Advances in Petri Nets 1991 (LNCS 524) pp. 29–47

Proving Temporal Properties of Petri Nets

J. C. Bradfield
Department of Computer Science
University of Edinburgh
The King's Buildings
EDINBURGH
U.K. EH9 3JZ
email: jcb@dcs.ed.ac.uk

ABSTRACT: We present a sound and complete tableau system for proving temporal properties of Petri nets, expressed in a propositional modal mu-calculus which subsumes many other temporal logics. The system separates the checking of fix-points from the rest of the logic, which allows the use of powerful reasoning, perhaps specific to a class of nets or an individual net, to prove liveness and fairness properties. Examples are given to illustrate the use of the system. The proofs of soundness and completeness are given in detail.

KEYWORDS: Petri nets, temporal logic, tableau systems, model-checking.

Contents

1. Introduction

1.1 Topic.

Temporal logics provide a very powerful means of expressing crucial properties of systems, such as liveness and safety properties. Typically, liveness properties (such as 'eventually') can be expressed as least fix-points of some temporal operator, and safety properties (such as 'always') as greatest fix-points; and in logics supporting nested fix-point operators, fairness properties such as 'infinitely often' can be expressed. A long-standing technique for proving temporal logic properties is that of *model-checking*, in which an algorithm is given for determining whether a state of a system (viewed as a model for a temporal logic) satisfies some desired property; for example, using CTL with transition systems ([Clarke *et al.* 83]), or propositional mu-calculus with transition systems or CCS ([Stirling & Walker 89]).

However, such techniques suffer from the 'state explosion problem': the number of states in a system may grow exponentially with the number of concurrent processes in the system, so even algorithms with low complexity quickly become unusable; and infinite state spaces are outwith the scope of the techniques altogether. Hence they are of somewhat limited use for Petri nets (where by Petri net we shall usually mean place/transition nets without capacities).

On the other hand, when dealing with nets there is a range of techniques for showing certain restricted safety and liveness properties; such as invariance analysis (see, e.g. [Lautenbach 86]), and deadlock and trap analysis (see, e.g. [Best 86]). We should therefore like to produce a method for proving temporal properties that allows us to incorporate at least some of the techniques employed on nets. In this paper we present a tableau-based system for proving properties expressed in the propositional mu-calculus. The system extends the ideas of traditional model-checking to infinite systems, but by dealing with sets of states the construction of the state space is avoided, and fix-point induction is used to prove maximal fix-points. Proof of minimal fix-points (which express liveness properties) is separated from the rest of the logic, in order to allow the use of whatever reasoning is necessary. It should be emphasized that the system provides a framework into which more powerful reasoning may be fitted to prove minimal fix-points—there is not, and cannot be, a general effective technique for liveness proofs, so the reasoning required may be specific to a class of nets, or to the single net under consideration.

1.2 Structure of this paper.

In section 2, we describe briefly a restricted propositional mu-calculus, and give a very simple example of a net and a property to be proved. In section 3, we describe the tableau system, and apply it to the example of section 2. In sections 4 and 5, we give more complex examples, and in section 6 we prove the soundness and completeness of the system.

2. Propositional Mu-Calculus

2.1 Syntax and semantics.

The logic we use is the propositional modal mu-calculus restricted to a single (empty) label and positive normal form. The modal mu-calculus was introduced by [Pratt 81] and [Kozen 83] as an extension of dynamic logics such as PDL. [Emerson & Lei 86] employed it as a temporal logic—it subsumes several standard temporal logics such as CTL*. In the slightly restricted form we use, the syntax for formulae Φ is:

$$\Phi \ ::= \ Z \mid \Phi_1 \wedge \Phi_2 \mid \Phi_1 \vee \Phi_2 \mid [\,]\Phi \mid \langle\,\rangle\Phi \mid \nu Z.\Phi \mid \mu Z.\Phi$$

where Z ranges over a (countably infinite) set Var of variable symbols. We shall, when convenient, allow negated variables $\neg Z$, subject to the condition that no negated variable may be bound by μ or ν (this restriction ensures that the fix-points have solutions, by making the bodies monotonic in the variables). We shall use σ to stand for either μ or ν. The precedence of operators is that $[\,]$ and $\langle\,\rangle$ bind most tightly, then $\wedge$ and $\vee$, and $\mu Z.$ and $\nu Z.$ bind least tightly so that their scope extends as far to the right as possible. Parentheses are used as usual.

Note that $\sigma Z.$ binds the variable Z; by renaming bound variables we can ensure that distinct variables have distinct names, and henceforth we assume that formulae are in this normal form. Note also that the absence of negation is not a restriction, since we have the dual operators and negated variables; we could equivalently omit $\langle\,\rangle$, $\vee$ and $\mu Z.$, add negation and say that bound variables must be within the scope of an even number of negations.

This logic may be interpreted on the markings of a place/transition net $\mathcal{N} = (\mathbf{S}, \mathbf{T}; \mathbf{F})$; we shall write $\mathcal{M}$ for the set of all markings (taking nets without initial markings—if nets have initial markings, $\mathcal{M}$ should be the set of reachable markings, but for our purposes it is cleaner not to consider initial markings as part of the net, especially since we do not wish to do reachability analysis unless we need to), and for markings M and M' we write $M \xrightarrow{t} M'$, where $t \in \mathbf{T}$ is a transition, to mean t fires at M to give M' (we do not use $M[t\rangle M'$, in order to avoid confusion with the modal operators), and $M \longrightarrow M'$ if there exists t such that $M \xrightarrow{t} M'$. This assumes an interleaving semantics, for the sake of simplicity in the examples, but we could just as well use a concurrent semantics with sets T of transitions.

The meanings of the Boolean connectives are as usual; the box operator is 'for all immediately reachable markings Φ holds' and the diamond is 'there exists an immediately reachable marking at which Φ holds'; the fix-point formulae mean the fix-points of Φ regarded as a function via the free variable Z. The semantics is given formally as follows.

A *model* for this logic is a net $\mathcal{N}$ together with a *valuation* $\mathcal{V}\colon \mathrm{Var} \to 2^{\mathcal{M}}$ assigning a set of markings to each variable symbol, and in a model $(\mathcal{N}, \mathcal{V})$ we assign to a formula Φ its *denotation* $\|\Phi\|_{\mathcal{V}}^{\mathcal{N}}$ (or just $\|\Phi\|_{\mathcal{V}}$ when $\mathcal{N}$ can be understood), the set of markings at which it holds, thus:

$$\begin{aligned}
\|Z\|^{\mathcal{N}}_{\mathcal{V}} &= \mathcal{V}(Z) \qquad (Z \in \mathrm{Var}) \\
\|\Phi_1 \wedge \Phi_2\|^{\mathcal{N}}_{\mathcal{V}} &= \|\Phi_1\|^{\mathcal{N}}_{\mathcal{V}} \cap \|\Phi_2\|^{\mathcal{N}}_{\mathcal{V}} \\
\|\Phi_1 \vee \Phi_2\|^{\mathcal{N}}_{\mathcal{V}} &= \|\Phi_1\|^{\mathcal{N}}_{\mathcal{V}} \cup \|\Phi_2\|^{\mathcal{N}}_{\mathcal{V}} \\
\|[\,]\Phi\|^{\mathcal{N}}_{\mathcal{V}} &= \{\, M \in \mathcal{M} \mid \forall M' .\, M \longrightarrow M' \;\Rightarrow\; M' \in \|\Phi\|^{\mathcal{N}}_{\mathcal{V}} \,\} \\
\|\langle\,\rangle\Phi\|^{\mathcal{N}}_{\mathcal{V}} &= \{\, M \in \mathcal{M} \mid \exists M' .\, M \longrightarrow M' \;\wedge\; M' \in \|\Phi\|^{\mathcal{N}}_{\mathcal{V}} \,\} \\
\|\nu Z.\Phi\|^{\mathcal{N}}_{\mathcal{V}} &= \bigcup \{\, \mathbf{M} \subseteq \mathcal{M} \mid \|\Phi\|^{\mathcal{N}}_{\mathcal{V}[Z:=\mathbf{M}]} \supseteq \mathbf{M} \,\} \\
\|\mu Z.\Phi\|^{\mathcal{N}}_{\mathcal{V}} &= \bigcap \{\, \mathbf{M} \subseteq \mathcal{M} \mid \|\Phi\|^{\mathcal{N}}_{\mathcal{V}[Z:=\mathbf{M}]} \subseteq \mathbf{M} \,\}
\end{aligned}$$

where $\mathcal{V}[Z := \mathbf{M}]$ is the valuation which assigns $\mathbf{M}$ to Z and otherwise agrees with $\mathcal{V}$. The boolean and modal clauses here reflect the intended meanings directly, and the fix-point clauses are correct by virtue of the Tarski fix-point theorem (our logic has the property that all formulae are monotonic in all variables).

