Can Skeletons Really Be Used to Detect Deadlocks of Nets?

Greg Findlow
Telecom Australia Research Laboratories
770 Blackburn Road
Clayton North, Victoria 3168
Australia

Abstract

Vautherin's results relating the behaviour of classes of coloured (high-level) Petri nets to ordinary Petri nets (skeletons for the classes) are examined, and their usefulness and ease of application are investigated. The feasibility of using skeletons for the purpose of 'first estimate' verifications of nets (in particular, for the early detection of deadlocks, i.e. reachable dead markings) is considered. It is found that because the relationship established by Vautherin between the dead markings of a coloured net and those of the corresponding skeleton does not differentiate between reachable and unreachable markings, deadlock-freeness of a coloured net and deadlock-freeness of its skeletons are essentially independent properties.

1 Introduction

Vautherin [7] describes a way of combining coloured Petri nets with algebraic specification techniques, for the purpose of specifying parallel systems. His approach is to use a 'schema' to represent a whole class of coloured Petri nets, in a way analogous to that in which one abstract data type signature represents a number of (many-sorted) algebras [2]. Indeed, each of Vautherin's schemas is written with respect to such a signature, and each corresponding algebra gives rise to an 'interpretation' of the schema as a coloured net.

The main results of Vautherin's paper relate the behaviour of all the coloured Petri nets which can be viewed as interpretations of a given schema to the behaviour of an ordinary Petri net obtained from the schema, called the skeleton. In particular, the possibility of detecting deadlocks of coloured nets simply by looking at their skeletons is suggested. The investigation of this possibility is the purpose of this paper.

It is very important to note at this point that whereas we will use the term 'deadlock' to mean a *reachable* dead marking of a net, Vautherin uses it to mean any dead marking, reachable or otherwise.

The concept of skeleton is old—it is, for example, introduced by Jensen for coloured nets in [4] and also [5]—he calls the skeleton the 'underlying place/transition net'. For simplicity, we will also work with skeletons obtained directly from coloured nets, rather than using schemas. Vautherin's main results can be easily transferred into this new context [3].

Section 2 gives the definition of coloured Petri nets we will be using, and defines skeletons of (some) coloured nets, as well as describing briefly how these relate to Vautherin's skeletons of schemas. In Section 3 we transfer the results of interest in Vautherin's paper to the case of skeletons derived directly from coloured nets. Section 4 points out the subtleties which need to be recognised when attempting to apply the results, while Section 5 illustrates how deceptive the results produced by a skeleton can be if the considerations of Section 4 are ignored. A possible way of making use of a (possibly deceptive) set of results generated by a skeleton is suggested in Section 6. In Section 7 we identify the additional work which needs to be performed in order that skeletons may be usefully applied to detect deadlocks of coloured nets.

2 Coloured nets and skeletons

We will use the following definition of coloured nets.

Definition 1 *A* **coloured Petri net** *is a tuple* $N = (P, T, C, W, M_0)$*, where:*

- P *is a finite set of places.*
- T *is a finite set of transitions, disjoint from* P*.*
- C *is a* $P \cup T$*-indexed family of nonempty sets (colour sets for the places and occurrence colour sets for the transitions).*
- W *is a* $(P \times T \cup T \times P)$*-indexed family of functions such that for each* $t \in T$*, each* $\sigma \in C(t)$*, and each* $p \in P$*,* $W(p,t)(\sigma)$ *and* $W(t,p)(\sigma)$ *are multisets over* $C(p)$*.*
- M_0 *is the initial marking—a function assigning to each place* $p \in P$ *a multiset over* $C(p)$*.*

When we draw a coloured Petri net, we shall draw an arc from each place p to each transition t for which $W(p,t)(\sigma)$ is nonempty for some $\sigma \in C(t)$. Similarly, if there is some $\sigma \in C(t)$ for which $W(t,p)(\sigma)$ is nonempty, an arc from t to p will be included in the drawing. These arcs will carry 'arc expressions', which may contain 'variables'. The colour set of each transition will not be explicitly specified, rather it will be understood to consist of all possible assignments to the variables on the surrounding arcs of the transition. It is possible to give a formal definition of the graphical form of a coloured net (see e.g. [1, 5]). We will

TH0386-3/91/0000/0198/$01.00 © 1991 IEEE

not be concerned with such formalisation, however—the pictures we will use later are not intended to be formal objects themselves, merely a convenient way of representing them.

