

Checking System Properties via Integer Programming*

Stephan Melzer and Javier Esparza

Institut für Informatik
Technische Universität München
Arcisstr. 21, D-80290 München
e-mail: {melzers,esparza}@informatik.tu-muenchen.de

Abstract. The marking equation is a well known verification method in the Petri net community. It has also been applied by Avrunin, Corbett *et al.* to automata models. It is a semidecision method, and it may fail to give an answer for some systems, in particular for those communicating by means of shared variables. In this paper, we complement the marking equation by a so called trap equation. We show that both together significantly extend the range of verifiable systems by conducting several case studies.

1 Introduction

The use of linear algebra and integer programming for verification purposes has a long tradition in Petri net theory [6, 19, 18, 20]. One of the best known techniques is the *state* or *marking equation* [6, 20]. This is a linear equation which can be easily derived from the description of the net and its initial marking (in linear time). It can be seen as a set of linear constraints \mathcal{L} that every reachable marking must satisfy. In other words, the solutions of \mathcal{L} are a superset of the reachable markings. In order to use the marking equation, we add to it new linear constraints \mathcal{L}_P , which specify the markings which do *not* satisfy a desirable property P.² Then, we use integer programming to solve the system $\mathcal{L} \cup \mathcal{L}_P$: if the system has no solution, every reachable marking satisfies P.

The disadvantage of this technique is the fact that the markings satisfying \mathcal{L} are only a superset of the reachable markings: the solutions of $\mathcal{L} \cup \mathcal{L}_P$ may or may not correspond to a reachable marking. Therefore, the marking equation is only a *semidecision method*. Its main advantage is that it does not explore the state space, and therefore it avoids the state explosion problem. It can also be used to verify systems having infinite state spaces.

The marking equation can be applied to many different models of concurrency, not only to Petri nets. Actually, the most comprehensive study of its

* This work was partially supported by the Sonderforschungsbereich SFB-342 A3 SAM.

² It is also possible to impose linear constraints on the occurrence sequence leading to those markings. This is a very useful feature, but we omit it here for the sake of simplicity.

applications for verification has been carried out by Avrunin, Corbett *et al.* using coupled automata as a model [2, 3, 8]. They have developed the Constrained Expression Toolset, later updated to the Inequality Necessary Condition Analyzer (INCA), a tool for the verification of a large class of safety and liveness properties. It is easy to see that the basis of the technique implemented in INCA is equivalent to the marking equation. In [7], Corbett shows that INCA is able to prove deadlock freedom for 19 different examples taken from different sources, and can compete with symbolic and partial order theorem provers.

One of the main limitations of the marking equation is that it tends to fail for systems which communicate via shared variables. For instance, it cannot prove mutual exclusion of any of the most popular mutual exclusion algorithms (Dekker's, Dijkstra's, Knuth's, Peterson's etc.) without user's help. The reason is that the method is not sensitive to the guards which allow to perform an action only if a variable has a certain value, in the sense that the systems with or without the guards are assigned the same set of constraints. Since the correctness of these algorithms crucially depends on these guards, the method fails.

In this paper, we show how to obtain a set of constraints which better approximate the set of reachable markings, and are sensitive to these guards. We then test the improved algorithm on a number of examples. In particular, we automatically prove mutual exclusion of five mutual exclusion algorithms.

This refined set of constraints is derived from some results of Petri net theory concerning so called *traps*. Therefore, it is convenient to present our results in Petri net terms. However, there would be no problem in recasting them for, say, the communicating automata of Corbett [7], the synchronized products of transition systems of Arnold and Nivat (see, for instance, [1]), or for CCS processes of the form $(P_1 \mid \dots \mid P_n) \setminus L$, where the P_i are regular. All of them can be easily translated into (1-safe) Petri nets. The common idea of the translations is simple: each sequential component is modelled by means of a Petri net, just mapping states to places and transitions of the transition system into transitions of the Petri net. Communication is then modelled by merging transitions.

Linear upper approximation of the set of reachable states have also been used by Cousot and Halbwachs and others in the field of abstract interpretation [9, 15]. The main difference with our approach is that we derive the linear approximation directly from our structure of the system, in one single step, and not by means of successive approximations, as in [9].

Our paper is organised as follows. In Section 2 we introduce some basic definitions. Section 3 describes the marking equation. In Section 4 we introduce traps, and present our improved method. In Section 5 we apply the results to examples. In Section 6 we present a result on checking deadlock freedom. Finally, we present our conclusions in Section 7.