We shall define a special subclass of variable symbols, sometimes referred to as atomic propositions, which have the form of linear (in)equalities on letters $\mathrm{A}, \mathrm{B}, \ldots$. Such a symbol is intended to be used with a model $(\mathcal{N}, \mathcal{V})$ in which the places of the net are labelled $\mathrm{A}, \mathrm{B}, \ldots$, and the valuation of the symbol is those markings satisfying the inequality. For example, if $\mathrm{B}, \mathrm{C} \in \mathbf{S}$, then $\mathcal{V}(\mathrm{B}+2\mathrm{C} > 1) = \{\, M \in \mathcal{M} \mid M_{\mathrm{B}} + 2M_{\mathrm{C}} > 1 \,\}$. Such a model is called an intended model for the symbol, and henceforth all models are assumed to be intended for any formula interpreted in them.

Some examples of mu-formulae and their informal meanings are (where tt and ff abbreviate $0 = 0$ and $0 = 1$):

$\mu Z.\Phi \vee ([\,]Z \wedge \langle\,\rangle \mathrm{tt})$	'eventually Φ'
$\nu Z.\Phi \wedge [\,]Z$	'always Φ'
$\mu Z.\Psi \vee (\Phi \wedge \langle\,\rangle Z)$	'perhaps Φ will hold until Ψ does'
$\nu Y.\mu Z.((\Phi \wedge [\,]Y) \vee [\,]Z) \wedge \langle\,\rangle \mathrm{tt}$	'infinitely often Φ'

For those unfamiliar with the mu-calculus, some explanation of these may be in order.

(i) $\mu Z.\Phi \vee ([\,]Z \wedge \langle\,\rangle \mathrm{tt}) \stackrel{\mathrm{def}}{=} \mu Z.\Psi$
Suppose a marking M satisfies $\mu Z.\Psi$. Then, since $\mu Z.\Psi$ is a fix-point of Ψ, M satisfies $\Phi \vee ([\,](\mu Z.\Psi) \wedge \langle\,\rangle \mathrm{tt})$, so either M satisfies Φ, or there is at least one successor marking ($\langle\,\rangle \mathrm{tt}$) and all successor markings satisfy $\mu Z.\Psi$. Since $\mu Z.\Psi$ is a minimal fix-point, we cannot have an unending sequence of markings not satisfying Φ, and so the formula is satisfied by markings at which Ψ is eventually true.

(ii) $\nu Z.\Phi \wedge [\,]Z \stackrel{\mathrm{def}}{=} \nu Z.\Psi$
Here we have that both M satisfies Ψ, and all successors satisfy $\nu Z.\Psi$, hence Φ always holds.

(i) and (ii) illustrate the general point that maximal fix-points represent safety properties, and minimal fix-points represent liveness properties. (It may be better to say that safety is represented by the idea of post-fix-points; then since the maximal fix-point is also the maximal post-fix-point, it is the biggest set satisfying the safety property.)

(iii) $\mu Z.\Psi \vee (\Phi \wedge \langle\,\rangle Z) \stackrel{\mathrm{def}}{=} \mu Z.\Upsilon$
Either Ψ holds, or Φ holds, and there is some successor at which $\mu Z.\Upsilon$ holds. Thus, there exists some run on which Φ holds until Ψ does (and Ψ eventually does hold, since this is a minimal fix-point—$\nu Z.\Upsilon$ would also be satisfied by the existence of a path on which Φ holds for ever, giving 'it may be that Φ holds as long as $\neg\Psi$').

(iv) Formulae with nested mixed fix-points are more difficult to understand intuitively.

$$\nu Y.\mu Z.((\Phi \wedge [\,]Y) \vee [\,]Z) \wedge \langle\,\rangle \mathrm{tt} \stackrel{\mathrm{def}}{=} \nu Y.\mu Z.\Psi$$

Here, either Φ holds now, in which case $\nu Y.\mu Z.\Psi$ holds at all successors, or $\mu Z.\Psi$ holds at all successors. Since this latter is a minimal fix-point, only finitely many steps may occur before $(\Phi \wedge [\,]Y)$ is true—so the formula means 'Φ holds eventually, and when it holds, it holds again', which is to say, 'Φ holds infinitely often'.

As can be seen from these examples, the mu-calculus has few primitives but great expressive power; since it has so few primitives, the proof system is relatively easy to understand, and in practice one would have a macro language for common complex properties such as those above.

The formulae above illustrate the inherent expressive power of the modal mu-calculus (that is, the properties expressible without considering the structure of the model or the power of atomic propositions). This inherent power is great (closely related to that of Rabin's SnS, the second order theory of n successors [Hüttel 90]), but limited. As a gross simplification, the modal mu-calculus 'knows about' looping and termination (and mixes of these such as fairness properties), but it cannot count—for example, the property 'a fires as often as b' is not generally expressible. However, in a given model (net), it may be possible for mu-calculus formulae to express complex (even non-arithmetic) sets of markings; so by adding new places A and B with $^\bullet\mathrm{A} = \mathrm{a}$ and $^\bullet\mathrm{B} = \mathrm{b}$ one can count firings of a and b, and if one makes A and B empty at the initial markings of interest, 'a as often as b' becomes just $\mu Z.(\mathrm{A} = \mathrm{B} \wedge [\,]\mathrm{ff}) \vee [\,]Z$.

The standard formulation of the propositional mu-calculus works with labelled transition systems $(\mathcal{M}, L, \{\, \stackrel{l}{\longrightarrow} \mid l \in L\,\})$ where L is a set of *labels* and each $\stackrel{l}{\longrightarrow}$ is a transition relation; then relativize box and diamond formulae to

$$\|[l]\Phi\|^{\mathcal{N}}_{\mathcal{V}} = \{\, M \in \mathcal{M} \mid \forall M' \,.\, M \stackrel{l}{\longrightarrow} M' \;\Rightarrow\; M' \in \|\Phi\|^{\mathcal{N}}_{\mathcal{V}} \,\}$$

and so on, so that we can, for example, express the firing of a particular event in a net. The extension of the tableau system to deal with this is entirely straightforward; we show an example in section 5, but do not discuss it further. (For a generalized tableau system for the full mu-calculus on transition systems, see [Bradfield & Stirling 91].)

2.2 An example.

Consider the net of figure 1. The elements A and a add tokens to B until prevented by the occurrence of b; then c may fire a number of times bounded by B. So one property that we should be able to prove is that c is enabled (that is, $(\mathrm{B} > 0) \wedge (\mathrm{C} > 0) \stackrel{\mathrm{def}}{=} \Psi$) only finitely often. The property may be expressed in the mu-calculus as

$$\mu Y.\nu Z.((\neg\Psi \vee \langle\,\rangle Y) \wedge \langle\,\rangle Z) \vee [\,]\mathrm{ff}$$

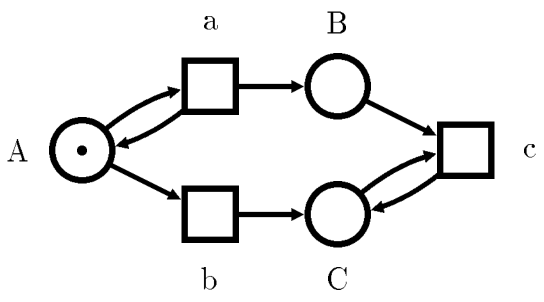

Figure 1

3. The Tableau System

Assume that the model $(\mathcal{N}, \mathcal{V})$ is fixed, and write $\mathbf{M} \models_{\mathcal{V}} \Phi$ for $\mathbf{M} \subseteq \|\Phi\|^{\mathcal{N}}_{\mathcal{V}}$: we wish to show that an initial marking, or more generally a set $\mathbf{M}$ of markings, satisfies some property Ψ. We present a tableau system for proving this, built on sequents of the form $\mathbf{M} \vdash \Phi$; a sequent represents a goal of showing $\mathbf{M} \models_{\mathcal{V}} \Phi$, and each rule of the system takes a rule and produces one or two subgoals, until goals are reached which have a variable, in which case the goal is a leaf of the tableau. Once built, we give criteria for determining whether a tableau is successful, that is, gives a valid proof of its root sequent.

