Temporal Logic and Categories of Petri Nets

Carolyn Brown* and Doug Gurr†
School of Cognitive and Computing Sciences
University of Sussex, Falmer, Brighton, BN1 9QH

Abstract. We present a novel method for proving temporal properties of the behaviour a Petri net. Unlike existing methods, which involve an exhaustive examination of the transition system representing all behaviours of the net, our approach uses morphisms dependent only on the static structure of the net. These morphisms correspond to simulations. We restrict the analysis of dynamic behaviours to particularly simple nets (test nets), and establish temporal properties of a complex net by considering morphisms between it and various test nets. This approach is computationally efficient, and the construction of test nets is facilitated by the graphical representation of nets. The use of category theory permits a natural modular approach to proving properties of nets.

Our main result is the syntactic characterisation of two expressive classes of formulae: those whose satisfaction is preserved by morphisms and those whose satisfaction is reflected.

1 Introduction

Proving properties of the operational behaviour of Petri nets is computationally expensive, as most existing techniques [1] involve an exhaustive examination of the labelled transition system representing all possible markings and behaviours of the net. In this paper we describe a novel

2 Definitions concerning Petri Nets

We recall some elementary definitions of Petri net theory: details may be found in [13].

Definition 1 *A*

condition relations of a net from events to computations in the evident way. For parallel composition, we define

$$pre(c_0 + c_1) = pre(c_0) + pre(c_1) \quad \text{and} \quad post(c_0 + c_1) = post(c_0) + post(c_1).$$

Defining pre and postcondition relations for sequential composition requires a little care. Note that sequential composition is associative, even

and so $\langle M_1, M_1' \rangle \in F^+$. That $\langle$

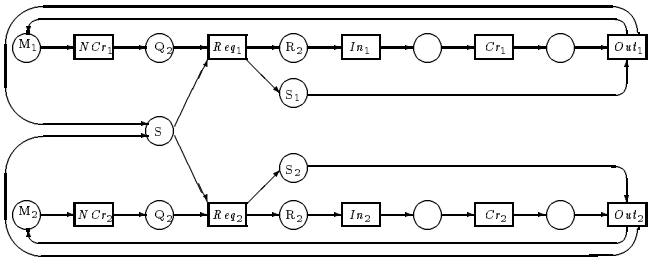

FIGURE 1. *The net N_1: mutual exclusion*

It is readily proved that every behaviour of the net $\langle C, N_2\rangle$ is a sequence of form $Beg_a\,; End_a\,; Beg_b\,; End_b\,; Beg_c\,; End_c\,; \ldots$ where a, b and c range over $\{1, 2\}$. We shall add to the net N_2 a trivial event $*$, which has empty pre- and post-condition set. The resultant net $\langle C, N_2 + \bot\rangle$ is the coproduct in $\mathbf{MNet}^+$ of N_2 with the marked net $\bot = \langle\emptyset, \{*\}, \{*\}, 0, 0\rangle$ (where 0 denotes the empty multirelation). There is a morphism $\langle f, F\rangle$ in $\mathbf{MNet}^+$ from $\langle M_1 + M_2 + S, N_1\rangle$ to $\langle C, N_2 + \bot\rangle$ given by:

$f(Req_i) = Beg_i \quad f(Out_i) = End_i \quad f(e) = *$ for all other events e
$F(C) = S \qquad F(C_i) = S_i$

By Corollary 3, the existence of this morphism shows that the net $\langle C, N_2 + \bot\rangle$ can simulate any behaviour of the net $\langle M_1 + M_2 + S, N_1\rangle$. Since the behaviour of the image net is so restricted (indeed, $\langle f, F\rangle$ is minimal in the sense of Definition 11), this proves an important feature of the marked net $\langle M_1 + M_2 + S, N_1\rangle$, that it can never reach a state in which Req_2 can occur if Req_1 has occurred and Out_1 has not. This, together with the analogous property for Req_2, ensures that $\langle M_1 + M_2 + S, N_1\rangle$ preserves mutual exclusion of the behaviours $In_1\,; Cr_1\,; Out_1$ and $In_2\,; Cr_2\,; Out_2$. This example is particularly simple. Note, however, for any net $\langle M, N\rangle$ intended as a mutual exclusion algorithm, the existence of a morphism from $\langle M, N\rangle$ to $\langle C, N_2 + \bot\rangle$ can be used to demonstrate that $\langle M, N\rangle$ preserves mutual exclusion.