2 Basic notations

A *net* is a triple $N = (P, T, W)$ where $P \cap T = \emptyset$ and $W: (P \times T) \cup (T \times P) \rightarrow \mathbb{N}$. P is the set of *places* (symbolized by circles), T the set of *transitions* (symbolized

by rectangles) and W is the *weight function*. The *pre-set* of $x \in PUT$ is ${}^*x = \{y \in PUT \mid W(y, x) > 0\}$. The *post-set* of $x \in PUT$ is $x^\bullet = \{y \in PUT \mid W(x, y) > 0\}$. The pre- and post-set of a subset of PUT are the union of the pre- and post-sets of its elements.

All the examples of Section 5 (and all the examples of [7]) can be modelled by *ordinary* nets, in which the weight function has codomain $\{0, 1\}$. However, more general weight functions play an important role in the development of the results of Section 4, and that is why we define nets in this generality.

A function $M : P \rightarrow \mathbb{IN}$ is called a *marking*. A *Petri net* is a pair (N, M_0) where N is a net and M_0 a marking of N called *initial marking*. A transition $t \in T$ is *enabled at M* iff $\forall p \in {}^*t : M(p) \geq W(p, t)$. If t is enabled at M , then t may *fire* or *occur*, yielding a new marking M' (denoted $M \xrightarrow{t} M'$), where $M'(p) = M(p) + W(t, p) - W(p, t)$.

A sequence of transitions, $\sigma = t_1 t_2 \dots t_r$ is an *occurrence sequence* of (N, M_0) iff there exist markings M_1, \dots, M_r such that $M_0 \xrightarrow{t_1} M_1 \xrightarrow{t_2} M_2 \dots \xrightarrow{t_r} M_r$. The marking M_r is said to be *reachable* from M_0 by the occurrence of σ (denoted $M \xrightarrow{\sigma} M_r$).

A Petri net (N, M_0) is *safe* iff $M(p) \leq 1$ for every reachable marking M and every place p .

A *linear programming problem* or linear problem is a system $A \cdot X \leq B$ of linear (in-)equalities called the *constraints*, plus maybe a linear function $C^T \cdot X$ called the *objective function*. A solution of the problem is a vector of rational numbers that satisfy the constraints. A solution is *optimal* if it maximises the value of the objective function (over the set of all solutions).

An *integer programming problem* consists of the same elements as a linear programming problem, but only integer solutions are allowed. In a *mixed programming problem*, some variables may take rational values, and some only integer ones.

A linear, integer or mixed programming problem is *feasible* if it has a solution. Otherwise it is infeasible.

3 The marking equation

Each place p of a net has associated a *token conservation* equation. Given an occurrence sequence $M_0 \xrightarrow{\sigma} M$, the number of tokens that p contains at the marking M is equal to the number of tokens it contains at M_0 , plus the tokens added by (the firings of) the input transitions of p , minus the tokens removed by the output transitions. If we denote by $\#(\sigma, t)$ the number of times that a transition t occurs in σ , we can write the *token conservation* equation for p as:

$$M(p) = M_0(p) + \sum_{t \in {}^*p} \#(\sigma, t)W(t, p) - \sum_{t \in p^\bullet} \#(\sigma, t)W(p, t)$$

The token conservation equations for every place are usually written in the following matrix form:

$$M = M_0 + N \cdot \vec{\sigma}$$

where $\vec{\sigma} = (\#(\sigma, t_1), \dots, \#(\sigma, t_m))$ is called the *Parikh vector* of σ , and \mathbf{N} denotes the *incidence matrix* of N , a $P \times T$ integer matrix given by

$$\mathbf{N}(p, t) = W(p, t) - W(t, p)$$

If a given marking M is reachable from M_0 , then there exists a sequence σ satisfying $M_0 \xrightarrow{\sigma} M$. So the following problem has at least one solution, namely $X := \vec{\sigma}$.

Variables: X , integer.

$$M = M_0 + \mathbf{N} \cdot X$$

$$X \geq 0$$

The equation $M = M_0 + \mathbf{N} \cdot X$ (and, by extension, the whole problem) is called the *marking equation*. If the marking equation has no solution, then M is not reachable from M_0 .

We wish to verify that every reachable marking satisfies a desirable property, or, equivalently, that no marking satisfying the negation of this desirable property is reachable. The negation of the property can often be expressed by means of *linear constraints* on the markings of the net. Here are two examples:

– Mutual exclusion.

In Petri net models of mutual exclusion algorithms the possible states of a process (idle, requesting, critical, ...) are modelled by places which can hold at most one token. The process is in the critical section if the corresponding place is marked. If s_1, \dots, s_n are the places corresponding to the critical sections, then the reachable markings that violate the mutual exclusion property are those satisfying

$$M(s_1) + \dots + M(s_n) \geq 2$$

– Deadlock freedom in safe Petri nets.