3.1 Building the tableau.

A *tableau* is a tree of sequents built according to the following rules (presented as inverted proof rules) from a root $\mathbf{M} \vdash \Psi$.

(And) $$\frac{\mathbf{M} \vdash \Phi_1 \wedge \Phi_2}{\mathbf{M} \vdash \Phi_1 \quad \mathbf{M} \vdash \Phi_2}$$

(Or) $$\frac{\mathbf{M} \vdash \Phi_1 \vee \Phi_2}{\mathbf{M}_1 \vdash \Phi_1 \quad \mathbf{M}_2 \vdash \Phi_2} \qquad \mathbf{M} = \mathbf{M}_1 \cup \mathbf{M}_2$$

(Box) $$\frac{\mathbf{M} \vdash [\,]\Phi}{\mathbf{M}' \vdash \Phi} \qquad \mathbf{M}' = \{\, M' \mid \exists M \in \mathbf{M} \,.\, M \longrightarrow M' \,\}$$

(Diamond) $$\frac{\mathbf{M} \vdash \langle\,\rangle\Phi}{\mathbf{M}' \vdash \Phi} \qquad \mathbf{M}' = \{\, M' \mid \exists M \in \mathbf{M} \,.\, M \xrightarrow{f(M)} M' \,\}$$
where $f\colon \mathbf{M} \to \mathbf{T}$ is a function

(Unfolding) $$\frac{\mathbf{M} \vdash \sigma Z.\Phi}{\mathbf{M} \vdash \Phi}$$

(Weakening) $$\frac{\mathbf{M} \vdash \Phi}{\mathbf{M}' \vdash \Phi} \qquad \mathbf{M}' \supseteq \mathbf{M}$$

Most of these rules are straightforward: (Box) takes all successors of markings in the premise; (Diamond) for each marking in the premise chooses a transition to fire (this assumes an interleaving firing strategy, which keeps the size of the examples down, but we could just as well select sets or multisets of transitions)—it is important to note that the function f is an integral part of the rule application: the rule demands specific successors, rather than just the existence of successors; (Unfolding) unfolds a fix-point formula; and finally (Weakening) allows us to enlarge the set being checked. (Weakening) is essential for the completeness of the system; however, it need only be applied to fix-point sequents.

3.2 Successful tableaux.

The above rules build the tableau: we now define when a tableau succeeds.

A node in a tableau is called *terminal* if no rule other than (Weakening) applies to it. If $\mathbf{M} \vdash Z$ is a terminal node, and Z is a fix-point variable, its *companion* is the (lowest) node $\mathbf{M}' \vdash \sigma Z.\Phi$ above it.

A terminal node $\mathbf{n} = \mathbf{M} \vdash \Phi$ is *successful* iff

(i) $\mathbf{M} = \emptyset$, or

(ii) $\Phi = Z \in \mathrm{Var}$ with no companion, and $\mathbf{M} \models_{\mathcal{V}} \Phi$, or

(iii) $\Phi = Z \in \mathrm{Var}$ with companion $\mathbf{n}' = \mathbf{M}' \vdash \nu Z.\Psi$ where $\mathbf{M}' \supseteq \mathbf{M}$, or

(iv) $\Phi = Z \in \mathrm{Var}$ with companion $\mathbf{n}' = \mathbf{M}' \vdash \mu Z.\Psi$ where $\mathbf{M}' \supseteq \mathbf{M}$, and $\mathbf{n}'$ satisfies the condition (mu-success) given below

and a tableau is successful iff every terminal node is successful.

(i)–(ii) are clear; (iii) encapsulates the principle of fix-point induction, that $\mathbf{M} \subseteq \|\Phi\|_{\mathcal{V}[Z:=\mathbf{M}]} \Rightarrow \mathbf{M} \subseteq \|\nu Z.\Phi\|_{\mathcal{V}}$; and (iv) shunts the burden of showing a least fix-point off to a test (mu-success) on a relation defined by the tableau.

A notion we shall use later is that of *valuation-independent success* (v-i success), which is defined as for success except that all free-variable terminals are v-i successful.

To define the condition (mu-success), we define notions which formalize the intuitive idea of the truth of a fix-point formula for a marking being dependent on the truth of the formula at another marking; for minimal fix-points, all such chains of dependency must terminate.

In a tableau, a *path* from a marking M at a node $\mathbf{n}$ to a marking M' at a node $\mathbf{n}'$ is a sequence of markings and nodes $(M, \mathbf{n}) = (M_0, \mathbf{n}_0), (M_1, \mathbf{n}_1), \ldots, (M_k, \mathbf{n}_k) = (M', \mathbf{n}')$ such that each $\mathbf{n}_{i+1}$ is a child of $\mathbf{n}_i$, and each M_i is in the set of the sequent of $\mathbf{n}_i$, and $M_i = M_{i+1}$ except where the rule applied to $\mathbf{n}_i$ is (Box), when $M_i \longrightarrow M_{i+1}$, or (Diamond), when $M \xrightarrow{f(M)} M_{i+1}$; that is, a path is the path through the tableau found by following a particular marking through the application of rules. We write $M@\mathbf{n} \dashrightarrow M'@\mathbf{n}'$ if there exists a path from M at $\mathbf{n}$ to M' at $\mathbf{n}'$.

Now we say that there is an *extended path* from M at $\mathbf{n}$ to M' at $\mathbf{n}'$, written $M@\mathbf{n} \dashrightarrow\!\!\!\!\rightarrow M'@\mathbf{n}'$, if either

(i) $M@\mathbf{n} \dashrightarrow M'@\mathbf{n}'$, or

(ii) there exists a node $\mathbf{n}'' = \mathbf{M}'' \vdash \sigma Z.\Phi$ and a finite sequence of markings $M_0, \ldots, M_n$ $(n \geq 0)$ such that $M@\mathbf{n} \dashrightarrow M_0@\mathbf{n}''$, and for each i, $M_i@\mathbf{n}'' \dashrightarrow\!\!\!\!\rightarrow M_{i+1}@\mathbf{n}_{i+1}$

where $\mathbf{n}_{i+1} = \mathbf{M}_{i+1} \vdash Z$ for some M_{i+1}, and $M_n@\mathbf{n}'' \dashrightarrow M'@\mathbf{n}'$.
(Note that this definition involves an implicit induction on the height of σ-nodes.) This notion is illustrated graphically in the example in the next subsection.

Then if $\mathbf{n}$ is labelled $\sigma Z.\Phi$ we say $M \sqsupset_{\mathbf{n}} M'$ iff $M@\mathbf{n} \dashrightarrow M'@\mathbf{n}'$ for some $\mathbf{n}'$ labelled by Z, and $\mathbf{n}$ satisfies the condition (mu-success) iff the relation $\sqsubset_{\mathbf{n}}$ is well-founded.

Theorem. *The tableau system as given above is sound and complete, that is,* $\mathbf{M} \vdash \Phi$ *is the root of a successful tableau iff* $\mathbf{M} \models_{\mathcal{V}} \Phi$.

Proof. See section 6. □

3.3 An example of the tableau system.

We give the tableau for the example of section 2.

First, we must have a method of representing sets of markings. Clearly we cannot represent arbitrary sets; instead, if Φ is a boolean combination of atomic propositions, we let $\{\,\Phi\,\}$ denote $\|\Phi\|_{\mathcal{V}}^{\mathcal{N}}$: this method suffices to represent all sets in our examples. (The completeness of the system relative to various presentations of sets is an interesting question for further investigation.)