The skeleton of a coloured net is the ordinary Petri net obtained in the following way. One replaces each bag of tokens (in the initial marking and the 'arc expressions' of the net) by its 'size', effectively removing the colours from all tokens. For any place p and transition t of a coloured net, however, the sizes of the bags $W(p,t)(\sigma)$ and $W(t,p)(\sigma)$ generally depend on the choice of transition occurrence colour $\sigma \in C(t)$. Thus Jensen makes (the equivalent of) the following definition in [5].

Definition 2 *A coloured net $N = (P,T,C,W,M_0)$ is* **uniform** *if for each $p \in P$ and each $t \in T$, the sizes of the multisets $W(p,t)(\sigma)$ and $W(t,p)(\sigma)$ are independent of the choice of $\sigma \in C(t)$.*

In other words, a uniform coloured net is one having the property that for any place p and transition t, the number of tokens removed from p by t is the same for any occurrence colour of t, and similarly for tokens added to p. The following definition formalises the concept of skeleton only for nets with this property [5]. A non-uniform net does not have a well-defined skeleton (because of the first two parts of the following definition), and thus must have its transitions unfolded enough to convert it into a uniform net before a skeleton can be obtained.

Definition 3 *Let $N = (P,T,C,W,M_0)$ be a uniform coloured Petri net. The* **skeleton** *of N is the ordinary Petri net given by $|N| = (P,T,|W|,|M_0|)$, where for each $p \in P$ and each $t \in T$:*

- $|W|(p,t) = |W(p,t)(\sigma)|$, *for any* $\sigma \in C(t)$.
- $|W|(t,p) = |W(t,p)(\sigma)|$, *for any* $\sigma \in C(t)$.
- $|M_0|(p) = |M_0(p)|$.

To conclude this section, we will now describe briefly the relationship between skeletons obtained from uniform coloured nets as per the above definition, and Vautherin's skeletons of schemas.

Vautherin defines his schemas to have 'arc expressions' which are multisets of terms. Because of this, every coloured net which may be obtained as an interpretation of a schema *is* uniform (Vautherin uses the term 'simple'), and every schema has a skeleton. Furthermore, it is easy to show that if a coloured net N is realizable as an interpretation of a schema, then any two such schemas *do* have the same skeleton, and that this skeleton is in fact the same as the skeleton of the coloured net as described in the above definition [3].

Thus the skeletons we are dealing with here are essentially the same as Vautherin's—if one were to take any schema, interpret it as a coloured net using an appropriate algebra, and then take the skeleton of that net, the result would simply be the skeleton of the original schema. This means that skeletons of coloured nets may inherit (adaptations of) Vautherin's results without any trouble.

3 Vautherin's results

We will use the notation $M[t\rangle M'$ to indicate that M' is a marking which may be reached from the marking M by firing (once) the transition t (for some occurrence colour $\sigma \in C(t)$, if the net in question is coloured). The notation $M[t\rangle$ will indicate simply that t is enabled in the marking M (again, for some occurrence colour, in the case of a coloured net).

Let $N = (P,T,C,W,M_0)$ be a uniform coloured Petri net, so that $|N|$ is defined. One may define a map h from the set of all possible markings (i.e. reachable or otherwise) of N onto the set of all possible markings of $|N|$ as follows. Set $h(M) = |M|$, where $|M|$ is the marking of $|N|$ obtained from M by stripping all tokens of their colours, i.e. $|M|(p) = |M(p)|$ for all $p \in P$.