A marking is a deadlock if it does not enable any transition. In safe Petri nets a place can hold at most one token, and therefore a transition is enabled if and only if the total number of tokens in its input places is at least equal to the number of input places. In other words, the reachable deadlocked markings satisfy

$$\sum_{s \in {}^*t} M(s) < |{}^*t|$$

for every transition t .

A *linear property* of N is a predicate \mathcal{P} on the markings of N (or, equivalently, a subset of the markings of N) such that

$$\mathcal{P}(M) \Leftrightarrow A \cdot M \leq b$$

for some matrix A and vector b . We can use the marking equation to verify properties whose negation is linear. If some marking satisfying \mathcal{P} is reachable from M_0 , then the *generalised marking equation*

$$\begin{aligned}
 &\text{Variables: } M, X: \text{integer} \\
 &M = M_0 + N \cdot X \\
 &A \cdot M \leq b \\
 &M, X \geq 0
 \end{aligned}$$

has a solution.³ Therefore, if the generalised marking equation is infeasible, every reachable marking satisfies the negation of \mathcal{P} . We can use integer programming to check infeasibility.

The implication “infeasibility $\Rightarrow \neg\mathcal{P}$ holds for every reachable marking” still holds if M and X are allowed to take rational values. So, in principle, one may try to use ordinary linear programming to check infeasibility. Unfortunately, the experiments show that in most cases even though the desired property holds, the marking equation has non-integer solutions, and therefore linear programming is of little use. Using integer programming leads to much better results [7, 8].

Unfortunately, the marking equation still fails very often when the Petri net models a distributed system with shared variables. The components of this kind of systems test the value of a variable to determine the flow of control. Now, consider the two Petri nets of Figure 1. The Petri net on the left models a component which may change state, from s_0 to s_1 , only if the variable x has value 0, which happens not to be the case. In the Petri net on the right, the component can change its state independently of the value of x . Obviously, the marking $\{s_1\}$ is not reachable on the left, and reachable on the right. However, the marking equations of these two nets *coincide*. Therefore, the generalised marking equation cannot be used to prove that $\{s_1\}$ is not reachable on the left.




Fig. 1. A limitation of the marking equation

We could of course prove this by constructing the reachability graph, which is very small in this example, but may grow exponentially in the size of the net (or be infinite). An alternative is the use of *traps* [24, 11].

Definition 1. Traps

A set R of places of a net is a trap if $R^\bullet \subseteq {}^\bullet R$. ■ 1

³ Since M is in fact a linear function of X , it would still be more general to add a constraint of the form $C \cdot X \leq d$, and this is in fact the approach of [8]. Since the examples of this paper only consider constraints on markings, we will use the constraint shown above for clarity.

In the sequel, we shall use the letter Θ to denote traps. Traps have the following fundamental property:

Proposition 2. *Marked traps remain marked*

Let (N, M_0) be a Petri net, and let Θ be a trap of N . If Θ is marked at M_0 (i.e., if $\sum_{p \in \Theta} M_0(p) > 0$), then Θ remains marked at every reachable marking. ■ 2

The set $\{s_0, s_2\}$ is a trap of the net on the left, and this trap is marked at the initial marking $\{s_0\}$. However, the trap is not marked at $\{s_1\}$. Therefore, the marking $\{s_1\}$ is not reachable.

If a marking marks every trap that is marked at M_0 we say that it satisfies the *trap property*. Proposition 2 states that, on top of the marking equation, a reachable marking must satisfy the trap property as well. We have thus a refined test of non-reachability.

In order to check that every marking satisfying a linear property \mathcal{P} violates the trap property we may compute all the traps marked at M_0 , say $\Theta_1, \dots, \Theta_n$, and then compute iteratively the subsets \mathcal{P}_i of \mathcal{P} that mark the traps $\Theta_1, \dots, \Theta_i$ for $1 \leq i \leq n$. However, this method is very inefficient, because the number of traps may be exponential in the size of the net⁴. In order to make traps useful for automatic verification, we have to find an alternative, which we present in the next section.

4 The trap equation

In this section we obtain the *generalised trap equation* for a linear property \mathcal{P} . This is a linear equation which has a solution if and only if no marking satisfies simultaneously \mathcal{P} and the trap property.