Let Φ be the equality $\mathrm{A} + \mathrm{C} = 1$, and recall that Ψ is $\mathrm{B} > 0 \;\wedge\; \mathrm{C} > 0$. The tableau is given in figure 2.

$$_1\{\,\mathrm{A}=1\wedge\mathrm{B}=0\wedge\mathrm{C}=0\,\}\vdash \mu Y.\nu Z.((\neg\Psi\vee\langle\rangle Y)\wedge\langle\rangle Z)\vee[\,]\mathrm{ff}$$

$$_2\{\,\Phi\,\}\vdash \mu Y.\nu Z.((\neg\Psi\vee\langle\rangle Y)\wedge\langle\rangle Z)\vee[\,]\mathrm{ff}$$

$$_3\{\,\Phi\,\}\vdash \nu Z.((\neg\Psi\vee\langle\rangle Y)\wedge\langle\rangle Z)\vee[\,]\mathrm{ff}$$

$$_4\{\,\Phi\,\}\vdash ((\neg\Psi\vee\langle\rangle Y)\wedge\langle\rangle Z)\vee[\,]\mathrm{ff}$$

$$_5\{\,\Phi\wedge(\Psi\vee\mathrm{A}>0)\,\}\vdash (\neg\Psi\vee\langle\rangle Y)\wedge\langle\rangle Z \qquad \mathcal{T}$$

$$_6\{\,\Phi\wedge(\Psi\vee\mathrm{A}>0)\,\}\vdash \neg\Psi\vee\langle\rangle Y \qquad {}_{10}\{\,\Phi\wedge(\Psi\vee\mathrm{A}>0)\,\}\vdash \langle\rangle Z$$

$$_7\{\,\Phi\wedge(\neg\Psi\wedge\mathrm{A}>0)\,\}\vdash \neg\Psi \quad {}_8\{\,\Phi\wedge(\Psi\wedge\mathrm{A}=0)\,\}\vdash \langle\rangle Y \quad {}_{11}\{\,\Phi\wedge\mathrm{A}=0\,\}\vdash Z$$

$$_9\{\,\Phi\wedge\mathrm{A}=0\wedge\mathrm{C}=1\,\}\vdash Y$$

$$\text{where } \mathcal{T} = \quad \frac{_{12}\{\,\Phi\;\wedge\;\mathrm{A}=0\;\wedge\;\neg\Psi\,\}\vdash[\,]\mathrm{ff}}{_{13}\{\;\}\vdash\mathrm{ff}}$$

Figure 2

Some rules require a little explanation: at 1, we apply (Weakening), and expand the set from the initial marking to the set $\{\,\Phi\,\}$; Φ (as is easily seen) is an invariant of the system, and so will be carried through to all sequents. At 4, we have a disjunction, so in 5 we put precisely those markings satisfying the left disjunct, and put the rest in 12. Similarly at 6. At 8 and 10, we have (Diamond), so we must specify the functions f_8

and f_{10} choosing a transition to fire. These functions are the same, given by $f(M) = \mathrm{b}$ if $M_\mathrm{A} = 1$, and $f(M) = \mathrm{c}$ if $M_\mathrm{A} = 0$.

It is now easy to see that the tableau is successful. Firstly, the fix-point terminals 9 and 11 satisfy the inclusion condition. Now we need to consider the relation $\sqsupset_2$. By definition, $M \sqsupset_2 M'$ if there is an extended path from M at 2 to M' at a Y terminal, and the only Y terminal is 9. So there is such an extended path if either we can find a path from 2 to 9 (along the solid line in the figure), or we can follow the dashed path down to 11, jump back up to the inner fix-point 3 and repeat a few times, and then continue down the solid path to 9. So we see that any extended path from 2 to 9 passes through 8, perhaps first passing several times through 10, so $M \sqsupset_2 M'$ iff $M' = f_8(f_{10}^k(M))$ for some $k \geq 0$; and since application of f_8 or f_{10} results in a strict decrease in the pair $(M_\mathrm{A}, M_\mathrm{B})$ (ordered lexicographically), we have that $\sqsubset_2$ is well-founded.

This tableau demonstrates how the system has both mechanical and non-mechanical aspects: application of the (Or) and (Diamond) rules requires making a choice, and while a correct choice always exists, there is no effective rule to determine it; and even when the correct choice is made, well-foundedness of $\sqsubset$ must be shown outside the system.

Note, incidentally, that properties of this kind are already more than can be shown by use of coverability graphs; if we add an arc from c to B, the property fails, but the coverability graph is unchanged. However, limited reachability analyses such as coverability graphs can be of value in building a tableau, if not in showing the success of the tableau, by suggesting values for the sets of markings to label sequents.

4. A Slot Machine

For a more complex example we have a slot machine and a foolish gambler. Consider the Petri net of figure 3. This represents a simple slot machine which accepts a coin, and then pays out a random amount between zero and the amount in its bank ; and a gambler who has an inexhaustible supply of credit—we shall assume that the creditor, though infinitely rich, is irrational enough to demand his money back eventually. Since the machine starts with an empty bank, the gambler is always in debt; we should like to encourage her to continue to play by proving to her that it is always possible for her to recover her losses, that is, it is always possible that $\mathrm{C} - \mathrm{W} = 0$, which in the mu-calculus is

$$\nu Y.(\mu Z.(\mathrm{C} - \mathrm{W} = 0) \vee \langle\rangle Z) \wedge [\,]Y$$

As always when applying this system, a crucial step is choosing the right set of states for the root sequent of the tableau, and a good heuristic is to choose the biggest sensible set of states, so let us start with the following:

$$\{\, \Phi \stackrel{\mathrm{def}}{=} (\mathrm{M} = 1) \;\wedge\; (\mathrm{P} + \mathrm{R} = 1) \wedge (\mathrm{C} = \mathrm{S} + \mathrm{B} + \mathrm{W}) \,\}$$

The tableau for this formula and set of markings is set out in figure 4.

Notes on the rules:

4 (Or) At a disjunction, a heuristic for producing successful tableaux is to put as many states as possible into the branch with no variables, and the remainder into

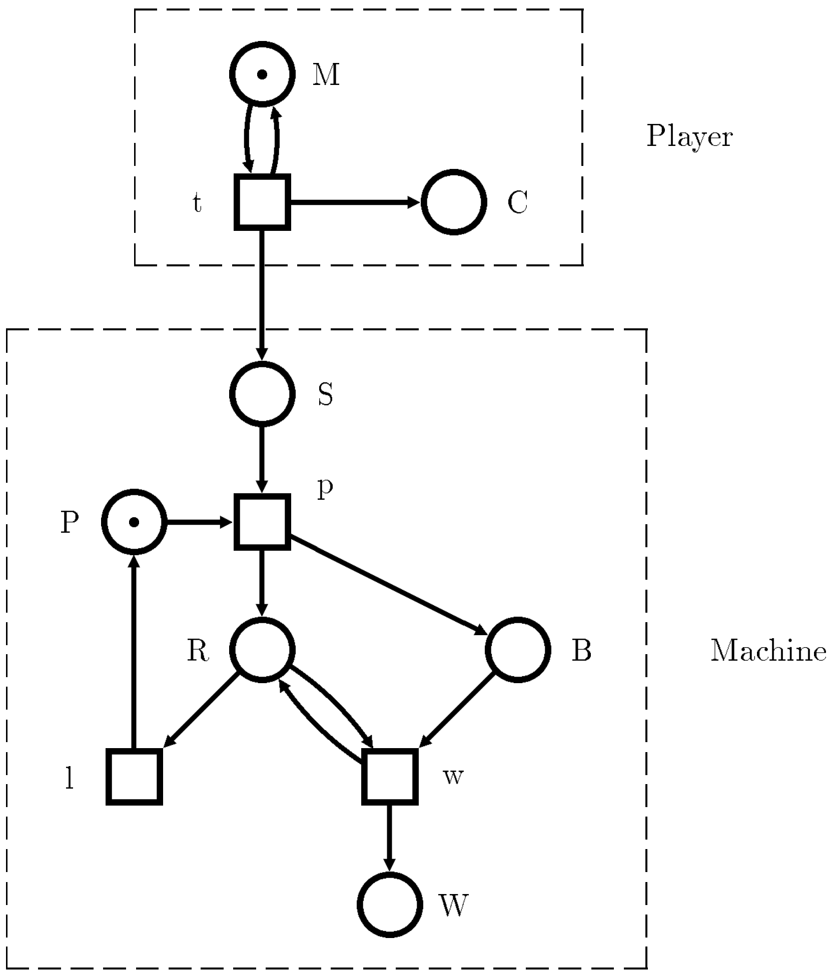

Legend

M	the player's source of credit
t	the player puts a coin in the slot
C	how many coins the player has spent
S	the slot (with unbounded storage capacity)
p	the PLAY button
P	the machine is ready to play
B	the machine's bank
R	the reels are spinning
W	the player's winnings
l	lose: the machine stops paying out and is ready to play again
w	win: the machine pays out a coin, and keeps spinning the reels