The net $\langle C, N_2\rangle$ describes the behaviour of the shared resource, abstracting away from the competing processes. A different abstraction is given in the net $\langle m_1 + m_2 + s, N_3\rangle$ below, which describes only the possible states of the processes (critical, requesting entry to the critical region, or neither of these) and how these interact. There is a morphism $\langle g, G\rangle$ in $\mathbf{MNet}^+$ from $\langle m_1 + m_2 + s, N_3\rangle$ to $\langle M_1 + M_2 + S, N_1\rangle$ given by:

$g(ncr_i) = Ncr_i \quad g(req_i) = Req_i \quad g(cr_i) = In_i\,;\, Cr_i\,;\, In_i$
$G(M_i) = m_i \quad G(Q_i) = q_i \quad G(R_i) = G(S_i) = r_i$
$G(S) = s \quad G(b) = *$ for all other conditions b of N_1

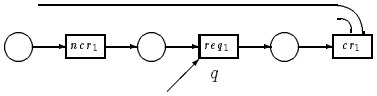

is preserved by morphisms then $\langle M', N' \rangle$ also has the property described by ϕ. If $\langle M', N' \rangle$ has the property described by ψ and satisfaction of ψ is reflected by morphisms then $\langle M, N \rangle$ also has the property described by ψ.

In Section 5.1 we reprove the results of this section in our formal set-

that $\langle M, N\rangle \models_\theta \neg E(t)$ if and only if $\langle M, N\rangle \models_\theta \langle t\rangle \mathbf{ff}$, that is, precisely when the computation interpreting t is not enabled. Observe that if α is interpreted by the identity step id_{nb} then $\langle M, N\rangle \models_\theta E(\alpha)$ whenever the condition b is marked in $\langle M, N\rangle$ with at least n tokens. In general such properties as mutual exclusion or freedom from deadlock can be expressed in terms of the enabling of events. For example, the fact that two events e_0 and e_1 cannot occur concurrently is expressed by the formula $\neg E(\alpha_0 + \alpha_1)$, where α_i is interpreted in $\langle M, N\rangle$ by e_i.

We wish to specify and reason about both the overall behaviour of a net and individual enabled steps: we therefore turn our attention from steps to step sequences, and extend $\mathcal{M}$ to the temporal logic $\mathcal{T}$ by considering the modal formulae which hold on computation paths rather than at individual states. $\mathcal{T}$ is given by:

$$\phi ::= \mathbf{tt} \;\Big|\; \neg\phi \;\Big|\; \phi \wedge \phi \;\Big|\; \forall x.\phi \;\Big|\; [t]\phi \;\Big|\; \Box\phi \quad \text{for } t \text{ a closed term.}$$

Definition 5 *The interpretation of a closed formula ϕ of $\mathcal{T}$ relative to an interpretation θ of $\mathcal{T}$ in a marked net $\langle M, N\rangle$ is a set of step sequences $\sigma = \sigma_0\,;\,\sigma_1\,;\,\ldots$ given as follows:*