The first step towards our goal is to find a link between traps and linear algebra. Fortunately, we can profit from several existing results. In [17], Lautenbach showed that there exists a tight relation between the traps of a net N and the solutions of the equation $Y^T \cdot N_\Theta = 0$, where N_Θ is obtained from N by means of a relatively complicated transformation. Later, Lautenbach's results were used and slightly improved by Esparza and Silva in [12]. Finally, Ezpeleta, Couvreur and Silva found another improvement [13]. They showed that Lautenbach's net N_Θ can be replaced by a simpler one. N and the new N_Θ have the same places, transitions and arcs: they only differ in the *weights* of some arcs leading from transitions to places.

⁴ In fact, it suffices to compute all minimal traps, which are the nonempty traps not included in any other trap. However, there may also be exponentially many minimal traps.

Theorem 3. Algebraic characterization of traps [13]

Let $N = (P, T, W)$ be a net. Let $N_\Theta = (P, T, W_\Theta)$, where

$$W_\Theta(p, t) = W(p, t)$$

$$W_\Theta(t, p) = \begin{cases} \sum_{p' \in \bullet t} W(p', t) & \text{if } p \in t^\bullet \\ 0 & \text{otherwise} \end{cases}$$

A set $\Theta \subseteq P$ is a trap of the net N if and only if the equation $Y^T \cdot N_\Theta \geq 0$ has a nonnegative solution Y such that $\|Y\| = \Theta$. ■ 3

We illustrate this result on the Petri net of Figure 2. The vectors Y_1^T and Y_2^T satisfy the equation of Theorem 3, and therefore $\{s_3, s_5\}$ and $\{s_3, s_4, s_6\}$ are traps of the net. The vector Y_3^T does not satisfy it, and in fact $\{s_1, s_2\}$ is not a trap.




Fig. 2. An example.

$$N_\Theta = \begin{pmatrix} -1 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 \\ 1 & 0 & 1 & -1 & 0 \\ 0 & 1 & 0 & -1 & 1 \\ 0 & 0 & -1 & 2 & 0 \\ 0 & 0 & 0 & 2 & -1 \end{pmatrix} \quad \begin{matrix} Y_1^T = (1, 0, 1, 0, 1, 0) \\ Y_2^T = (0, 0, 1, 1, 0, 1) \\ Y_3^T = (1, 1, 0, 0, 0, 0) \end{matrix}$$

We can use Proposition 2 to test if a marking M violates the trap property.

Proposition 4.

Let (N, M_0) be a Petri net, and let M be a marking of N . M satisfies the trap property if and only if the problem below is infeasible.

Variables: Y : rational.

$$Y^T \cdot N_\Theta \geq 0$$

$$Y \geq 0 \ (\Theta = \|Y\| \text{ is a trap})$$

$$Y^T \cdot M_0 > 0 \ (\Theta \text{ is initially marked})$$

$$Y^T \cdot M = 0 \ (\Theta \text{ is not marked at } M)$$

Proof: By Theorem 3, a solution of the problem corresponds to a trap initially marked, but unmarked at M , and vice versa. ■ 4

Now, in order to test if M violates the trap property we solve a linear programming problem instead, which intensionally checks if *every* initially marked trap remains marked at M .

However, Proposition 4 is not directly useful when we consider linear properties. If M becomes a variable subject to the linear condition $A \cdot M \leq b$, then the equation $Y^T \cdot M = 0$ becomes non-linear, which very much complicates the verification. To remove this difficulty we shall use one of the many versions of the Minkowski-Farkas Lemma (see, for instance, [25]).

Theorem 5. Minkowski-Farkas Lemma

One and only one of the following two problems is feasible:

$$\begin{array}{ll} \text{Variables: } X: \text{ rational.} & \text{Variables: } Y: \text{ rational.} \\ A \cdot X \leq b & Y^T \cdot A \geq 0 \\ X \geq 0 & Y^T \cdot b < 0 \\ & Y \geq 0 \end{array}$$

In order to apply this theorem, we first have to modify the problem of Proposition 4. We observe that, since M is a nonnegative vector and any solution Y must also be nonnegative, the constraint $Y^T \cdot M = 0$ can be safely replaced by $Y^T \cdot M \leq 0$. So the problem is equivalent to (i.e., has the same solutions as):

$$\begin{array}{l} \text{Variables: } Y: \text{ rational.} \\ Y^T \cdot (\mathbf{N}_\ominus | -M) \geq 0 \\ Y^T \cdot (-M_0) < 0 \\ Y \geq 0 \end{array}$$

where $(\mathbf{N}_\ominus | -M)$ denotes the matrix obtained by adding $-M$ to \mathbf{N}_\ominus as rightmost column.