Figure 3

$$
\dfrac{_{(1)}\,\{\,\Phi\,\} \vdash \nu Y.(\mu Z.(\mathrm{C}-\mathrm{W}=0) \vee \langle\,\rangle Z) \wedge [\,]Y}{\dfrac{_{(2)}\,\{\,\Phi\,\} \vdash (\mu Z.(\mathrm{C}-\mathrm{W}=0) \vee \langle\,\rangle Z) \wedge [\,]Y}{\dfrac{_{(3)}\,\{\,\Phi\,\} \vdash \mu Z.(\mathrm{C}-\mathrm{W}=0) \vee \langle\,\rangle Z}{\dfrac{_{(4)}\,\{\,\Phi\,\} \vdash (\mathrm{C}-\mathrm{W}=0) \vee \langle\,\rangle Z}{_{(5)}\,\{\,\Phi \,\wedge\, (\mathrm{S}=0) \,\wedge\, (\mathrm{B}=0)\,\} \vdash (\mathrm{C}=\mathrm{W}) \qquad \dfrac{_{(6)}\,\{\,\Phi \,\wedge\, (\mathrm{S}+\mathrm{B}>0)\,\} \vdash \langle\,\rangle Z}{\begin{array}{ll} \{ & (\Phi \,\wedge\, (\mathrm{S}+(\mathrm{B}+1)>0) \,\wedge\, (\mathrm{B}\geq 0)) \\ & \vee\ (\Phi \,\wedge\, (\mathrm{S}+\mathrm{B}>0) \,\wedge\, (\mathrm{B}=0)) \\ & \vee\ (\Phi \,\wedge\, ((\mathrm{S}+1)+(\mathrm{B}-1)>0) \\ & \qquad \wedge\, (\mathrm{B}-1=0) \,\wedge\, (\mathrm{S}\geq 0)) \\ & \vee\ (\Phi \,\wedge\, ((\mathrm{S}-1)+\mathrm{B}>0)) \\ _{(7)} & \} \vdash Z \end{array}}}}} \qquad \dfrac{_{(8)}\,\{\,\Phi\,\} \vdash [\,]Y}{_{(9)}\,\{\,\Phi''\,\} \vdash Y}}}
$$

Figure 4

the other branch; so we put in the left branch precisely those markings that satisfy C = W.

6 (Diamond) To apply this rule, for each marking in the premise, we must select a transition to fire. The heuristic is always to choose the transition that gives a new marking 'nearest' the goal C = W (and indeed, when 'nearest' and 'goal' are made precise, this gives an algorithm that yields the completeness of the system). To do this, we partition the markings of the premise into the following four disjoint subsets, where $\Phi' \stackrel{\text{def}}{=} \Phi \,\wedge\, (\mathrm{S}+\mathrm{B})>0$:

(i) $\{\,\Phi' \,\wedge\, (\mathrm{R}=1) \,\wedge\, (\mathrm{B}>0)\,\}$

(ii) $\{\,\Phi' \,\wedge\, (\mathrm{R}=1) \,\wedge\, (\mathrm{B}=0)\,\}$

(iii) $\{\,\Phi' \,\wedge\, (\mathrm{P}=1) \,\wedge\, (\mathrm{S}>0)\,\}$

(iv) $\{\,\Phi' \,\wedge\, (\mathrm{P}=1) \,\wedge\, (\mathrm{S}=0)\,\}$

and for all states in class (i) we fire w, in (ii) l, in (iii) p, and in (iv) t; this gives the four disjuncts in the conclusion.

7 This is a successful mu-terminal: for every disjunct implies Φ, so the node satisfies the inclusion condition; and $\sqsubset_3$ is well-founded, for if $M@3 \dashrightarrow M'@7$, the path must be $M@3$, $M@4$, $M@6$, $M'@7$, so we need only consider M satisfying Φ'. Then we observe

(a) if M is in class (i), then either $\mathrm{B}>1$, in which case M' is in class (i) with strictly smaller B, or $\mathrm{B}=1$, in which case M' is in class (ii) ; and in either case, S is unchanged;

(b) if M is in class (ii), then M' is in class (iii) (because $\mathrm{B}=0 \Rightarrow \mathrm{S}>0$ since $\mathrm{S}+\mathrm{B}>0$); and S is unchanged;

(c) if M is in class (iii), then M' is in class (i) and S decreases by one;

(d) if M is in class (iv), then either M' has $\mathrm{S} = \mathrm{B} = 0$, in which case $M' \not\sqsupset_3$, or M' is in class (iii).

If we now impose an order $\prec$ on markings by the lexicographic ordering on the pair (S, B), we see from the above that if $M \sqsupset_3 M' \sqsupset_3 M''$ then $M'' \prec M$; and since $\prec$ is well-founded, so is $\sqsubset_3$.

8 (Box) Here we must consider all possible firings of transitions for each marking. But since all we want is that 9 should satisfy the inclusion condition, it suffices to note that Φ remains true under any firing, and so

9 is a successful ν-terminal, since Φ'', whatever it be, implies Φ.

5. MRSW Interlock

Our final example serves two purposes: it illustrates the use of mu-calculus with labels: the step $M \xrightarrow{t} M'$ is labelled by t, and we have relativized modalities $[t]$ as well as the general modality $[\,]$; and it also gives a very simple example of using the system to reason about a class of nets. Specifically, we present a net which implements multiple-read–single-write interlock for m readers and n writers, and give a tableau parametrized by m and n for a liveness property.

The net is shown in figure 5. Proving the safety of the system is left as an exercise for the reader; we shall prove a liveness property, namely that a writer process that wants to write eventually will. The net ensures this by arranging that a writer that wishes to write removes the token from place P, which prevents any new read actions (note that actually we are only ensuring that readers cannot indefinitely delay writers—there is no fairness amongst readers, or amongst writers: these could be added, but would make the example bigger). The property may be informally expressed as:

$$\mathrm{always}([\mathrm{u}_{j_0}](\mathrm{eventually}(\mathrm{W}_{j_0} = 1)))$$

for a particular writer W_{j_0}, and formalized as

$$\nu Y.[\,]Y \wedge [\mathrm{u}_{j_0}](\mu Z.(\mathrm{W}_{j_0} = 1) \vee [\,]Z)$$

The main invariant for this system is

$$\begin{aligned}\Phi \stackrel{\mathrm{def}}{=} & \bigwedge_j (\mathrm{U}_j + \mathrm{V}_j + \mathrm{W}_j = 1) \\ \wedge & \bigwedge_i (\mathrm{R}_i + \mathrm{S}_i = 1) \\ \wedge & \ (\mathrm{Q} + m \sum_j \mathrm{W}_j + \sum_i \mathrm{R}_i = m) \\ \wedge & \ (\mathrm{P} + \sum_j \mathrm{U}_j = 1)\end{aligned}$$

and a tableau is formed thus, where $\Psi \stackrel{\mathrm{def}}{=} (\mathrm{U}_{j_0} = 1) \wedge (\mathrm{P} = 0)$:

