nodes, with two resource each $(T_A = T_B = 2)$:





Consider two nodes, with two resource each $(T_A = T_B = 2)$:





Consider two nodes, with two resource each $(T_A = T_B = 2)$:





Consider two nodes, with two resource each $(T_A = T_B = 2)$:





Consider two nodes, with two resource each $(T_A = T_B = 2)$:





Revisiting the Invariant

Lemma: The following is an **invariant**:

The number of processes running methods with annotation i or higher is at most $T_A - i$.

$$act_{A,\geq i} \leq T_A - i$$
 for all notes A and i

where

 $act_{A,i}$: number of active processes in A with annotation i $act_{A,\geq i} = \sum_{k\geq i} act_{A,i}$

Revisiting the Invariant

Lemma: The following is an **invariant**:

The number of processes running methods with annotation i or higher is at most $T_A - i$.

$$act_{A,\geq i} \leq T_A - i$$
 for all notes A and i

where

 $act_{A,i}$: number of active processes in A with annotation i $act_{A,\geq i} = \sum_{k\geq i} act_{A,i}$

The weakest precondition on allowing a request for

 $\begin{bmatrix} n & A \end{bmatrix}^{i}$

$$\varphi' = \bigwedge_{k} \begin{cases} act_{A,\geq k} \leq T_A - k & \text{if } k > i \\ act_{A,\geq k} + 1 \leq T_A - k & \text{if } k \leq i \end{cases}$$



Theorem (Deadlock Avoidance): If α is acyclic, then LIVE-P guarantees absence of deadlock.

Theorem (Liveness): If α is acyclic, then LIVE-P guarantees that every waiting process is eventually enabled.

Consider two nodes, with two resource each $(T_A = T_B = 2)$:



BASIC-P







Consider two nodes, with two resource each $(T_A = T_B = 2)$:



Basic-P







Consider two nodes, with two resource each $(T_A = T_B = 2)$:



Basic-P







Consider two nodes, with two resource each $(T_A = T_B = 2)$:





BASIC-P







Consider two nodes, with two resource each $(T_A = T_B = 2)$:



Basic-P







Consider two nodes, with two resource each $(T_A = T_B = 2)$:



Basic-P







Conclusions

- Distributed Deadlock Avoidance is possible without communication
- provided call-graphs are known
- Using static annotations + runtime protocols
- If cycles are allowed (e.g. by uncontrolled resource allocation), then deadlocks are unavoidable, provided enough resources
- Individual liveness is also enforceable
- ► Future work:
 - is deadlock avoidance enforceable for any amount of initial resources?
 - can this be adapted to composable conveyor systems?
Conclusions

- Distributed Deadlock Avoidance is possible without communication
- provided call-graphs are known
- Using static annotations + runtime protocols
- If cycles are allowed (e then deadlocks are una
- Individual liveness is al
- Future work:
 - is deadlock avoidan resources?
 - can this be adapted to composable conveyor systems?



Distributed Dinning Philosphers



Distributed Dinning Philosphers



Distributed Dinning Philosphers



do not the take last fork if going in increasing order. For your second pick,

do as you wish.



Thank you for your attention!