Now, by Proposition 4 and the Minkowski-Farkas Theorem, M satisfies the trap property if and only if the following problem is feasible:

$$\begin{array}{l} \text{Variables: } X: \text{ rational.} \\ (\mathbf{N}_\ominus | -M) \cdot X \leq -M_0 \\ X \geq 0 \end{array}$$

Notice that the dimension of X is equal to the number of transitions of N plus 1, because of the addition of the column M . Define $X = (X' | x)$, i.e., X' is the vector containing all the components of X but the last, and x is the last component of X . With these notations, we can rewrite the problem as:

$$\begin{array}{l} \text{Variables: } X', x: \text{ rational.} \\ xM \geq M_0 + \mathbf{N}_\ominus \cdot X' \\ X', x \geq 0 \end{array}$$

Assume that this problem has a solution for $x = 0$. Then, since M is nonnegative, it also has a solution for every $x > 0$. So we can replace $x \geq 0$ by $x > 0$, and the resulting problem is still feasible if and only if M satisfies the trap property. Now, since $x > 0$, we can divide the first inequality by it. Redefining $X := \frac{1}{x}X'$ and then $x := \frac{1}{x}$, we finally get the *trap equation*:

$$\begin{aligned} \text{Variables: } & M \text{ :integer; } X, x \text{ :rational.} \\ & M \geq xM_0 + \mathbf{N}_\Theta \cdot X \\ & X \geq 0 \\ & x > 0 \end{aligned}$$

We have reached our goal: the trap equation is linear, and M appears isolated on the left side, as in the marking equation. We can thus generalise it to linear properties by adding the constraint $A \cdot M \leq b$.

Theorem 6. *Generalised trap equation*

Let (N, M_0) be a Petri net, and let P be a linear property of the markings of N , characterised by the equation $A \cdot M \leq b$. If the problem below is infeasible, then no marking satisfies both P and the trap property.

$$\begin{aligned} \text{Variables: } & M \text{ :integer; } X, x \text{ :rational} \\ & M \geq xM_0 + \mathbf{N}_\Theta \cdot X \\ & A \cdot M \leq b \\ & M, X \geq 0 \\ & x > 0 \end{aligned}$$

■ 6

Finally, putting together the marking and trap equations we obtain a negative test for linear properties:

Corollary 7.

Let (N, M_0) be a Petri net, and let P be a linear property of the markings of N , characterised by the equation $A \cdot M \leq b$. If the problem below is infeasible, then every reachable marking satisfies the negation of P .

$$\begin{aligned} \text{Variables: } & M, X_1 \text{ :integer; } X_2, x \text{ :rational} \\ & M = M_0 + \mathbf{N} \cdot X_1 \\ & M \geq xM_0 + \mathbf{N}_\Theta \cdot X_2 \\ & A \cdot M \leq b \\ & M, X_1, X_2 \geq 0 \\ & x > 0 \end{aligned}$$

■ 7

This problem can be solved using *mixed programming*, a combination of linear and integer programming. Mixed programming solves systems of the form $A \cdot X \leq b$, where part of the variables are required to take integer values, while others may be rational. The constraint $x > 0$ does not fit in this format, but this

problem can be easily solved making use of the optimization facilities of mixed programming solvers: we solve the system with $x \geq 0$ as constraint, but search for the solution with maximal value of x . If this value is 0, then the original problem is infeasible.

5 Examples

In this section we show that a number of properties of several systems that could not be verified by the marking equation alone can be verified by the combination of the marking equation and the trap equation.

As a first case study, we consider five popular mutual exclusion algorithms taken from [23], namely those by De Bruijn, Dekker, Dijkstra, Knuth and Peterson. For each of them we verify deadlock freeness and mutual exclusion.

The algorithms are easily encoded in $B(PN)^2$ (Basic Petri net Programming Notation), an imperative language designed to have a simple Petri net semantics [5]. 1-safe Petri nets are then automatically generated by the PEP-tool [4]. We then generate the corresponding mixed problems, which are solved using CPLEXTM (version 3.0) [10] on a SUN SPARC 20/712.

None of the properties can be proved using linear programming. However, we do not have to require both M and X_1 to be integer in Corollary 7: it suffices to require it for M . The results of the two tables below correspond to this case.

In the table on the left we have considered algorithms for two processes. On the right we have considered Dijkstra's algorithm for n processes.

