

Distributed Data Base

Abstract

This is a small toy example which is well-suited as an introduction to occurrence graphs. The analysis of the occurrence graph is described in great detail.

The CPN model describes the communication between a set of data base managers in a distributed system. The model is identical to the “Distributed Data Base” presented in “Introductory Examples”(which we recommend to study before this example).

The example is taken from Sect. 1.5 of Vol. 2 of the CPN book.

Developed and Maintained by:

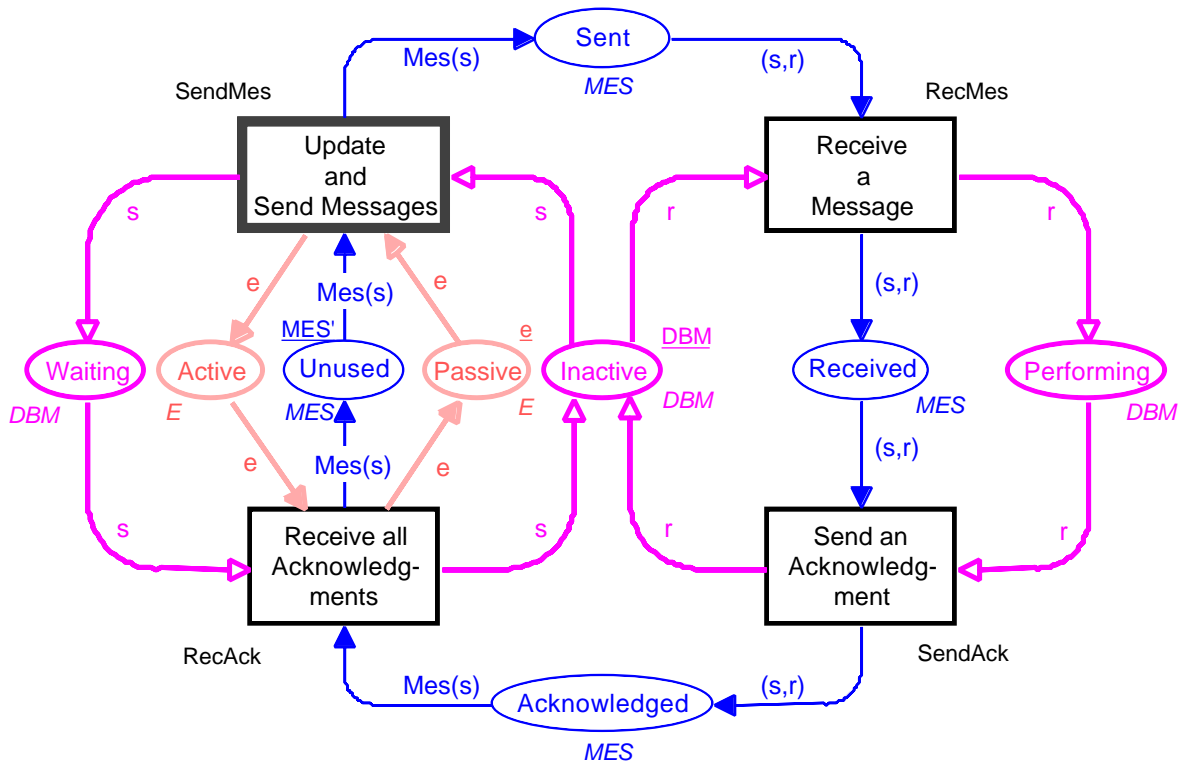
Kurt Jensen, Aarhus University, Denmark (kjensen@daimi.aau.dk).

Graphical Quality

The figures in this document are inserted via PICT format. This is why some of the arcs and place borders look a bit ragged. A postscript printout from Design/CPN (and the screen image in Design/CPN) has much higher graphical quality.