The following proposition corresponds to part of Vautherin's Proposition 3.

Proposition 1 *Let N, h be as described above. For any transition t of N (and $|N|$), and any markings M,M' of N, $M[t\rangle M'$ implies $h(M)[t\rangle h(M')$.*

Intuitively, Proposition 1 says simply that if a transition t has tokens of all the right colours available so that it may fire in the coloured net, then certainly it has the right *numbers* of tokens available, hence may also fire in the skeleton, and furthermore, the subsequent skeletal marking has the same number of tokens in each place as the subsequent coloured marking. All this should be fairly clear from the way the skeleton is obtained.

One may also see immediately from the definition of the skeleton that h sends the initial marking of N to that of $|N|$. Having made this observation, the following (which corresponds mostly to Vautherin's Proposition 4) can then be obtained as a fairly straightforward consequence of Proposition 1.

Proposition 2 *Let N,h be as above. For any marking M, transition t, and place p of N :*

1. *If $h(M)$ is dead in $|N|$, then M is dead in N.*
2. *If p is bounded by k (some integer) in $|N|$, then p is bounded by k in N.*
3. *If t is non quasi-live (may not fire in any reachable marking) in $|N|$, then it is non quasi-live in N.*
4. *If $|N|$ has the finite termination property (i.e. has no infinite transition sequence in its reachability graph), then N has the finite termination property.*

Vautherin needs the observation that h sends the initial marking of N to that of $|N|$ for (his versions of) parts 2, 3, and 4 of the above proposition, because these parts relate to properties defined in terms of reachability graphs. It is not needed for the first part, however, since the 'deadness' of a marking may be considered independently of its reachability.

Thus, in applying the results of this section, we need to be aware that Proposition 1 really gives a projection of the 'whole state space graph'[1] of N onto that of its skeleton, while the last three parts of Proposition 2 really concern only the projection of the reachability graph of N *into* that of $|N|$.

This subtlety is not pointed out by Vautherin—the only thing in his paper which suggests that there may be something different about the first part of the above proposition is that although he includes a version of it in his Proposition 3, it is absent from his Proposition 4, despite the fact that it follows from his Proposition 3 in exactly the same way as the other three parts.

4 Applying the results

This section illustrates how the distinction between the reachability graph mapping and the extended mapping of whole state space graphs affects the application of the results of the previous section to the problem of deadlock detection in coloured nets.

There are coloured nets for which Vautherin's results seem to work very well. For example, a coloured net (see [6]) for the well known 'Dining Philosophers' system known to have exactly one deadlock was found to have exactly one deadlock in its skeleton, and furthermore, the latter deadlock was the image under h of the deadlock in the coloured net.

Such an example is misleading, however, because the existence of a deadlock in the coloured net was already known. In practice, one wants to gain information about a coloured net by reachability analysis of its skeleton. Since a marking of the skeleton generally has many preimages in the coloured net, finding a skeletal deadlock reveals a number of dead (by Proposition 2) markings of the coloured net. What the analyst still does not know, however, is whether some, all, or none of these dead markings are actually reachable.

In the dining philosophers net which was analysed, all but one of the skeletal deadlock's preimages were already known to be unreachable (mainly because the coloured net was known to possess only one deadlock). Such knowledge should clearly not have to be relied upon—the idea of using skeletons is to perform faster deadlock-freeness checking in nets whose behaviour is unknown.

We will now use a simple example to illustrate graphically how the map h of the previous section really acts between whole state space graphs rather than reachability graphs, and the consequent difficulty in interpreting deadlocks of the skeleton.

Figure 1 shows a very small coloured Petri net and its skeleton (without initial markings). In the coloured net, $\{a, b\}$ is the colour set of the place.

