Decomposition of Processes and Factorization of Cubical Area

Nicolas Ninin
CEA - Université Paris Sud
nicolas.ninin@lix.fr

Introduction

We model concurrent programs without loop nor branching by certain subsets of $\mathbb{R}^n$. One major hurdle met in static analysis of such programs is that the size of their state spaces exponentially depend on the number of processes they are made of. The size of the model can be made much smaller if one can write the program as a parallel composition of independent groups of processes. We provide an efficient and intuitive algorithm that returns such a factorization.

PV Programs

A *resource* of arity n can be simultaneously used by at most $n-1$ processes. A *PV* program is a parallel composition of sequential processes that take and release any resource a through the instructions $P(a)$ and $V(a)$ *e.g.* $P(a)P(b)V(b)V(a) \mid P(b)P(a)V(a)V(b)$.

Geometric model

The geometric model of $P = P_1|...|P_n$ is a subset of $\mathbb{R}^n$ obtained as follows:

- Associate each process P_i with an axis of $\mathbb{R}^n$
- Choose a finite set of points of the axis of P_i and label them with the instructions of P_i
- Any state where a resource of arity α is held by (at least) α processes is forbidden
- The forbidden region is a finite union of cubes $C_1 \cup ... \cup C_k$.

$P = Pa.Pb.Vb.Va|Pb.Pa.Vb.Va$

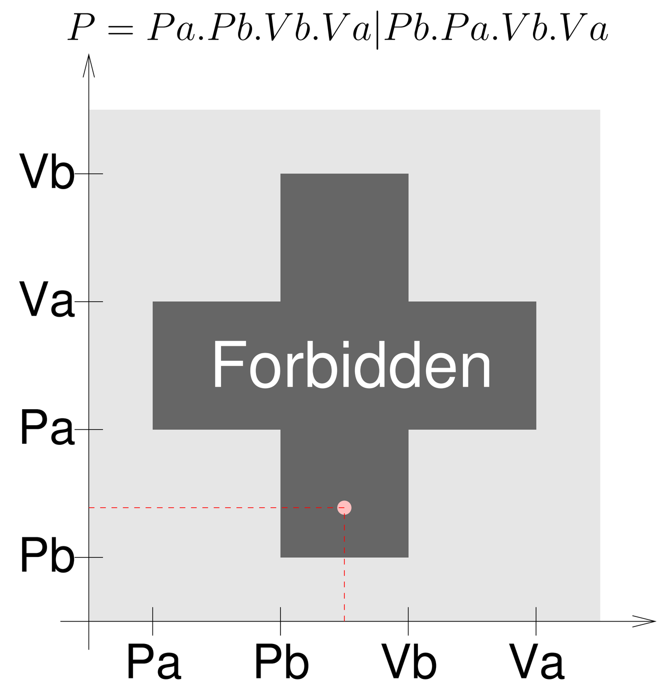

Two forbidden rectangles
$C_1 = [2,3] \times [1,4]$ and $C_2 = [2,3] \times [1,4]$.
The state space is $X = \mathbb{R}^2 \setminus (C_1 \cup C_2)$