CPN Model

In this example we study the O-graph for the data base system, i.e., the O-graph for the following CP-net:

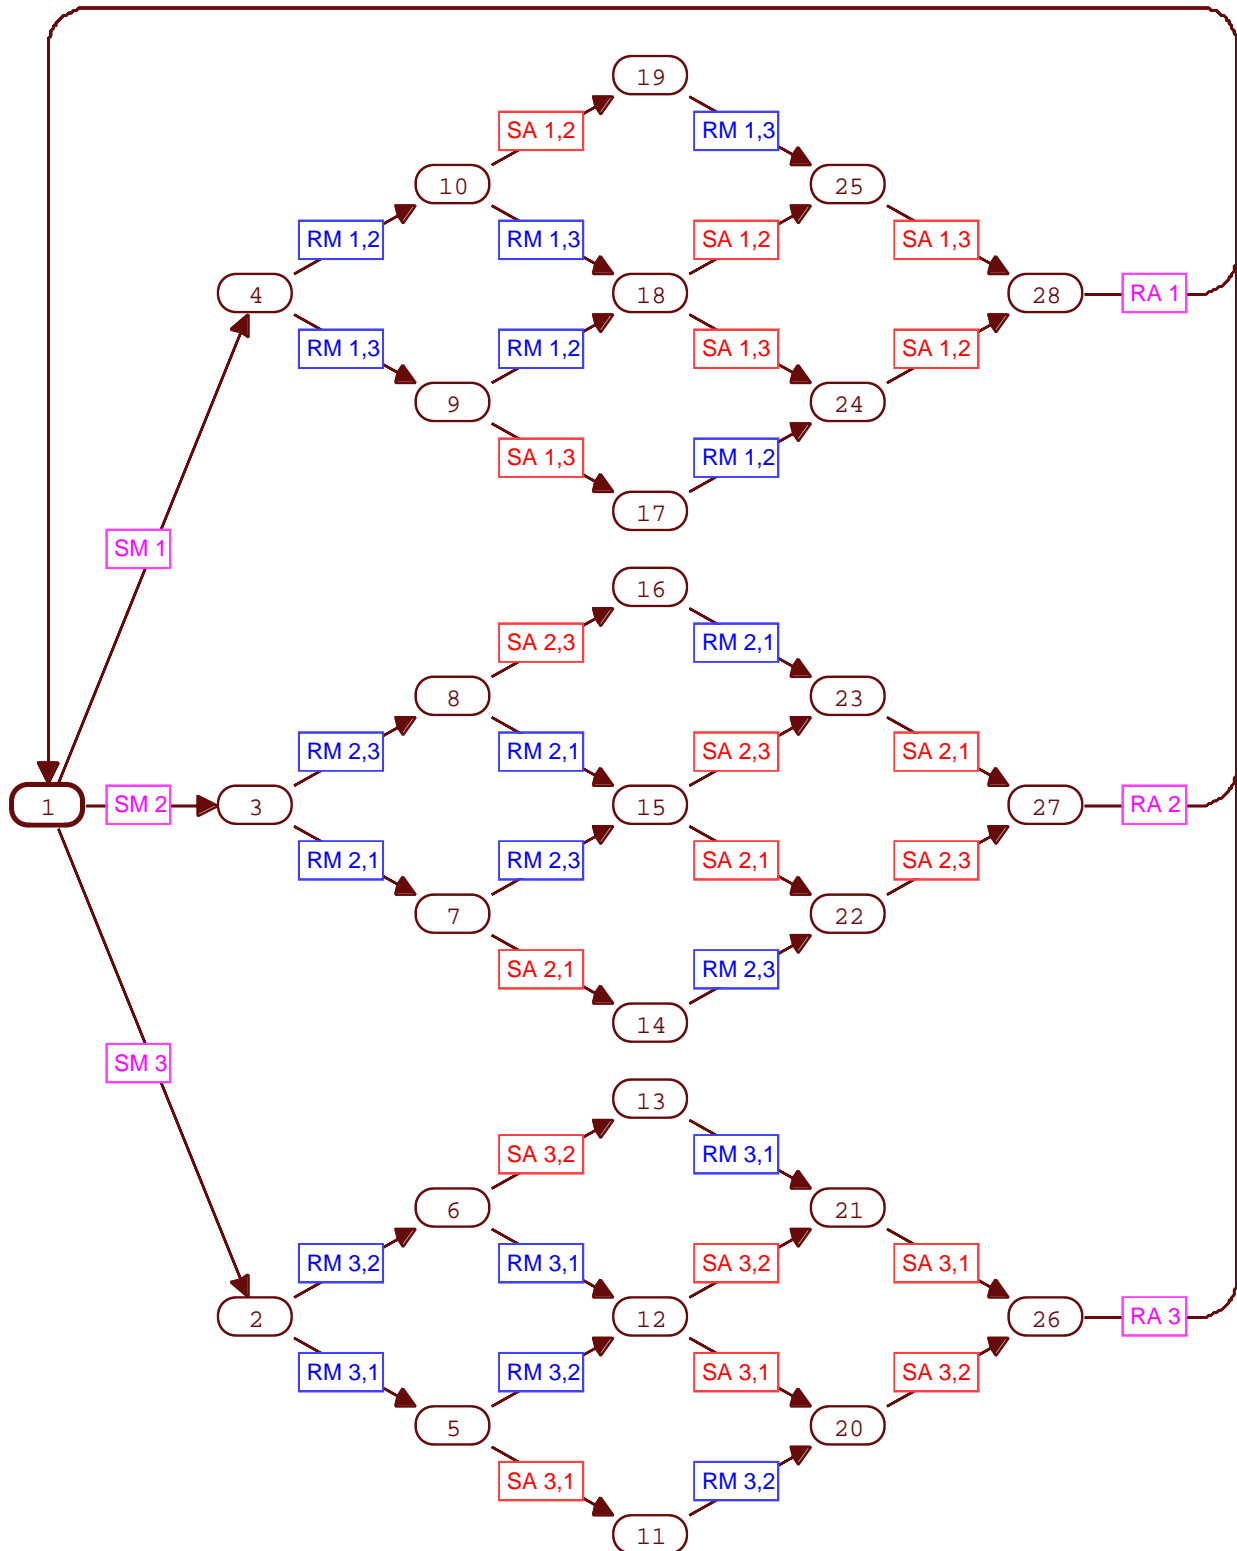


```

val n = 3;
color DBM = index d with 1..n declare ms;
color PR = product DBM * DBM declare mult;
fun diff(x,y) = (x<>y);
color MES = subset PR by diff declare ms;
color E = with e;
fun Mes(s) = mult'PR(1`s,DBM-1`s);
var s, r : DBM;