[1] By the 'whole state space graph', we mean a graph which has as its vertices all markings of N, reachable or otherwise, and edges corresponding to possible transition firings—in other words, a sort of extended reachability graph in which no initial marking is singled out. Such a graph would include as a subgraph the reachability graph corresponding to any particular choice of initial marking.

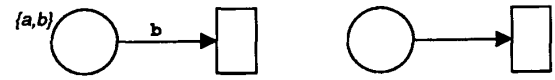

Figure 1: A very small coloured Petri net, and its skeleton (without initial markings)

Figure 2 shows (a representative portion of) the 'whole state space graph' of the coloured net, and that of its skeleton. Each node in each of these two graphs is labelled by the marking of the single place p of the two nets. The action of the map h of the previous section is to send each node of the top graph to the node of the bottom graph directly underneath it, so that, for example, $a + 2b$ is mapped to 3 (since $h(a + 2b) = |a + 2b| = 3$).

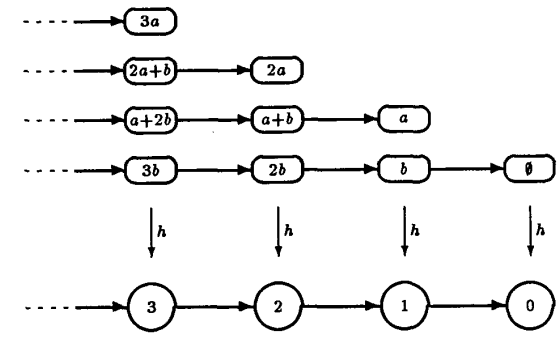

Figure 2: The mapping between two 'whole state space graphs'

Suppose now that the coloured net is given the initial marking $2a+b$. Then there is exactly one deadlock of the net, namely the marking $2a$. The corresponding marking of the skeleton is 2, which is not a deadlock. On the other hand, the skeleton does contain a deadlock (the marking 0), but the only preimage of this under h is the (dead) marking $\emptyset$ of the coloured net, which is not reachable from the chosen initial marking.

One can also see from this example that the relationship between what one detects in the skeleton and the behaviour of the coloured net is quite sensitive to the choice of initial marking of the coloured net. By way of illustration, the skeletal marking 1 of Figure 2 corresponds to the non-dead marking b reachable from $3b$, the dead marking a reachable from $a + 2b$, but no marking which is reachable from $2a + b$.

The next section looks at a more complicated example which shows how deceptive skeletons may be, in terms of how little the deadlocks of a coloured net may be related to those of its skeleton.

5 Skeletal deception

The coloured Petri net in Figure 3 is a net for a shared resource system. We will consider three differ-

ent initial markings of this net, hence no initial marking is specified in the figure.

The shared resource system operates as follows. A set Q of processes move from an IDLE state through three different stages (STG1, STG2, STG3), before returning to the IDLE state. The first two transitions in this cycle require the input of a resource R1. At the time of the third transition, a resource R2 is needed (and one of the two R1 resources already gathered is no longer needed, so becomes free again). The final transition in the cycle returns the resources still being held by a process to the free resources place RES.

```
Colour sets    Q={Q1,Q2,Q3}
               R={R1,R2}

Variables      q:Q
```

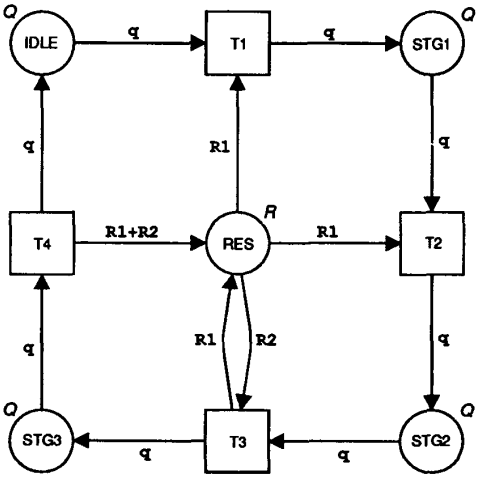

Figure 3: A net for a shared resource system

In each of the initial markings of the net which we consider, IDLE will contain three processes Q1, Q2, and Q3, while the places STG1, STG2, and STG3 will all be empty. The only difference between the three initial markings considered will be the number of resource tokens (R1 and R2) that will be made available in place RES. The three markings of RES are shown in the following table.

