Deadlocks and dihomotopy in mutual exclusion models

Martin Raussen

Department of mathematical sciences
Aalborg Universitet

MSRI-seminar, 9.10.2006

Mutual exclusion

Mutual exclusion occurs, when n processes P_i compete for m resources S_j.

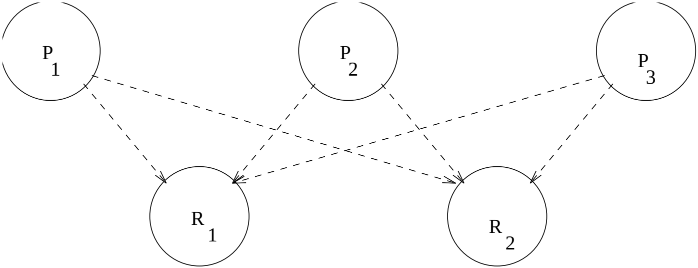

Only k processes can be served at any given time.
Semaphores!
Semantics: A processor has to lock a resource and relinquish the lock later on!
Description/abstraction $P_i : \dots PR_j \dots VR_j \dots$ (Dijkstra)

Schedules in ”progress graphs”

The Swiss flag example

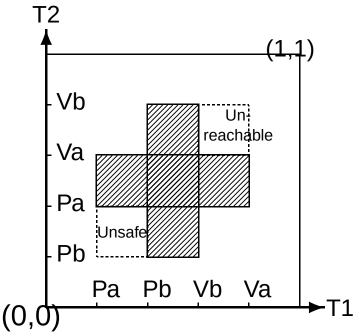

$P_1 : P_a P_b V_b V_a \qquad P_2 : P_b P_a V_a V_b$

Executions are directed paths avoiding a forbidden region F.

Deadlocks: no directed path with that source.

Unsafe regions: Every inextendible dipath ends in a deadlock.

Deadlocks in higher dimensions

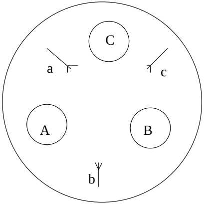

A=Pa.Pb.Va.Vb

B=Pb.Pc.Vb.Vc

C=Pc.Pa.Vc.Va

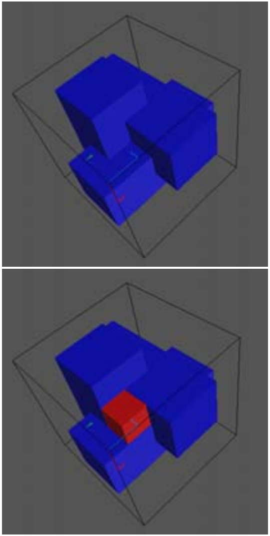

Higher dimensional complex with a forbidden region consisting of isothetic hypercubes and an unsafe region.

Geometric deadlock detection 1

Deadlocks may occur at the
lower corners of intersections of *n* hypercubes –
unless contained in the interior of the forbidden region.

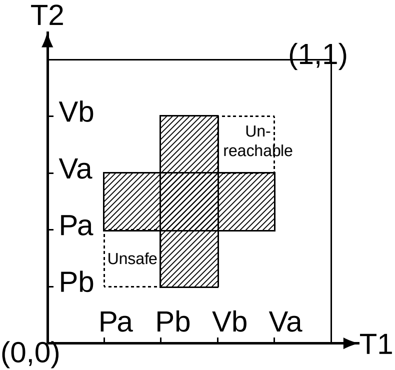

Dimension 3: The front right upper corner of a room is the intersection of 3 (forbidden) walls.

Geometric deadlock detection 2

published in CONCUR'98 proceedings

Theorem

A (non-final) point $\mathbf{x} \in X = I^n \backslash int(F)$ *is a deadlock if and only if*

- *there is an n-element index set* $J = \{i_1, \ldots, i_n\}$ *with*

$$R^J = R^{i_1} \cap \cdots \cap R^{i_n} \neq \emptyset$$

- $\mathbf{x} = \mathbf{a}^J = (a_1^J, \ldots, a_n^J) = \min R^J \notin int(F)$.

Remark The coordinate a_j^J is then maximum of the j-th coordinates of the lower corners of the participating hypercubes R^i – easy to find algorithmically.

(Mininal) unreachable points can be found analogously.

Unsafe regions?