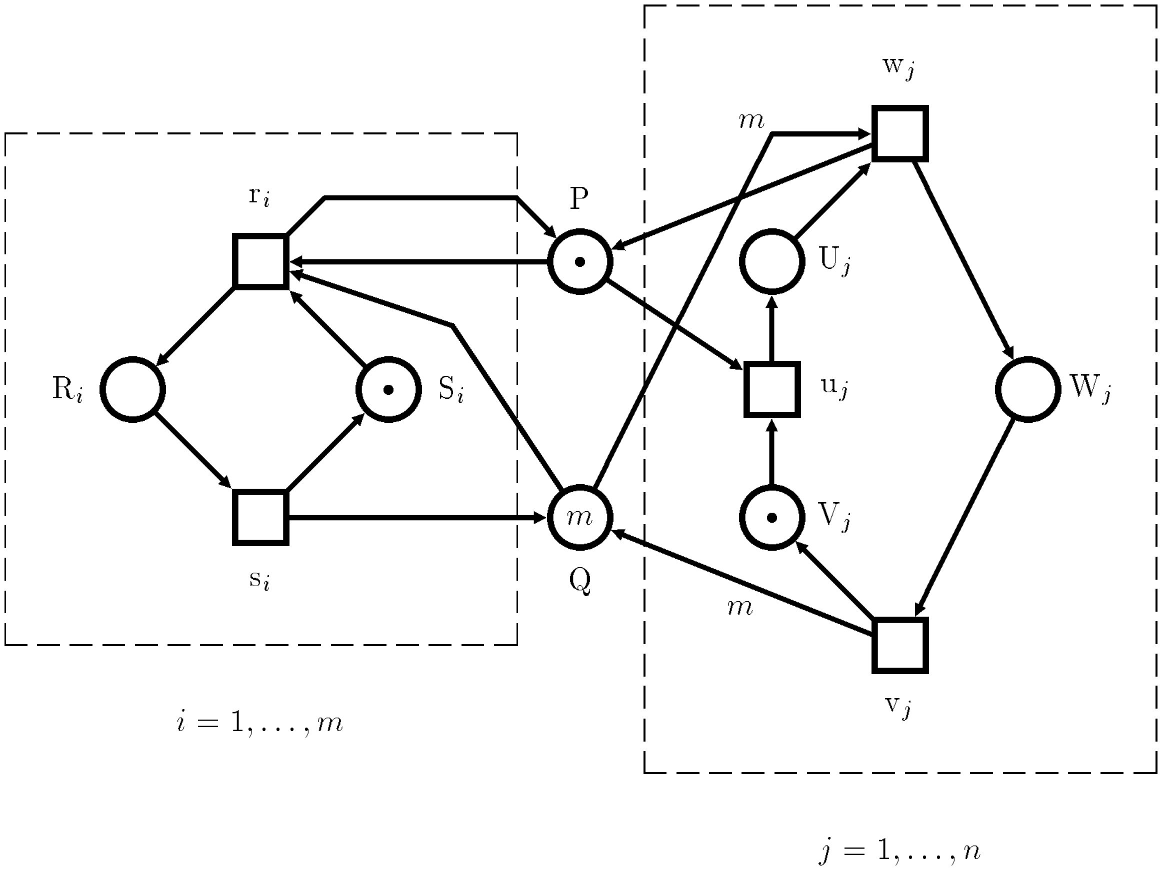

Legend

For m reader processes		and n writer processes	
S_i	process not reading	V_j	process not writing
r_i	process starts reading	w_j	process starts writing
R_i	process reading	W_j	process writing
s_i	process stops reading	v_j	process stops writing
		u_j	process requests write access
		U_j	process waiting to write
Q	the resource		
P	an interlock		

Figure 5

$$\frac{{}_1\{\,\Phi\,\} \vdash \nu Y.[\,]Y \wedge [\mathrm{u}_{j_0}](\mu Z.(\mathrm{W}_{j_0} = 1) \vee [\,]Z)}{{}_2\{\,\Phi\,\} \vdash [\,]Y \wedge [\mathrm{u}_{j_0}](\mu Z.(\mathrm{W}_{j_0} = 1) \vee [\,]Z)}$$

$$\frac{{}_3\{\,\Phi\,\} \vdash [\,]Y}{{}_4\{\,\Phi \wedge \Phi'\,\} \vdash Y} \qquad \frac{{}_5\{\,\Phi\,\} \vdash [\mathrm{u}_{j_0}](\mu Z.(\mathrm{W}_{j_0} = 1) \vee [\,]Z)}{{}_6\{\,\Phi \wedge \Psi\,\} \vdash \mu Z.(\mathrm{W}_{j_0} = 1) \vee [\,]Z}$$

$$\frac{{}_6\{\,\Phi \wedge \Psi\,\} \vdash \mu Z.(\mathrm{W}_{j_0} = 1) \vee [\,]Z}{{}_7\{\,\Phi \wedge (\Psi \vee (\mathrm{W}_{j_0} = 1))\,\} \vdash \mu Z.(\mathrm{W}_{j_0} = 1) \vee [\,]Z}$$

$$\frac{{}_7\{\,\Phi \wedge (\Psi \vee (\mathrm{W}_{j_0} = 1))\,\} \vdash \mu Z.(\mathrm{W}_{j_0} = 1) \vee [\,]Z}{{}_8\{\,\Phi \wedge (\Psi \vee (\mathrm{W}_{j_0} = 1))\,\} \vdash (\mathrm{W}_{j_0} = 1) \vee [\,]Z}$$

$$\frac{{}_8\{\,\Phi \wedge (\Psi \vee (\mathrm{W}_{j_0} = 1))\,\} \vdash (\mathrm{W}_{j_0} = 1) \vee [\,]Z}{{}_9\{\,\Phi \wedge \neg\Psi \wedge (\mathrm{W}_{j_0} = 1)\,\} \vdash (\mathrm{W}_{j_0} = 1) \qquad {}_{10}\{\,\Phi \wedge \Psi \wedge (\mathrm{W}_{j_0} = 0)\,\} \vdash [\,]Z}$$

$$\frac{{}_{10}\{\,\Phi \wedge \Psi \wedge (\mathrm{W}_{j_0} = 0)\,\} \vdash [\,]Z}{{}_{11}\{\,\Phi \wedge \Phi''\,\} \vdash Z}$$

Some notes on the tableau:

At 4, the value of Φ' is irrelevant: the terminal is successful.

At 11, the enthusiastic reader should check that $\Phi'' = \Psi \vee (\mathrm{W}_{j_0} = 1)$ and so 11 satisfies the inclusion condition.

The main point is to show that $\sqsubset_7$ is well-founded. For this, suppose that $M \sqsupset_7 M' \sqsupset_7 M''$. Then we have that $\mathrm{P} = 0$ and $\mathrm{U}_{j_0} = 1$ for both M and M' (from node 10, since M goes to M' via node 10). Combining these conditions with the invariants in Φ tells us that the only transitions that may fire to give M' from M are v_j and s_i, for any i, j; now if we assign a non-negative measure δ to markings by $\delta(M) = \sum_i \mathrm{R}_i + \sum_j \mathrm{W}_j$, we see that $\delta(M') < \delta(M)$, and so $\sqsubset_7$ is well-founded.

The preceding examples have all used place-invariants to describe the sets about which fix-point properties are to be proved, so one may wonder whether this is a general rule. In fact it is not; it is not hard to construct nets with fix-point properties that cannot be described by any combination of linear invariants, or indeed by any invariant. However, many systems one meets in practice are designed in such a way that interesting sets of states are describable by invariants, and so finding a likely-looking invariant is frequently a good way of starting to build a tableau.

6. Proof of Soundness and Completeness

6.1 Approximants.

We use the notion of ordinal approximants of fix-points, so we now define these.

Let $\sigma Z.\Phi$ be a fix-point formula, and $\mathcal{V}$ be a valuation giving values to the free variables of $\sigma Z.\Phi$. For an ordinal α, the denotation wrt $\mathcal{V}$ of the *α'th approximant* $\sigma^\alpha Z.\Phi$ is defined inductively by
(for zero)

$$\|\mu^0 Z.\Phi\|_{\mathcal{V}} = \varnothing \qquad \|\nu^0 Z.\Phi\|_{\mathcal{V}} = \mathcal{M}$$

(for a successor ordinal $\alpha + 1$)

$$\|\sigma^{\alpha+1} Z.\Phi\|_{\mathcal{V}} = \|\Phi\|_{\mathcal{V}[Z:=\|\sigma^\alpha Z.\Phi\|_{\mathcal{V}}]}$$

(for a limit ordinal λ)

$$\|\mu^\lambda Z.\Phi\|_\mathcal{V} = \bigcup_{\alpha<\lambda} \|\mu^\alpha Z.\Phi\|_\mathcal{V} \qquad \|\nu^\lambda Z.\Phi\|_\mathcal{V} = \bigcap_{\alpha<\lambda} \|\nu^\alpha Z.\Phi\|_\mathcal{V}$$

By the Knaster–Tarski theorem,

$$\|\mu Z.\Phi\|_\mathcal{V} = \bigcup_{\alpha\in\mathrm{Ord}} \|\mu^\alpha Z.\Phi\|_\mathcal{V} \qquad \|\nu Z.\Phi\|_\mathcal{V} = \bigcap_{\alpha\in\mathrm{Ord}} \|\nu^\alpha Z.\Phi\|_\mathcal{V}$$

where Ord is the class of ordinals, or, to avoid foundational problems, the set of ordinals $\{\,\alpha \mid \mathrm{card}\,\alpha \leq \mathrm{card}\,\mathcal{M}\,\}$. (And hence we shall use the notation that $\sigma^\infty Z.\Phi$ means $\sigma Z.\Phi$.)