	M_0(RES)
First initial marking	4.R1 + 1.R2
Second initial marking	3.R1 + 4.R2
Third initial marking	2.R1 + 2.R2

In the first initial marking, the coloured net is deadlock-free, but the skeleton has one deadlock, namely the marking (0,1,2,0,0) (here we are listing the markings of individual places in the order IDLE, STG1, STG2, STG3, RES). This marking can be reached for example by firing T1 three times and then T2 twice. Note that the presence of a skeletal deadlock does not contradict the first part of Proposition 2—it is just that none of the preimages of (0,1,2,0,0) under h are reachable in the coloured net.

From the second initial marking, the system may reach a deadlock in which all three processes have moved into STG1 (by removing one R1 token each from RES), but none may move into STG2 (because RES contains no more R1 tokens). The corresponding marking of the skeleton is not a deadlock; the skeleton turns out to be deadlock-free.

Now let us consider the third initial marking. The table below shows the three reachable markings of the coloured net which are dead, as well as the two deadlocks that the skeleton possesses. Note that the skeletal image (1,2,0,0,2) of the coloured net's deadlocks is not dead, but that both the skeletal deadlocks are subsequent markings of (i.e. reachable from) this marking.

	IDLE	STG1	STG2	STG3	RES
N	Q1	Q2+Q3	$\emptyset$	$\emptyset$	2.R2
	Q2	Q1+Q3	$\emptyset$	$\emptyset$	2.R2
	Q3	Q1+Q2	$\emptyset$	$\emptyset$	2.R2
$\lvert N\rvert$	0	2	1	0	0
	1	0	2	0	0

In summary, the failure of part 1 of Proposition 2 to distinguish between reachable and unreachable markings can cause some rather inconvenient phenomena, when one attempts to detect deadlocks in coloured nets by generating the reachability graphs of their skeletons. In the case where the net contains a deadlock, the skeleton may contain no deadlocks, or may contain deadlocks which don't match the deadlock of the net. On the other hand, a net without deadlocks may have a skeleton with deadlocks.

6 Skeleton-guided reduced reachability analysis

At this point it may appear that because of their deceptive nature, skeletons are of little use in deadlock detection, due to the fact that there is no easy way of knowing whether a skeletal deadlock corresponds to a deadlock in the coloured net, without further knowledge of the latter net's behaviour. This section suggests one possible approach which may be useful in making use of the results obtained from skeletal reachability analysis.

Suppose one is given a (uniform) coloured net which one wishes to test for deadlocks. After generating the reachability graph of its skeleton, what should one do? If the skeleton is found to be deadlock-free, then the results as they stand do not allow any conclusions to be drawn about the coloured net.

On the other hand, if a deadlock is found in the skeleton, a set of dead markings of the coloured net is revealed. Even though one does not know which of these are reachable, Proposition 1 can be applied to deduce that *if there is a path to any one of them in the coloured net's reachability graph, then the image*

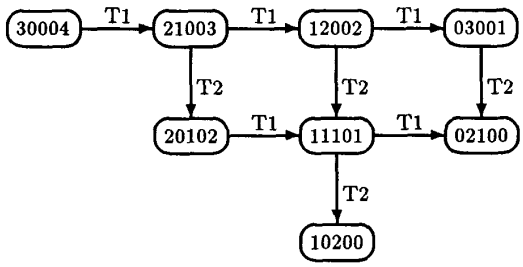

Figure 4: A portion of the reachability graph of the skeleton of the shared resource system net (with the third initial marking)

of this path under h will be present in the reachability graph of the skeleton.