Both tables have the same structure. The first column shows the name of example, e.g. *Dijkstra 5* means Dijkstra's mutex algorithm for 5 processes. The next two numbers indicate the number of places and transitions of the Petri net. PEP generates a number of redundant places and transitions, which have not been removed for the case study. The fourth column describes the verified property: Deadlock (actually deadlock-freedom) or Mutex (mutual exclusion). The next column shows which constraints were needed to verify the property: ME (marking equation) or ME + TE (marking equation plus trap equation). The last column gives the CPU time in seconds.

Example	P	T	Property	Program	Time
Dekker	50	75	Mutex	TE + ME	0.27
			Deadlock	TE + ME	0.61
Peterson	40	69	Mutex	TE + ME	0.31
			Deadlock	ME	0.44
Dijkstra 2	64	89	Mutex	TE + ME	0.22
			Deadlock	ME	0.25
Knuth 2	74	140	Mutex	TE + ME	0.67
			Deadlock	ME	0.67
De Bruijn 2	80	166	Mutex	TE + ME	0.91
			Deadlock	ME	1.09

Example	P	T	Property	Program	Time
Dijkstra 2	64	89	Mutex	TE + ME	0.22
			Deadlock	ME	0.25
Dijkstra 3	98	160	Mutex	TE + ME	5.02
			Deadlock	ME	0.88
Dijkstra 4	134	257	Mutex	TE + ME	28.50
			Deadlock	ME	1.55
Dijkstra 5	172	386	Mutex	TE + ME	120.12
			Deadlock	ME	10.45
Dijkstra 6	212	553	Mutex	TE + ME	144.37
			Deadlock	ME	53.30

The next table shows results for a slotted ring protocol described in [21], in which n processes are placed in a ring. In [21] the state space of the example was encoded into BDDs – Binary Decision Diagrams – and then used to check

different properties, one of which was deadlock freedom. The construction of the BDD for a ring of 9 processes (the largest ring considered in [21]) took 4080 seconds. Using our method we can prove deadlock-freedom in 0.68 seconds. The trap equation is not needed in this case. The example shows that linear constraint methods can compete with symbolic model checkers (there exist other examples (see [7]) in which BDD methods are more efficient).

Example	$ P $	$ T $	Property	Program	Time
Slotted Ring 2	20	20	Deadlock	ME	0.02
Slotted Ring 3	30	30	Deadlock	ME	0.03
Slotted Ring 4	40	40	Deadlock	ME	0.03
Slotted Ring 5	50	50	Deadlock	ME	0.07
Slotted Ring 6	60	60	Deadlock	ME	0.20
Slotted Ring 7	70	70	Deadlock	ME	0.32
Slotted Ring 8	80	80	Deadlock	ME	0.63
Slotted Ring 9	90	90	Deadlock	ME	0.68
Slotted Ring 10	100	100	Deadlock	ME	2.72

Finally, we consider a less academic example. We prove deadlock freedom of two versions of a call handling for intelligent telephone networks which is closely related to a *Basic Call State Model* [22] of the ITU-T (former CCITT) standardization committee. The systems are described in [16]. We have used the $B(PN)^2$ translations of [14]. The first version (Telephone) is the original protocol, while the second version (Telephone (par)) is a refinement which allows parallel communications.

Example	$ P $	$ T $	Property	Program	Time
Telephone	87	188	Deadlock	ME + TE	10.82
Telephone(par)	232	672	Deadlock	ME + TE	705.68

6 Siphons

In Petri net theory, traps are usually studied together with siphons [24, 11]. The results of Section 4 lead to ‘dual’ results about siphons. We study their possible applications in this section.

Definition 8. *Siphons, proper siphons*

A set R of places of a net is a siphon if ${}^*R \subseteq R^*$. A siphon is called proper if it is not the empty set. ■ 8

In the sequel, we shall use the letter Σ to denote siphons. Since a transition which puts tokens in the places of a siphon also removes tokens from them, we have the following fundamental property:

Proposition 9. *Unmarked siphons remain unmarked*

Let (N, M_0) be a Petri net, and let Σ be a siphon of N . If Σ is unmarked at M_0 , then Σ remains unmarked at every reachable marking. ■ 9

Proposition 9 provides a further negative test for reachability: if M marks some siphon unmarked at M_0 , then M is not reachable. Using another version of the Alternatives Theorem we can obtain a siphon equation, which may be added to the marking and trap equations. However, the siphon equation has little interest. The reason is the following: since a siphon Σ unmarked at M_0 remains unmarked, no transition of Σ^\bullet can ever occur. This is usually undesirable and a very serious design error. In all the Petri net models we have considered so far (correct or incorrect), the initial marking marks every siphon, and so the siphon equation does not add discriminating power.