```

For three data base managers the O-graph looks as shown below. Node number one is the initial marking. To save space the transition names are abbreviated to SM, RM, SA, and RA. Moreover, we write SM i and RM i,k instead of $(SM, \langle s = d_i \rangle)$ and $(RM, \langle s = d_i, r = d_k \rangle)$, and analogously for SA and RA. These abbreviations are obtained, quite easily, by using the replace operation in the Find command (which must be used in the CPN editor).



The standard report looks as shown below.

From the statistics we see that there only is one strongly connected component. This means that all reachable states are reachable from each other.

Statistics	

Occurrence Graph	
Nodes:	28
Arcs:	42
Secs:	2
Status:	Full
Scc Graph	
Nodes:	1
Arcs:	0
Secs:	0

All the integer bounds are as expected (see below). In particular, we see that at most one process can be *Waiting*. This tells us that a new update cannot start until all data base managers have finished the processing of the previous one. Also the multi-set bounds are as expected. To improve the readability, we have substituted *DBM* for the multi-set $1^{d(1)} + 1^{d(2)} + 1^{d(3)}$ and *MES* for $1^{(d(1),d(2))} + 1^{(d(1),d(3))} + 1^{(d(2),d(1))} + 1^{(d(2),d(3))} + 1^{(d(3),d(1))} + 1^{(d(3),d(2))}$.