Thus, generating that portion of the coloured net's reachability graph which h projects to the path(s) leading to the skeletal deadlock is an approach which is guaranteed to tell us whether any of the skeletal deadlock's coloured preimages are in fact reachable.

Now, as an illustration of something else the above approach may reveal, let us turn our attention back to the shared resource net of the previous section, with the third initial marking—the one for which both coloured net and skeleton contain deadlocks.

Figure 4 shows the subgraph of the skeleton's reachability graph consisting of all paths (without repeated markings) leading from the initial marking (3,0,0,0,4) to the two skeletal deadlocks (0,2,1,0,0) and (1,0,2,0,0).

Generating the corresponding subgraph of the coloured net's reachability graph (i.e. the reachable portion of the preimage under h of Figure 4) would reveal that none of the preimages of these two deadlocks are reachable. However, it would also reveal the three dead markings of the coloured net which *are* reachable, since we have already noted in the previous section that these three deadlocks are preimages of (1,2,0,0,2), which appears as one of the markings 'on the way' to the deadlocks in Figure 4.

In other words, although the skeletal deadlocks do not have reachable preimages, they are still useful, because they may be viewed as 'postponed versions' of the deadlocks in the coloured net. These coloured deadlocks may therefore be found by performing 'skeleton-guided' reduced reachability analyis of the coloured net.

The nets of Figure 1 provide another example where the conversion from coloured net to skeleton 'shifts' a deadlock rather than eliminating it—one can see from the graphs of Figure 2 that for *any* initial marking of the coloured net, there is exactly one deadlock, and that it will always appear as a postponed deadlock in the skeleton.

Of course, since a coloured net with deadlocks may have a deadlock-free skeleton, while a skeleton may contain deadlocks which are not related to any deadlocks in the corresponding coloured net, there will be some examples for which the above approach will not help to detect deadlocks of the coloured net.

It does appear, however, that a skeleton-guided reduced reachability analysis approach should be able to aid deadlock detection in at least some examples.

7 Future Work

What we have revealed in the last three sections indicates that there are many questions which need to be answered before skeletons can really be used practically for detecting deadlocks of coloured nets. This section discusses these questions.

Regarding the skeleton-guided reduced reachability analysis approach suggested in the previous section, we really need some more solid information. As it stands, this approach can provide no guarantee of results; any usefulness it might have is due only to the fact that it may detect (some) deadlocks of a coloured net more quickly than a conventional reachability analysis approach.

One possible topic for investigation is whether a class of nets which exhibit postponed deadlocks could be determined. This problem is not likely to be easy, however, since the presence of postponed deadlocks in skeletons is very sensitive to the choice of initial marking—it is possible to have a coloured net with a deadlock which appears as a postponed version in the skeleton, but such that adding a token to some place in the (coloured) initial marking causes the skeletal deadlock to disappear, without removing the coloured one.

What would be markedly more valuable than an answer to the problem posed in the above paragraph, is a necessary and sufficient condition for a coloured net to have the property that a marking is dead *if and only if* the corresponding skeletal marking is dead. For any net with this property, Proposition 1 could be applied to guarantee that any reachable dead marking will have a reachable dead skeletal image, with the consequence that a deadlock-free skeleton implies that the coloured net is also deadlock-free.

Vautherin describes a sufficient condition in Proposition 5 of his paper, the gist of which is that the input token requirements of each transition must be (in Vautherin's words) 'independent and non-selective with regard to the colours'. 'Independent' means that the net in Figure 5 would be unacceptable, because the x token taken in by the transition from the place on the left is required to match the x token removed from the place on the right. 'Non-selective with regard to the colours' means that the requirement that one of the tokens removed from the right-hand place in that figure must be an 'a' is also unacceptable. Figure 6 shows a net which does satisfy Vautherin's condition.

There are simple examples which demonstrate that Vautherin's condition is not necessary to obtain the aforementioned property. At the time of writing, the author has made significant progress on the problem of determining a *necessary and sufficient* condition [3].