6.2 Soundness.

Lemma 1. *If a marking M fails the premise of a modal, boolean or (Weakening) rule, there is a path from M to a marking M' in a conclusion, and M' fails the conclusion.*

Proof. By definition of the rules and semantics. □

Lemma 2. *Let $\mathbf{n} = \mathbf{M} \vdash \sigma Z.\Phi$ be the root of a v-i successful tableau, and without loss of generality suppose that the rule applied to $\mathbf{n}$ is (Unfolding). If $\mathbf{n}$ is a false node, so that there exists $M \in \mathbf{M}$ such that $M \not\models_\mathcal{V} \sigma Z.\Phi$, then there exist a marking M', a node $\mathbf{n}' = \mathbf{M}' \vdash Z'$ such that Z' is free in $\sigma Z.\Phi$, and a finite sequence of markings $M = M_0, M_1, \ldots, M_n$ $(n \geq 0)$ such that $M_i \sqsupset_\mathbf{n} M_{i+1}$, $M_n@\mathbf{n} \dashrightarrow M'@\mathbf{n}'$, and $M' \not\models_\mathcal{V} Z'$.*

Proof. The proof is by induction on the σ-height of the tableau (σ-height is the maximum depth of nesting of fix-point formulae in the root formula). The argument is in two parts (ν) and (μ); each part has cases (a) and (b): the (a) cases apply in the base case of the induction (when Φ has no embedded fix-points), and either (a) or (b) may apply in the inductive step.

(ν) $\mathbf{n} = \mathbf{M} \vdash \nu Z.\Phi$

If $M \not\models_\mathcal{V} \nu Z.\Phi$, then there exists a least ordinal, which must be a successor ordinal, say $\alpha_0 + 1$, such that $M \not\models_\mathcal{V} \nu^{\alpha_0+1} Z.\Phi$ (by definition of approximants and Knaster–Tarski). Let $\mathcal{V}' = \mathcal{V}[Z := \|\nu^{\alpha_0} Z.\Phi\|_\mathcal{V}]$. Now, the child of $\mathbf{n}$ is $\mathbf{M} \vdash \Phi$, and by definition of approximant, $M \not\models_{\mathcal{V}'} \Phi$. Apply Lemma 1 repeatedly to get either

(a) (only this case occurs in the base case of the induction) a path $M@\mathbf{n} \longrightarrow M_1'@\mathbf{n}_1'$, where $\mathbf{n}_1' = \mathbf{M}_1' \vdash Z_1'$ and $M_1' \not\models_{\mathcal{V}'} Z_1'$; then $M@\mathbf{n} \dashrightarrow M_1'@\mathbf{n}_1'$; or

(b) a path $M@\mathbf{n} \longrightarrow M_0''@\mathbf{n}''$, where $\mathbf{n}''$ is a σ-node, and M_0'' is false at $\mathbf{n}''$ wrt $\mathcal{V}'$. In this case, apply the lemma inductively to the tableau rooted at $\mathbf{n}''$ wrt the valuation $\mathcal{V}'$ to get $\mathbf{n}_1' = \mathbf{M}_1' \vdash Z_1'$, and a sequence of markings $M_0'' \sqsupset_{\mathbf{n}''} \cdots \sqsupset_{\mathbf{n}''} M_k''$ such that $M_k''@\mathbf{n}'' \dashrightarrow M_1'@\mathbf{n}_1'$, and $M_1' \not\models_{\mathcal{V}'} Z_1'$; and by the definition of extended path and $\sqsupset$, here again $M@\mathbf{n} \dashrightarrow M_1'@\mathbf{n}_1'$.

Now in either of these cases, if Z_1' is free in $\nu Z.\Phi$, set M' to M_1' and Z' to Z_1', and the proof is complete. Otherwise, we have $Z_1' = Z$ (since no other variables are possible, and non-variable terminals cannot be false), so $M_1' \not\models_{\mathcal{V}'} Z$, so $M_1' \not\models_\mathcal{V} \nu^{\alpha_0} Z.\Phi$. But since the tableau is v-i successful, $M_1' \in \mathbf{M}$, and $M_1' \not\models_{\mathcal{V}} \nu^{\alpha_1+1}Z.\Phi$ for some $\alpha_1+1 \le \alpha_0$, so we may repeat the whole argument to get M_2'. If the sequence ever terminates with some Z_n' being free in $\nu Z.\Phi$, the proof is complete; otherwise we have an infinite descending chain of ordinals $\alpha_0 > \alpha_1 > \cdots$, which cannot happen.

(μ) $\mathbf{n} = \mathbf{M} \vdash \mu Z.\Phi$
If $M \not\models_{\mathcal{V}} \mu Z.\Phi$, then let $\mathcal{V}' = \mathcal{V}[Z := \|\mu Z.\Phi\|_{\mathcal{V}}]$, and then we have $M \not\models_{\mathcal{V}'} \Phi$ since least fix-points are fixed. The rest of the argument is as for (ν) omitting the ordinals α_i, for in this case an infinite sequence $M_1', M_2', \ldots$ directly contradicts the well-foundedness of $\sqsubset_{\mathbf{n}}$. □

Proposition. *The root of a successful tableau is true.*

Proof. A successful tableau is v-i successful, so by Lemma 2 if its root is false, it has a false free variable node; but the free variable nodes in a successful tableau are true by definition of success. □

6.3 Completeness.

We describe a (non-effective) function to build a tableau for $\mathbf{M}_0 \models_{\mathcal{V}} \Phi_0$, and show that it produces a successful tableau.

Firstly, we define a notion of signature for a marking wrt a subformula Φ of Φ_0 appearing at a node $\mathbf{n}$ in a tableau.

Let $Z_1, \ldots Z_n$ be the variables free in Φ but bound in Φ_0, with each Z_i bound by $\sigma Z_i.\Phi_i$ (and assume each $\sigma Z_{i+1}.\Phi_{i+1}$ is a subformula of Φ_i, so that the Z_i are bound in order), and suppose $Z_{i_1}, \ldots Z_{i_k}$ are the μ-variables among the Z_i. If $\alpha_1, \ldots, \alpha_k$ are ordinals, we define a valuation $\mathcal{V}_{\mathbf{n}}(\alpha_1, \ldots, \alpha_k)$ by $\mathcal{V}_{\mathbf{n}}(\alpha_1, \ldots, \alpha_k) = \mathcal{V}_n$, where

$$\begin{aligned}
\mathcal{V}_0 &= \mathcal{V} \\
\mathcal{V}_{i_j} &= \mathcal{V}_{i_j-1}[Z_{i_j} := \|\mu^{\alpha_j} Z_{i_j}.\Phi_{i_j}\|_{\mathcal{V}_{i_j-1}}] \\
\mathcal{V}_i &= \mathcal{V}_{i-1}[Z_i := \|\nu Z_i.\Phi_i\|_{V_{i-1}}] \quad \text{for } i \notin \{i_1, \ldots, i_k\}
\end{aligned}$$

and the *signature* $\mathrm{sig}(M, \mathbf{n})$ of a marking M at $\mathbf{n}$ to be the lexicographically least sequence of α_i such that $M \models_{\mathcal{V}_{\mathbf{n}}(\alpha_1,\ldots,\alpha_k)} \Phi$. We take the lexicographic ordering on signatures.

We now build a tableau by applying the following procedure to the root until there are no non-terminal leaves in the tableau:

Choose a non-terminal leaf node $\mathbf{n} = \mathbf{M} \vdash \Phi$. If Φ is

- $\Phi_1 \wedge \Phi_2$, then apply (And) to $\mathbf{n}$ with consequences both labelled by $\mathbf{M}$.
- $\Phi_1 \vee \Phi_2$, then apply (Or), creating consequent nodes $\mathbf{n}_1$ and $\mathbf{n}_2$, let $\mathbf{M}_1 = \{\, M \in \mathbf{M} \mid \mathrm{sig}(M, \mathbf{n}_1) \le \mathrm{sig}(M, \mathbf{n}_2) \,\}$ and $\mathbf{M}_2 = \mathbf{M} - \mathbf{M}_1$, and label $\mathbf{n}_i$ with $\mathbf{M}_i$.
- $[\,]\Phi_1$, then apply (Box), creating node $\mathbf{n}_1$, and label $\mathbf{n}_1$ with $\mathbf{M}_1 = \{\, M' \mid \exists M \in \mathbf{M} \,.\, M \longrightarrow M' \,\}$.
- $\langle\rangle\Phi_1$, then apply (Diamond), creating node $\mathbf{n}_1$, let $f\colon \mathbf{M} \to \mathbf{T}$ be such that if $M \xrightarrow{f(M)} M'$, then $\forall M''.\, M \longrightarrow M'' \Rightarrow \mathrm{sig}(M', \mathbf{n}_1) \le \mathrm{sig}(M'', \mathbf{n}_1)$, and label $\mathbf{n}_1$ according to f and the (Diamond) rule.