An algorithm detecting unsafe regions

published in CONCUR'98 proceedings, illustrations courtesy Eric Goubault

From the PV-program

- Compute the forbidden region $F \subset I^n$,
- The intersections R^J of n forbidden hyperrectangles $R^i = [a_1^i, b_1^i] \times \cdots [a_n^i, b_n^i]$ create deadlocks.
- Forbid successively the hyperrectangles $[\tilde{x}, x]$, where $x = \min J = (\max_i a_1^i, \cdots \max_i a_n^i)$ and $\tilde{x} = (\text{2nd } \max_i a_1^i, \cdots, \text{2nd } \max_i a_n^i) \leadsto$ secondary deadlocks, unsafe regions.

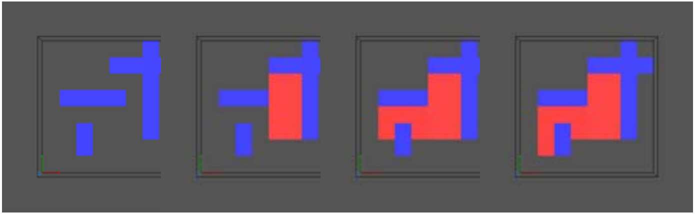

Dipaths and dihomotopy

combinatorial and geometric

Definition Two dipaths $f_0, f_1 : I \to X$ from **a** to **b** are dihomotopic if there is a one-parameter family $H : I \times I \to X$ such that $H_t = H(t, -)$ is a dipath for every t, $H_0 = f_0, H_1 = f_1$, $H(0, s) =$ **a**, and $H(1, s) =$ **b**.

Definition Combinatorial dipath: Concatenation of dipaths in $X \subset I^n$ parallel to one of the axes.

Elementary dihomotopy:

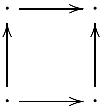

Combinatorial dihomotopy: Congruence relation generated by elementary dihomotopies.

Theorem

(L. Fajstrup, 05): In a cubical complex, combinatorial dipaths/combinatorial dihomotopy $\simeq$ *dipaths/dihomotopy.*

Dihomotopy is finer than homotopy with fixed endpoints

Example: Two wedges in the forbidden region

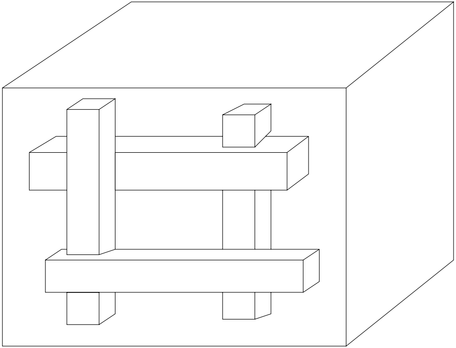

All dipaths from minimum to maximum are homotopic.
A dipath through the “hole” is not dihomotopic to a dipath on the boundary.

Why bother? Database management as an example
Serial and serializable executions

An execution is

serial if only one proces is accessing databases at a given time;
$P_{i_1}.P_{i_2}.\cdots.P_{i_n}$

serializable if the result of a schedule is always equivalent to a serial exedution (safe).

Correctness is

- often easy to check for serial executions
- difficult or impossible to check for general executions

Serializable executions have advantages:

- Check correctness for serial executions only!
- Can be much faster than a serial execution!

2-phase locking protocols

Which schedules (protocols) are known to be serializable?
Data engineers often use 2-phase locked protocols.

For those, every proces P_i should

- first do all the lock operations
- then the computations
- finally all the unlock operations

$$PPP \dots PVVV \dots V$$

2-phase locking is safe

Theorem

Every diapth in a 2-phase locking protocol is serializable (thus "safe" and correct).

Proofs using

graph theory came first, but were quite complicated

topological methods more transparent

- J. Gunawardena (1994)
- L. Fajstrup, E. Goubault, M.R. (1999, finally published in TCS, 2006)

Idea For a 2-phase locked protocol, the forbidden region F has a particular geometric structure ("blockwise starshaped"). This property can be used to prove geometrically, that every dipath in X is dihomotopic to a dipath on the edges of I^n – modelling a serial execution.

Conclusion Every execution is equivalent to a serial execution!