$P = Pa.Va|Pa.Va|Pb.Vb$

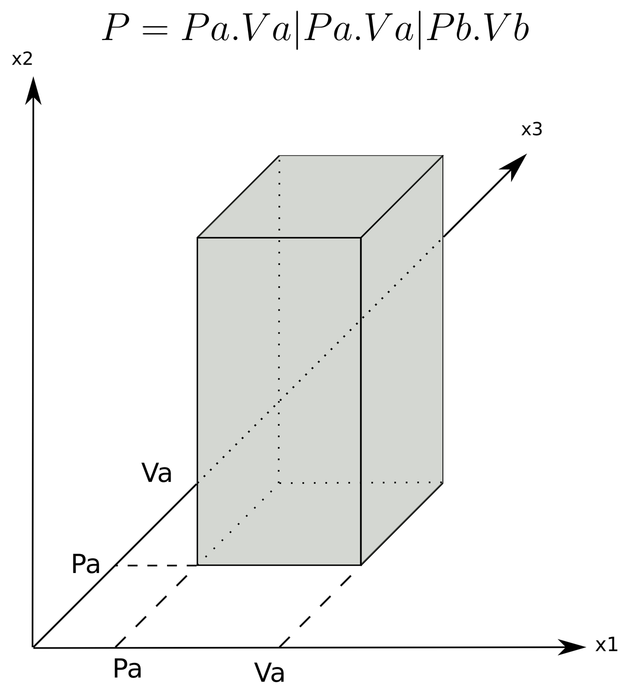

One forbidden cube and
$X = \mathbb{R}^3 \setminus ([1,2] \times [1,2] \times \mathbb{R})$

Factorization

A space $X \subset \mathbb{R}^n$ can be factorized if one can write X as $X_1 \times X_2$ (up to permutation of the coordinates), with X_1 and X_2 subspaces of $\mathbb{R}^p$ and $\mathbb{R}^{n-p}$.