For coloured nets which do not have the above property, it would be useful to know whether one can easily determine the set of initial markings for which every

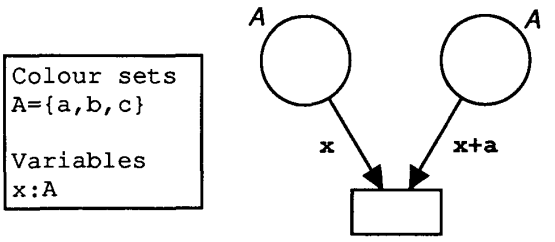

Figure 5: A coloured Petri net which does not satisfy Vautherin's condition

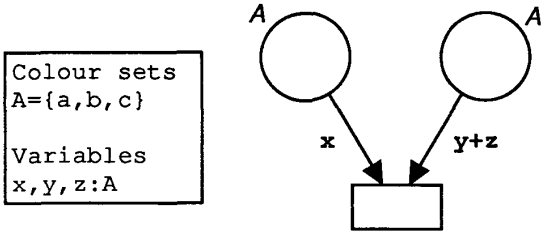

Figure 6: A coloured Petri net which does satisfy Vautherin's condition

reachable dead marking can be guaranteed to have a dead skeletal image.

Another problem on which the author is currently working [3] is the following: Given a coloured Petri net for which some dead markings have live skeletal images, i.e. without the property discussed above, how much does it need to be unfolded to get a coloured net which *does* have that property?

8 Conclusion

Although Vautherin [7] only defines skeletons for schemas, they can be satisfactorily defined for (and easily obtained from) uniform coloured Petri nets, as indicated in [5].

A relationship (given by a simply defined mapping) between the 'whole state space graph' of a coloured Petri net and that of its skeleton is easily established. Vautherin showed that under this mapping, a dead marking of the skeleton can have only dead preimages in the coloured net. However, a dead marking of the coloured net need not have a dead skeletal image, so that the mapping only sends the reachability graph of a coloured net *into* that of its skeleton.

We have shown that as a result, deadlock-freeness (i.e. the property of having no reachable dead markings) of a coloured net and deadlock-freeness of the corresponding skeleton are independent properties—a deadlock-free coloured net may have a skeleton containing deadlocks, or vice-versa.

Because of such phenomena, Vautherin's result as it stands (i.e. with no distinction made between reachable and unreachable markings) is much less useful than it might appear, for the purpose of using skeletons of coloured Petri nets to detect their deadlocks. However, a skeleton-guided reduced reachability analysis approach should be able to be assist in deadlock detection for some nets.

Section 7 indicates the work which still needs to be done to determine whether skeleton-based, or at least skeleton-assisted, deadlock detection can be made into a more profitable approach.

Acknowledgements

The author would like to acknowledge the fruitful discussions held with his colleagues, Jon Billington, Geoff Wheeler, and Brian Keck, during the work on this paper.

The anonymous reviewers' comments on the originally submitted version of the paper have also been valuable.

The permission of the Executive General Manager Research, Telecom Australia, to present this paper is hereby acknowledged.

References

[1] J. Billington. Many-Sorted High-Level Nets. In *Proceedings of the Third International Workshop on Petri Nets and Performance Models, Kyoto, Japan, 11–13 December 1989*, IEEE CS Press, Washington , D.C., USA, 1989.

[2] H. Ehrig and B. Mahr. *Fundamentals of Algebraic Specification 1, Equations and Initial Semantics.* Volume 6 of *EATCS Monographs on Theoretical Computer Science*, Springer-Verlag, Berlin, 1985.

[3] G. Findlow. *Using Skeletons for Deadlock Detection in TORAS.* Draft Telecom Australia Research Laboratories Report.