A single hypercube gives rise to nontrivial dihomotopy

but only between points in specified regions

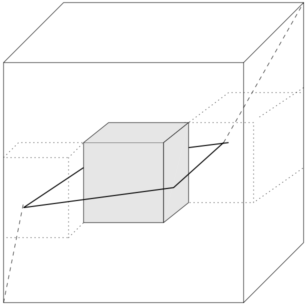

In dim. n: $n-2$ coordinates “forbidden”, 2 coordinates ”free”.

The two dipaths pass through forbidden intervals in reverse orders.

Nontrivial dihomotopy only “persists” if source and target live in the dotted boxes.

Conditions for persistent nontrivial dihomotopy?

Projections and forbidden regions

A tool linking to the deadlock situation

Projections to the first $(n-1)$, resp. last coordinate:
$\pi^n : \mathbf{R}^n \to \mathbf{R}^{n-1}, \pi_n : \mathbf{R}^n \to \mathbf{R}, A \mapsto A^n = \pi^n(A), A_n = \pi_n(A)$

$$\begin{aligned} A &= (a_1, b_1) \times \cdots \times (a_{n-1}, b_{n-1}) \times (a_n, b_n) \\ A^n &= (a_1, b_1) \times \cdots \times (a_{n-1}, b_{n-1}) \\ A_n &= (a_n, b_n) \end{aligned}$$

What happens to the forbidden region under projection?
New forbidden region $F^n \subset I^{n-1}$, new state space
$\bar{X} \subset I^{n-1}, \quad \bar{X} \neq X^n!$
$X = I^n \setminus F \subset I^n \quad \bar{X} = I^{n-1} \setminus F^n \subset I^{n-1}$
$\mathbf{x} \in I^{n-1}$ forbidden $\Leftrightarrow \mathbf{x} \in F^n \Leftrightarrow \exists x_n \in I$ with $(\mathbf{x}, x_n) \in F$.
$\bar{X}$ can have deadlocks even if X is deadlockfree.

Forbidden region and projection
from $n-1$ intersecting hyperrectangles

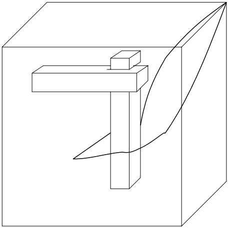

$J=(i_1,\dots,i_{n-1}), R_J=\bigcap_1^{n-1} R^i=(a_1,b_1)\times\cdots\times\times(a_n,b_n)$ the intersection of $n-1$ forbidden hyperrectangles.

$$\begin{aligned} R_J &= (a_1,b_1)\times\cdots\times(a_{n-1},b_{n-1})\times \quad (a_n,b_n)\\ R_J^n &= (a_1,b_1)\times\cdots\times(a_{n-1},b_{n-1})\\ R_{J,n} &= \qquad\qquad\qquad\qquad\qquad\qquad\quad (a_n,b_n)\end{aligned}$$

$R_J^n\subset I\ \ n-1$ gives rise to the deadlock $(a_1,\dots,a_{n-1})\in\bar{X}$ and $(b_1,\dots,b_{n-1})$ is unreachable.

Nontrivial dihomotopy
from $n-1$ intersecting hyperrectangles

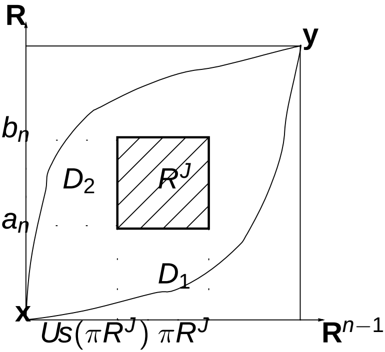

For a dipath $\alpha = (\alpha^n, \alpha_n)$ from $\mathbf{x} \in Us(R^n_J)$ either

- α^n waits in $Us(R^J_n)$ until $\alpha_n(t) > b_n$ (through D_2) or
- α^n passes R^J before $\alpha_n(t) > a_n$ (through D_1)

A dihomotopy respects this choice: D_1, D_2 disconnected!