- $\sigma Z.\Phi_1$, let $\mathbf{M}_1 = \|\sigma Z.\Phi_1\|_{\mathcal{V}_{\mathbf{n}}(\infty,\ldots,\infty)}$, apply (Weakening) to $\mathbf{n}$ to create $\mathbf{n}_1$ and label it with $\mathbf{M}_1$, then apply (Unfolding) to $\mathbf{n}_1$.

Lemma 1. *Each node of the resulting tableau is true.*

Proof. By inspection of the procedure, if it is called on a true node, it creates true nodes. □

Lemma 2. *The bound variable terminals of the tableau satisfy the inclusion condition for success.*

Proof. When a fix-point formula Φ appears in the tableau, it is immediately weakened to $\|\Phi\| \vdash \Phi$ before being unfolded to $\|\Phi\| \vdash U$, and since the terminals $S \vdash U$ are true, $S \subseteq \|\Phi\|$. (Note: in fact, it suffices to weaken to that part of the whole denotation that is reachable from $\mathbf{M}$, since the set of reachable markings is closed under application of the other rules.) □

Lemma 3. *If $M@\mathbf{n} \longrightarrow \mathbf{M}'@\mathbf{n}'$ is a path in the tableau from the premise to the conclusion of any rule other than μ-unfolding, then* $\mathrm{sig}(M', \mathbf{n}') \leq \mathrm{sig}(M, \mathbf{n})$.

Proof. In the case of (Box), from the definitions, the signature of M is the maximum of the signatures of its true successors; in the case of (And), similarly; for (Diamond) and (Or), we chose in the construction the true successor with least signature, and hence $\mathrm{sig}(M', \mathbf{n}') = \mathrm{sig}(M, \mathbf{n})$. (Weakening) is trivial. For ν-unfolding, we have $M \vDash_{\mathcal{V}_{\mathbf{n}}(\alpha_1,\ldots,\alpha_k)} \nu Z.\Phi$ iff $M \vDash_{\mathcal{V}'_{\mathbf{n}}(\alpha_1,\ldots,\alpha_k)} \Phi$ (by definition), and therefore $\mathrm{sig}(M, \mathbf{n}') = \mathrm{sig}(M, \mathbf{n})$. □

Lemma 4. *If $M@\mathbf{n} \longrightarrow \mathbf{M}@\mathbf{n}_1$ is a path across a μ-unfolding rule, and* $\mathrm{sig}(M, \mathbf{n}) = (\alpha_1, \ldots, \alpha_k)$, *then* $\mathrm{sig}(M, \mathbf{n}_1) = (\alpha_1, \ldots, \alpha_k, \alpha_{k+1})$ *for some α_{k+1}.*

Proof. From the definitions of $\mathcal{V}_{\mathbf{n}}$ and of approximants, $M \vDash_{\mathcal{V}_{\mathbf{n}}(\alpha_1,\ldots,\alpha_k)} \mu Z.\Phi$ iff $M \vDash_{\mathcal{V}_{\mathbf{n}}(\alpha_1,\ldots,\alpha_k)} \mu^{\alpha_{k+1}+1} Z.\Phi$ for some $\alpha_{k+1}+1$ iff $M \vDash_{\mathcal{V}_{\mathbf{n}_1}(\alpha_1,\ldots,\alpha_k,\alpha_{k+1})} \Phi$. □

Lemma 5. *If $\mathbf{n} = \mathbf{M} \vdash \mu Z.\Phi$ is the premise of an (Unfolding) rule, with child node $\mathbf{n}_1$, and $M \sqsupset_{\mathbf{n}} M'$, then* $\mathrm{sig}(M', \mathbf{n}_1) < \mathrm{sig}(M, \mathbf{n}_1)$.

Proof. Let $\mathrm{sig}(M, \mathbf{n}) = (\alpha_1, \ldots, \alpha_k)$. If $M \sqsupset_{\mathbf{n}} M'$ then there exists $\mathbf{n}' = \mathbf{M}' \vdash Z$ such that $M@\mathbf{n}_1 \dashrightarrow M'@\mathbf{n}'$. By Lemmas 3 and 4, the first k components of the signature do not increase across any of the rules traversed by the extended path, and so $\mathrm{sig}(M', \mathbf{n}') = (\beta_1, \ldots, \beta_k + 1, 0, \ldots, 0)$ for $(\beta_1, \ldots, \beta_k + 1) \leq (\alpha_1, \ldots, \alpha_k)$. (The components of $\mathrm{sig}(M', \mathbf{n}')$ after β_k are all zero, since they are associated with variables bound after Z, and so the signature does not depend on them.) Now, from the definitions of signature and approximant, $\mathrm{sig}(M', \mathbf{n}_1) = (\beta_1, \ldots, \beta_k) < \mathrm{sig}(M, \mathbf{n}_1)$. □

Corollary. *Since signatures are well-founded, $\sqsubset_{\mathbf{n}}$ is well-founded.* □

Proposition. *The tableau is v-i successful, and hence successful.*

Proof. By Lemmas 1 and 2 and the Corollary to Lemma 5. □

7. Conclusions and Further Work

We have presented a system for reasoning about nets in a powerful temporal logic. By separating fix-point detection from modal properties, we have allowed our system to use reasoning that may be specific to a particular net or class of nets (or more generally, to nets: this system is a particular case of one for general transition systems, and the system can as well be applied to CCS (see [Bradfield & Stirling 91]) or any other suitable model). We have given examples which demonstrate some of the styles of reasoning, such as the use of place invariants as candidate fix-points, and inductive arguments on the size of nets.

Work in progress is investigating questions of decidability, and completeness of the system with respect to particular finite presentations of sets of markings. (It is easily shown that the full mu-calculus is so expressive that there are small formulae on small nets which are undecidable; whether there are any non-trivial classes of nets and/or formulae which are decidable, is being examined.)

Future work will look at the application of compositional reasoning techniques to our system, and the use of other structural features of nets.

8. Acknowledgements

This system was developed from the finite model-checking tableau system developed with David Walker by my supervisor Colin Stirling, for whose advice and encouragement throughout I am deeply indebted.

I thank the referees of the International Conference on Petri Nets and four referees of *Advances in Petri Nets* for their helpful comments and questions.

This work is part of a Ph.D. supported by a Research Studentship from the U.K. Science and Engineering Research Council.

9. References

Best, E. [1986]. Structure theory of Petri nets: the free choice hiatus. *Advances in Petri Nets 1986, Part I.* LNCS 254.

Bradfield, J. C., Stirling, C. P. [1991]. Local model-checking for infinite state spaces. To appear in *Theoretical Computer Science.*

Clarke, E. M., Emerson, E. A., Sistla, A. P. [1983]. Automatic verification of finite-state concurrent systems using temporal logic specifications: a practical approach. *Proc. Tenth ACM Symposium on Principles of Programming Languages.* ACM, Austin, Texas.

Emerson, E. A., Lei, C.-L. [1986]. Efficient model checking in fragments of the propositional mu-calculus. *Proc. First IEEE Symposium on Logic in Computer Science* (267–278). Computer Society Press.

Hüttel, H. [1990] *SnS* can be modally characterized. *Theoretical Computer Science* **74** (239–248).

Kozen, D. [1983]. Results on the propositional mu-calculus. *Theoretical Computer Science* **27**.

Lautenbach, K. [1986]. Linear algebraic techniques for place/transition nets. *Advances in Petri Nets 1986, Part I.* LNCS 254.

Pratt, V. [1981]. A decidable mu-calculus. *Proc. 22nd IEEE Foundations of Computer Science.*

Stirling, C. P., Walker, D. J. [1989]. Local model checking in the modal mu-calculus. *Proc. International Joint Conference on Theory and Practice of Software Development* (369–382). LNCS 351.