Boundedness Properties		

Best Integers Bounds		
	Upper	Lower
Acknowledged	2	0
Active	1	0
Inactive	3	0
Passive	1	0
Performing	2	0
Received	2	0
Sent	2	0
Unused	6	4
Waiting	1	0
Best Upper Multi-set Bounds		
Acknowledged	MES	
Active	1^e	
Inactive	DBM	
Passive	1^e	
Performing	DBM	
Received	MES	

Sent	MES
Unused	MES
Waiting	DBM
Best Lower Multi-set Bounds	
Acknowledged	empty
Active	empty
Inactive	empty
Passive	empty
Performing	empty
Received	empty
Sent	empty
Unused	empty
Waiting	empty

The home properties tell us that all reachable markings are home markings. From the drawing of the occurrence graph, we can actually deduce that the system has a much stronger property. It is not only *possible* to return to the initial marking. This will *always* happen – whenever $2 * n$ transitions have occurred.

Home Properties

Home Markings: All

Also the liveness properties are as expected. There are no dead markings and all transitions are live.

Liveness Properties

Dead Markings: None
Dead Transitions Instances: None
Live Transitions Instances: All

Finally, the fairness properties tells us that all transitions are impartial. This is also easy to see from the drawing of the occurrence graph. Whenever $2 * n$ transitions have occurred *SendMes* and *RecAck* have occurred exactly one time each, while *RecMes* and *SendAck* have occurred exactly $n - 1$ times each.

Fairness Properties	

SendMes	Impartial
RecMes	Impartial
SendAck	Impartial
RecAck	Impartial

Now let us look at some model dependent properties. First we investigate whether the transitions are strictly live. For *SendMes* and *RecMes* the queries look as shown below. They show us that *SendMes* is strictly live, while *RecAck* is not – because binding elements such as $(\text{RecAck}, \langle s=d(2), r=d(2) \rangle)$ are dead. If we add a guard, $[s \langle \rangle r]$, to *RecAck*, the transition becomes strictly live.

Strict Liveness

<pre> BEsStrictlyLive [Bind.Top'SendMes (1, {s=d(1)}), Bind.Top'SendMes (1, {s=d(2)}), Bind.Top'SendMes (1, {s=d(3)})]; BEsStrictlyLive [Bind.Top'RecMes (1, {s=d(1), r=d(1)}), Bind.Top'RecMes (1, {s=d(1), r=d(2)}), Bind.Top'RecMes (1, {s=d(1), r=d(3)}), Bind.Top'RecMes (1, {s=d(2), r=d(1)}), Bind.Top'RecMes (1, {s=d(2), r=d(2)}), Bind.Top'RecMes (1, {s=d(2), r=d(3)}), Bind.Top'RecMes (1, {s=d(3), r=d(1)}), Bind.Top'RecMes (1, {s=d(3), r=d(2)}), Bind.Top'RecMes (1, {s=d(3), r=d(3)})]; </pre>	<pre> > true : bool > false : bool </pre>
--	---

Next we investigate the fairness properties of some typical binding elements. We see that the binding element of *SendMes* is just, while those of the other three transitions are fair.