Nontrivial dihomotopy from source to target

The double wedge example

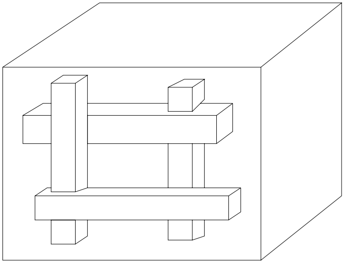

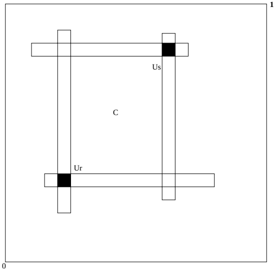

The dipath through the hole is not dihomotopic to a dipath on the boundary:
The projection to I^{n-1} exhibits intersection of an unsafe and unreachable region that is disconnected from source and target.

Nontrivial dihomotopy from source **0** to target **1**

for two arrangements of $n-1$ pairwise intersecting hyperrectangles:

$I=(i_1,\dots,i_{n-1}) \qquad R^I=(a_1,b_1)\times\cdots\times(a_n,b_n)$

$J=(j_1,\dots,j_{n-1}) \qquad R^J=(c_1,d_1)\times\cdots\times(c_n,d_n); a_n<d_n$

$(c_1,\dots,c_{n-1})$ deadlock in $\bar{X}$, $\quad (b_1,\dots,b_{n-1})$ unreachable in $\bar{X}$.

Theorem

Let $C=Us(R^n_I)\cap Us(R^n_J)$ be disconnected from $\mathbf{0}$ and $\mathbf{1}$.

If $\alpha\in\vec{P}_1(X)(\mathbf{0},\mathbf{1})$ has property

$$\text{(P)}\quad a_n\le\alpha_n(t)\le d_n\Rightarrow\alpha^n(t)\in C$$

then so has every $\beta\in\vec{P}_1(X)(\mathbf{0},\mathbf{1})$ dihomotopic to α.

Proof uses directed van Kampen theorem (M. Grandis, '03)

Corollary 1 A dipath $\alpha\in\vec{P}_1(X)(\mathbf{0},\mathbf{1})$ satisfying (P) is not serializable (dihomotopic to a dipath on the 1-skeleton).

Corollary 2 $\vec{\pi}_1(X)(\mathbf{0},\mathbf{1})$ is not trivial.

Trivial dihomotopy for other intersection patterns 1

Theorem

Assume that the forbidden region $F = \bigcup R^i$ satisfies:
$R^J = \bigcap_{i\in J} R^i = \emptyset$ for all index sets J with $|J| \geq n-1$.
*Then any two dipaths in the model space $X = I^n \setminus F \subset I^n$ from **0** to **1** are dihomotopic: $\vec{\pi}_1(X)(\mathbf{0},\mathbf{1})$ is trivial.*

Tool: σ_i: one edge step along x_i-axis. Every combinatorial dipath can be written in the form $\sigma_{i_L} * \cdots * \sigma_{i_2} * \sigma_{i_1}$.
For a vertex $x \in X$, let $Out(x) = \{\sigma_{i_1}, \cdots, \sigma_{i_k}\} \subseteq \{\sigma_1, \cdots, \sigma_n\}$ denote the set of edges with source x.
Proposition. Assume that X has no deadlocks and that for every vertex $x \in X$ and all directed edges $\sigma_{i_1}, \sigma_{i_2} \in Out(x)$:

1. $\sigma_{i_1}, \sigma_{i_2}$ homotopy commute[1] or
2. $\exists j \neq i_1, i_2 : \sigma_j$ homotopy commutes with both σ_{i_1} and σ_{i_2}.

Then $\vec{\pi}_1(X)(\mathbf{0},\mathbf{1})$ is trivial.

[1] Exists a 2-cube filling $\sigma_{i_1} * \sigma_{i_2}, \sigma_{i_2} * \sigma_{i_1}$

Trivial dihomotopy for other intersection patterns 2

Proof by induction on the "length" of dipaths.
Why is the condition on homotopy commutativity satisfied for forbidden regions in which at most $n-2$ hyperrectangles intersect nontrivially?
Look at the "local future" of a vertex $\mathbf{x}$. It is always of the form $\partial_-\vec{I}^{j_1} \times \cdots \times \partial_-\vec{I}^{j_k} \times \vec{I}^{n-j}$, $j := j_1 + \cdots + j_k$.
In our case $k \leq n-2$. Hence either

- $j < n$ (factor $\vec{I}$) or
- at least one $j_i \geq 3$ (lower boundary of a 3 cube) or
- there exist $i_1 \neq i_2$ with $j_{i_1} = j_{i_2} = 2$ (product of two Ls containing enough rectangles).