Siphons do help in a different way. In Section 3 we showed that the set of deadlocked markings of a Petri net that put at most one token on a place is linear. It is easy to see that this property ceases to hold if the deadlocks may put more than one token. In general, all we can say is that the set of deadlocks is the union of a finite number of linear sets, namely those characterised by equations of the form

$$M(s_1) + \dots + M(s_n) = 0$$

where the set $\{s_1, \dots, s_n\}$ contains exactly one input place of each transition. So in principle we could verify deadlock freedom by solving as many integer problems as linear sets. However, this is very inefficient, because the number of linear sets may be exponential in the size of the net.

The following observation is the key to a better method:

Proposition 10.

Let $N = (P, T, W)$ be a net, and let M be a deadlocked marking of N . The set $\Sigma = \{p \in P \mid M(p) = 0\}$ is a proper siphon of N . ■ 10

By this proposition, in order to check deadlock freedom it suffices to verify that every proper siphon remains marked at every reachable marking. Moreover, this new property is not too strong: most correct systems satisfy it, because the input transitions of an unmarked siphon cannot occur anymore, and, once again, this is undesirable in all the examples we have examined.

We borrow again a result from [13] :

Theorem 11. *Algebraic characterization of siphons [13]*

Let $N = (P, T, W)$ be a net. Let $N_\Theta = (P, T, W_\Sigma)$, where

$$W_\Sigma(p, t) = \begin{cases} \sum_{p' \in t^\bullet} W(t, p') & \text{if } p \in \bullet t \\ 0 & \text{otherwise} \end{cases}$$

$$W_\Sigma(t, p) = W(t, p)$$

A set $\Sigma \subseteq P$ is a siphon of the net N if and only if the equation $Y^T \cdot \mathbf{N}_\Sigma \leq 0$ has a nonnegative solution Y such that $\|Y\| = \Sigma$. ■ 11

So a marking M of N satisfies the siphon property iff the problem

Variables: Y : rational.

$$Y^T \cdot N_{\Sigma} \leq 0$$

$$Y \geq 0 \text{ } (\Sigma = \|Y\| \text{ is a siphon.})$$

$$Y^T \cdot M = 0 \text{ } (\Sigma \text{ is not marked at } M.)$$

is feasible. Using another version of the Alternatives Theorem and following a procedure similar to the one we used for the trap equation, we obtain that the markings satisfying the siphon property are the solutions of the equation $M > N_{\Sigma} \cdot X$, where $X \leq 0$. Then, the markings which violate the property are those satisfying $M_i \leq (N_{\Sigma})_i \cdot X$, where M_i is the i -th component of M , and $(N_{\Sigma})_i$ the i -th row of N_{Σ} . So we have:

Theorem 12.

Let (N, M_0) be a Petri net. If none of the problems below is feasible, then every reachable marking marks all siphons, and (N, M_0) is deadlock free.

Variables: M, X_1 : integer; X_2 : rational

$$M = M_0 + N \cdot X_1$$

$$M_i \leq (N_{\Sigma})_i \cdot X_2$$

$$M, X_1 \geq 0$$

$$X_2 \leq 0$$

where M_i is the i -th component of M , and $(N_{\Sigma})_i$ the i -th row of N_{Σ} .

■ 12

The number of inequation systems to solve is equal to the number of places of the net. So we have reduced the possibly exponential number of systems to linearly many.

7 Conclusion

We have extended the range of systems that can be verified using linear constraints by adding to the marking equation a new trap equation. The new equation proves to be very useful for the analysis of systems communicating by means of shared variables. We have proved properties of five mutual exclusion algorithms and a telephone communication protocol, none of which could be automatically proved before by linear methods.

We have also given a natural solution to a limitation of the method, namely the fact that deadlock-freedom is not a linear property for arbitrary Petri nets. We have introduced a slightly stronger property, in practice as desirable as deadlock freedom, which can be computed more easily.