[4] K. Jensen. High-level Petri Nets. Proceedings of the 3rd European Workshop on Applications and Theory of Petri Nets, Varenna, Italy, 1982. In A. Pagnoni and G. Rozenberg (eds.), *Applications and Theory of Petri Nets*, Inf.-Fachberichte Vol. 66, Springer-Verlag, 1983.

[5] K. Jensen. Coloured Petri Nets. In W. Brauer, W. Reisig, and G. Rozenberg (eds.), *Petri Nets: Central Models and Their Properties, Advances in Petri Nets 1986, Part I.* Lecture Notes in Computer Science, Vol. 254, pp. 248–299, Springer-Verlag, Berlin, 1987.

[6] A. Valmari. Stubborn Sets of Coloured Petri Nets. In *Proceedings of the 12th International Conference on Application and Theory of Petri Nets*, Aarhus, Denmark, 26–28 June 1991.

[7] J.Vautherin. Parallel systems specifications with Coloured Petri Nets and algebraic specifications. In G. Rozenberg, editor, *Advances in Petri Nets 1987.* Lecture Notes in Computer Science, Vol. 266, pp. 293–308, Springer-Verlag, Berlin, 1987.

SEARCHING BEST PATHS TO WORST STATES

G. Florin + C. Fraize+* S. Natkin +
+ CEDRIC - Centre d'Etudes et de recherche
en Informatique CNAM
292 rue Saint-Martin 75141 Paris Cedex 03
* GEC-ALSTHOM Tour Neptune La Defense Paris

Abstract

Probabilistic validation is a new approach to deal with large state transitions systems. The user's need is to prove that, for a given period of operations, that a given assertion on the reached states is true with a sufficient level of probability.

The system to be validated is modeled by a stochastic Petri net. The analysis relies on a partial exploration of the reachability set and tries to reach as quickly as possible critical states (states in which the assertion is not verified). An exact linear program solution allows to "travel" through the graph. The main goal of this paper is to present the principles of this searching algorithm.

This method can be used in probability computations in two ways. The first one is related to acyclic graphs. A breadth or a depth first search traversal can be done without considering all the trajectories but only those leading to critical states. The second one is related to importance sampling simulation.

Keywords : safety critical systems, distributed systems, probability, simulation, stochastic Petri nets, validation.

1 Introduction

The validation of complex systems such as safety-critical or distributed systems is a fundamental goal of dependable system design. It needs a formal model (defined by a user) of the system behavior and a software tool able to prove assertions on this model.

Unfortunately exhaustive analysis of the transition systems state space is limited by the exponential growth of the model complexity. Simulation of the behavior does not give an indicator of the validation level [3].

The aim of probabilistic validation [7] is to develop a new approach based on a partial analysis of a system model. The model and the analysis method must allow to prove that assertions are verified with an acceptable probability level.

Probabilistic validation tries to use the fact that "with a very low probability in a given period" is sometimes simpler to prove than "never". For the designer or the user, a system is often sufficiently dependable if undesirable events are never observed during an operational period. An operational validation method tries to prove that "bad events" will occur with a sufficiently low probability during a given period.

Several works are related to our approach. Some authors try to relate simulation and proof [8]. Our approach is also related to approximate solution of stochastic models and in particular stochastic Petri nets [9], [1]. Principles and methods able to evaluate complex user oriented goals, called performability measures, of a system have been developed by Meyer [12]. Maxemchuk presents a computation method for protocol probabilistic validation [11].

The principles of probabilistic validation are presented in [7]. Only basic principles are recalled in the second section. The goal of this paper is to present the techniques used to limit this scanning (classified in two classes): qualitative techniques presented in section 3 and quantitative techniques defined in section 4. In the last section a particular algorithm is presented.

2 Probabilistic validation

2.1 Time dependent validation

In state transition models or even temporal logic, the time is expressed by ordering conditions between events. Operations duration and frequencies are not explicitly described.

It has been formally proved that several deterministic distributed protocols do not tolerate failures with-

TH0386-3/91/0000/0204/$01.00 © 1991 IEEE