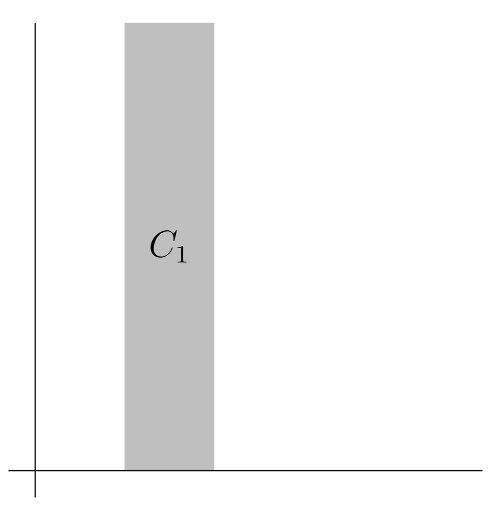

$X = (\mathbb{R} \setminus ([2,3]) \times \mathbb{R}$

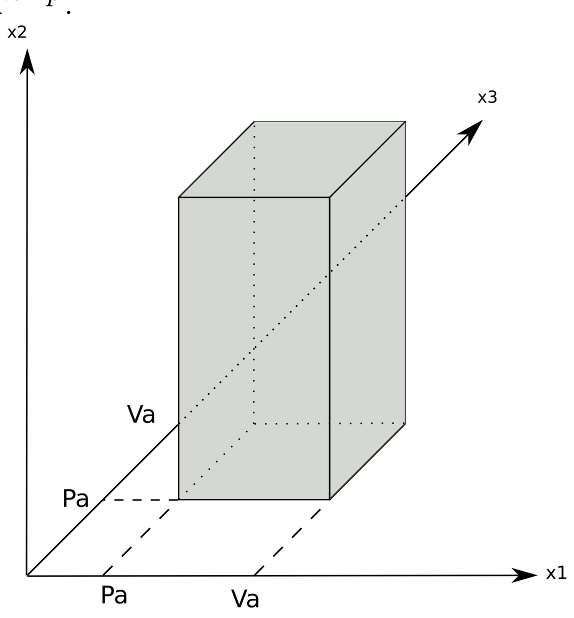

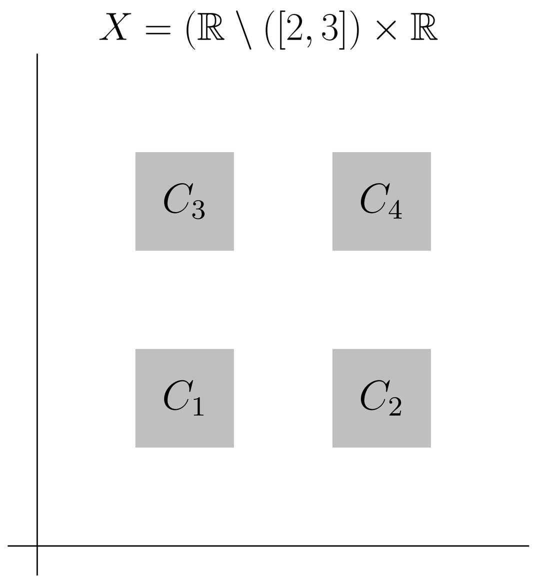

This space cannot be factorized

$X = (\mathbb{R}^2 \setminus ([1,2]^2) \times \mathbb{R}$

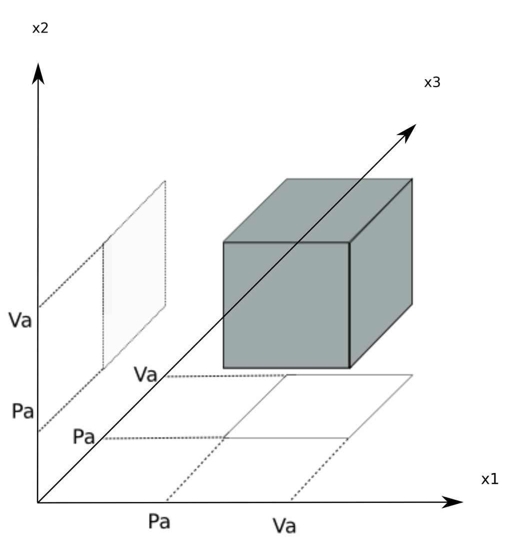

This space cannot be factorized

Syntactic Factorization

Observe the following code (the forbidden region is a pillar):

$$P = Pa.Va|Pa.Va|Pb.Vb$$

An easy homemade analysis proves that the first two processes are independent from the third one which will be denoted by $\{\{1,2\},\{3\}\}$. The syntactical factoring algorithm is as follows

- Gather all the processes that share a given resource in a single block
- Gather two blocks when they intersect (repeat inductively until a partition is obtained)

$PaVa|PaVa|PaVa$

Its the floating cube, the resource a is shared by all three processes P_1, P_2 and P_3, we'll say that the syntactic factorization is $\{1,2,3\}$ which means that the space doesnt really factorize

$Pa.Pb.Va.Vb|Pb.Pc.Vb.Vc|$
$Pc.Pd.Vc.Vd|Pd.Pa.Va.Vd$

It's the four philosophers example. Here we have four resources a,b,c,d all shared among two processes. From a we group $\{1,4\}$, from b we group $\{2,3\}$, and c group 3 and 4, since 4 is already with 2 we end up with $\{1,2,3,4\}$, again it doesn't factorize.

$Pa.Pb.Vb.Va|Pa.Pb.Vb.Va|$
$Pb.Vb$

Here b links the three processes and so we again have no factorization

Improving the Syntactic Factoring Algorithm

Let's come back to our last example and look at the space corresponding in $\mathbb{R}^3$.

$$P = Pa.Pb.Vb.Va|Pa.Pb.Vb.Va|Pb.Vb$$

The state space X is $\mathbb{R}^3$ from which a pillar and a cube have been removed. However the forbidden cube given by b is included in the forbidden pillar of a, therefore:

$$X = (\mathbb{R}^2 \setminus [1,2) \times [1,2]) \times \mathbb{R}$$

Which means that X is actually factorizable.