References

1. André Arnold. Verification and comparison of transition systems. In M.C. Gaudel and J.P. Jouannaud, editors, *TAPSOFT '93: Theory and Practice of Software Development*, volume 668 of *Lecture Notes in Computer Science*, pages 121–135. Springer-Verlag, 1993.
2. G. S. Avrunin, J. C. Corbett, and U. A. Buy. Integer Programming in the Analysis of Concurrent Systems. In K.G. Larsen and A. Skou, editors, *Computer Aided Verification*, volume 575 of *Lecture Notes in Computer Science*, pages 92–102, 1991.
3. G.S. Avrunin, U.A. Buy, J.C. Corbett, L.K. Dillon, and J.C. Wileden. Automated Analysis of Concurrent Systems with the Constrained Expression Toolset. *IEEE Transactions in Software Engineering*, 17(11):1204–1222, 1991.
4. E. Best and H. Fleischhack (eds.). *Pep: Programming environment based on nets*. Technical report, University of Hildesheim, Germany, 1994.
5. E. Best and R. P. Hopkins. $B(PN)^2$ – A Basic Petri Net Programming Notation. In *Proc. of PARLE-93*, volume 694 of *Lecture Notes in Computer Science*, pages 379–390. Springer-Verlag, 1993
Also: Hildesheimer Informatik Fachbericht 27/92 (1992).
6. G.V. Brams. *Réseaux de Petri: Theorie et Pratique, Vols. I and II*. Masson, 1982.
7. J.C. Corbett. Evaluating Deadlock Detection Methods for Concurrent Software. In T. Ostrand, editor, *Proceedings of the 1994 International Symposium on Software Testing and Analysis*, pages 204–215, New York, 1994.
8. J.C. Corbett and G.S. Avrunin. Using Integer Programming to Verify general Safety and Liveness properties. *Formal Methods in System Design*, 6(1):97–123, 1995.
9. P. Cousot and N. Halbwichs. Automatic discovery of linear restraints among variables of a program. In *5th ACM Symposium on Principles of Programming Languages*. ACM-Press, 1978.
10. CPLEX Optimization Inc. *Using the CPLEXTM Callable Library and CPLEXTM Mixed Integer Library*.
11. J. Desel and J. Esparza. *Free-choice Petri Nets*, volume 40 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1995.
12. J. Esparza and M. Silva. A Polynomial-Time Algorithm to Prove Liveness of Bounded Free Choice Nets. *Theoretical Computer Science*, 102:185–205, 1992.
13. J. Ezpeleta, J. M. Couvreur, and M. Silva. A New Technique for Finding a Generating Family of Siphons, Traps and ST-Components. Application to Colored Petri Nets. In G. Rozenberg, editor, *Advances in Petri Nets*, volume 674 of *Lecture Notes in Computer Science*, pages 126–147. Springer Verlag, 1993.
14. B. Grahlmann. Verifying telecommunication protocols with pep (draft). Technical report, University of Hildesheim, Germany, 1995.
15. N. Halbwichs. About synchronous programming and abstract interpretation. In B. Le Charlier, editor, *SAS '94: Static Analysis Symposium*, volume 864 of *Lecture Notes in Computer Science*, pages 179–192. Springer-Verlag, 1994.
16. Stephan Kleuker. A gentle introduction to specification engineering using a case study in telecommunications. In P.D. Mosses, M. Nielsen, and M.I. Schwartzbach, editors, *TAPSOFT '95*, volume 915 of *Lecture Notes in Computer Science*. Springer-Verlag, 1995.

17. K. Lautenbach. Linear algebraic calculation of deadlocks and traps. In H.J. Genrich K. Voss and G. Rozenberg, editors, *Concurrency and Nets*, pages 315–336. Springer-Verlag, 1987.
18. K. Lautenbach. Linear Algebraic Techniques for Place/Transition Nets. In W. Brauer, W. Reisig, and G. Rozenberg, editors, *Petri Nets: Central Models and Their Properties, Advance in Petri Nets 1986*, volume 254 of *Lecture Notes in Computer Science*, pages 142–167. Springer-Verlag, 1987.
19. G. Memmi and G. Roucairol. Linear Algebra in Net Theory. In W. Brauer, editor, *Net Theory and Applications*, volume 84 of *Lecture Notes in Computer Science*, pages 213–223. Springer-Verlag, 1980.
20. Tadao Murata. Petri nets: Properties, analysis and applications. *Proc. of the IEEE*, 77(4):541–580, 1989.
21. Enric Pastor, Oriol Roig, Jordi Cortadella, and Rosa M. Badia. Petri net analysis using boolean manipulation. In Robert Valette, editor, *Application and Theory of Petri Nets 1994*, volume 815 of *Lecture Notes in Computer Science*, pages 416 – 435. Springer-Verlag, 1994.
22. CCITT Recommendations Q.1200. Intelligent networks, final version. Technical report, 1992.
23. M. Raynal. *Algorithms for Mutual Exclusion*. North Oxford Academic, 1986.
24. W. Reisig. *Petri Nets*, volume 4 of *EATCS Monographs on Theoretical Computer Science*. Springer Verlag, 1985.
25. A. Schrijver. *Theory of Linear and Integer Programing*. Series in Discrete Mathematics. Wiley, 1986.