Fairness of Binding Elements

<pre> BEsFairness [Bind.Top'SendMes (1, {s=d(1)})]; BEsFairness [Bind.Top'RecMes (1, {s=d(1), r=d(3)})]; BEsFairness [Bind.Top'SendAck (1, {s=d(1), r=d(3)})]; BEsFairness [Bind.Top'RecAck (1, {s=d(1)})]; </pre>	<pre> > Just : FairnessProperty > Fair : FairnessProperty > Fair : FairnessProperty > Fair : FairnessProperty </pre>
--	--

Finally, let us demonstrate that occurrence graphs also can be used to check whether place invariants are fulfilled. It should, however, be stressed that the best way to check place invariants (for complex systems) is by checking the place flow property, which is a static and local property that can be checked *without* generating all possible system states. We want to check the following two place invariants:

$$\begin{aligned}
 M(\text{Performing}) &= \text{Rec}(M(\text{Received})) \\
 \text{Mes}(\text{Waiting}) &= M(\text{Sent}) + M(\text{Received}) + M(\text{Acknowledged}).
 \end{aligned}$$

We first define a function *Rec* that maps a message into its receiver.

A Projection Function

```
fun Rec((s,r):MES)=r; > val Rec = fn : MES -> DBM
```

Then we extend *Rec* and *Mes* to two new functions *Rec'* and *Mes'* which can be applied to *multi-sets* of data base managers. This is done by means of two predeclared functions *ext_col* and *ext_ms*. The first of these extends a function $[A \ B]$ to a function $[A_{MS} \ B_{MS}]$, while the second extends a function $[A \ B_{MS}]$ to a function $[A_{MS} \ B_{MS}]$.

Extensions to Multi-sets

```
val Rec' = ext_col Rec; > val Rec' = fn : (MES ms) -> (DBM ms)
val Mes' = ext_ms Mes; > val Mes' = fn : (DBM ms) -> (MES ms)
```

By definition the result of using an extended function to a multi-set is obtained by using the original function to each element in the multi-set – adding the results. This is illustrated by the following examples. The two *mkst_ms* functions map the ML representation of a DBM/MES multi-set into a much more readable string representation.

Examples

```
mkst_ms'DBM (Rec'(1^(d(1),d(3))+ > "1`d(2)+ 1`d(3)" :
                1^(d(1),d(2))))); string
mkst_ms'MES (Mes'(1`d(1)+1`d(2))); > "1^(d(1),d(2))
+ 1^(d(1),d(3))
+ 1^(d(2),d(1))
+ 1^(d(2),d(3))
" : string
```

Finally, we use a predeclared search function called *PredAllNodes* to list all nodes violating the two invariants. There are no such nodes, and hence we have proved that the invariants are fulfilled in all reachable markings. Please note that the $\langle \rangle$ operator checks whether two multi-sets differ from each other (if you replace $\langle \rangle$ by $\langle \rangle$ you only check whether the *representations* of the two multi-sets differ from each other).

Check of Two Place Invariants

```
PredAllNodes(fn n =>
(Mark.Top'Performi 1 n) <><>
Rec'(Mark.Top'Received 1 n));
```

```
PredAllNodes(fn n =>
Mes'(Mark.Top'Waiting 1 n) <><>
(Mark.Top'Sent 1 n)+
(Mark.Top'Received 1 n)+
(Mark.Top'Acknowle 1 n));
```

```
> [] : Node list
> [] : Node list
```

For the data base system it is rather easy to calculate how fast the O-graph grows – when we increase the number of data base managers. The results are as shown below. They illustrate the **space complexity** of the O-graph algorithm:

DBM	Nodes	Arcs
$O(n)$	$O(n * 3^n)$	$O(n^2 * 3^n)$
2	7	8
3	28	42
4	109	224
5	406	1,090
6	1,459	4,872
7	5,104	20,426
8	17,497	81,664
9	59,050	314,946
10	196,831	1,181,000
15	71,744,536	669,615,690
20	23,245,229,341	294,439,571,680

As illustrated above, it is often the case that the O-graph of a CP-net grows very fast when the sizes of the involved colour sets increase. However, in practice, it is fortunately often sufficient to consider rather small colour sets in order to verify the logical correctness of a given CP-net. Having convinced ourselves that the data base system has the correct behaviour for 4 or 5 managers, we can feel pretty sure that the design also works correctly for any larger number of managers. Sadly, a similar statement is not true when we try to evaluate the performance of a given system.