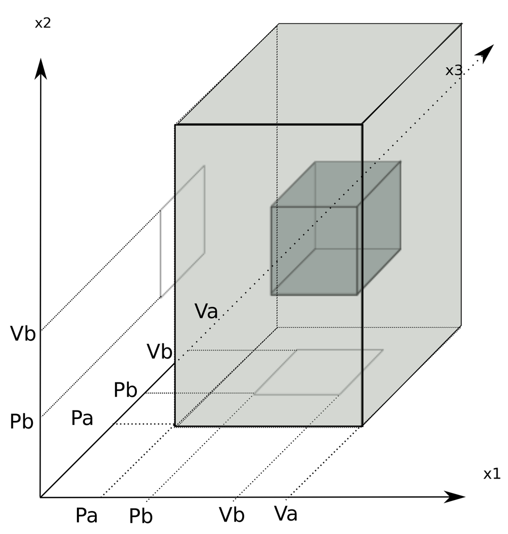

The resource b was indeed semantically useless. This phenomenon is observed each time a forbidden cube is contained in the union of the others, so the syntactic algorithm is *not optimal*.

An optimal algorithm

Suppose we are given the *maximal* cubes of the forbidden region (instead of the source code):

- Associate each maximal cube C with the block $\{i \in \{1,\ldots,n\} \mid pr_i(C) \neq \mathbb{R}\}$
- Gather two blocks when they intersect (repeat inductively until a partition is obtained)

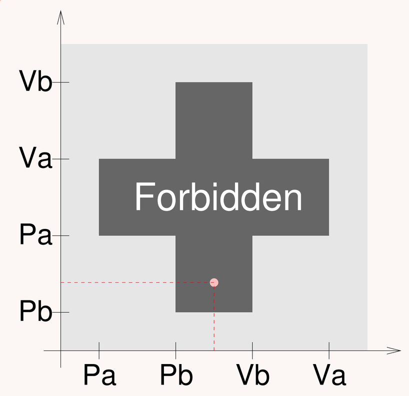

$[2,3] \times [1,4] \mapsto \{1,2\}$
$[1,4] \times [2,3] \mapsto \{1,2\}$
Factorization: $\{\{1,2\}\}$

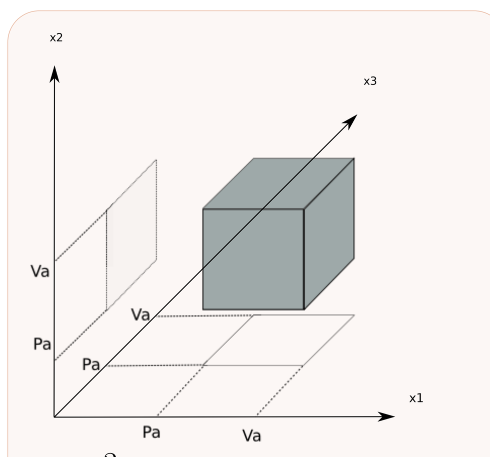

$[1,2]^3 \mapsto \{1,2,3\}$
Factorization: $\{\{1,2,3\}\}$

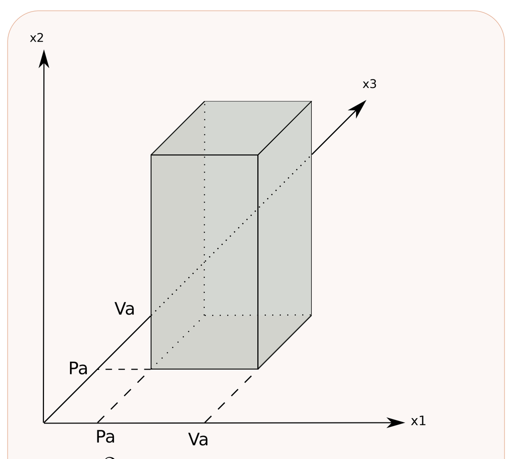

$[1,2]^2 \times \mathbb{R} \mapsto \{1,2\}$
Factorization: $\{\{1,2\},\{3\}\}$

Problems Under The Rug ?

Some difficulty have been avoided, we list a few of them:

- The decomposition of regions is only defined up to permutation of the coordinates.
- The syntactic algorithm cannot be easily modified to deal with the semaphore overlap problem.
- The algorithm that provides the maximal cubes is exponential.