$\sigma \in [\![\,\mathbf{tt}\,]\!]_\theta$ *for any* σ
$\sigma \in [\![\,\neg\phi\,]\!]_\theta$ *iff it is not the case that* $\sigma \in [\![\,\phi\,]\!]_\theta$
$\sigma \in [\![\,\phi \wedge \psi\,]\!]_\theta$ *iff* $\sigma \in [\![\,\phi\,]\!]_\theta \cap [\![\,\psi\,]\!]_\theta$
$\sigma \in [\![\,\forall x.\phi\,]\!]_\theta$ *iff for all* $\alpha \in dom(\theta)$ *we have* $\sigma \in [\![\,\phi[\alpha/x]\,]\!]_\theta$
$\sigma \in [\![\,[t]\phi\,]\!]_\theta$ *iff whenever there exists* k *such that* $\sigma_0\,;\,\sigma_1\,;\,\ldots\,;\,\sigma_k = \theta(t)$ *then* $\overline{\sigma}_{k+1} \in [\![\,\phi\,]\!]_\theta$
$\sigma \in [\![\,\Box\phi\,]\!]_\theta$ *iff for each* i *we have* $\sigma_i\,;\,\sigma_{i+1}\,;\,\ldots \in [\![\,\phi\,]\!]_\theta$.

The satisfaction relation $\models$ between marked nets and closed formulae of $\mathcal{T}$ relative to θ is given by $\langle M, N\rangle \models_\theta \phi$ iff every computation of $\langle M, N\rangle$ is an element of $[\![\,\phi\,]\!]_\theta$.

This interpretation gives the usual meaning to the derived operators. Thus $\langle M, N\rangle \models_\theta \Diamond\phi$ precisely when every computation of $\langle M, N\rangle$ eventually satisfies ϕ, while $\langle M, N\rangle \models_\theta \langle t\rangle\phi$ precisely when $\langle M, N\rangle$ can evolve under $\theta(t)$ to $\langle M', N\rangle$ and $\langle M', N\rangle \models_\theta \phi$. We could define $\models_\theta$ relative to certain fairness or liveness assumptions, considering, for example, only those step sequences which are *weakly* or *strongly fair* [7]. In his temporal logic for occurrence nets [14], Reisig restricts attention to behaviours in which no condition ever contains more than one token.

The language $\mathcal{T}$ expresses many interesting properties of nets, both positive (what can be enabled) and negative (what cannot be enabled).

For example, mutual exclusion of events interpreting α_0 and α_1 is expressed by satisfaction of the formula $\Box\neg E(\alpha_0 + \alpha_1)$ while freedom from deadlock is expressed by satisfaction of the formula $\Box\exists x.E(x)$.

In practice, the graphical representation of nets facilitates the creative process of constructing test nets. It appears difficult to find an algorithm

Let $\mathcal{L}$

- Suppose $\sigma \in [\![\forall x.\phi]\!]_\theta$. Then for each $\alpha \in dom(\theta)$ we have $\sigma \in [\![\phi[\alpha/x]]\!]_\theta$ and, since $dom(f\theta) = dom(\theta)$ and ϕ-computations are preserved, $f\sigma \in [\![\phi[\alpha/x]]\!]_{f\theta}$ for each $\alpha \in dom(f\theta)$. Hence $f\sigma \in [\![\forall x.\phi]\!]_{f\theta}$ and $\forall x.\phi$-computations are preserved.

□

Proposition 17 *If $\langle f, F\rangle$ reflects ϕ- and ψ-computations then $\langle f, F\rangle$ reflects*

- *$\phi \wedge \psi$-computations,*
- *$[t]\phi$-computations,*
- *$\Box\phi$-computations and*
- *$\forall x.\phi$-computations.*

Proof:

- Suppose $f\sigma \in [\![\phi \wedge \psi]\!]_{f\theta} = [\![\phi]\!]_{f\theta} \cap [\![\psi]\!]_{f\theta}$. Then $f\sigma \in [\![\phi]\!]_{f\theta}$ and $f\sigma \in [\![\psi]\!]_{f\theta}$ and since both ϕ- and ψ-computations are reflected, $\sigma \in [\![\phi]\!]_\theta$ and $\sigma \in [\![\psi]\!]_\theta$, whence $\sigma \in [\![\phi \wedge \psi]\!]_\theta$. Thus $f^{-1}([\![\phi \wedge \psi]\!]_{f\theta}) \subseteq [\![\phi \wedge \psi]\!]_\theta$ and $\langle f, F\rangle$ reflects $\phi \wedge \psi$-computations.
- Suppose $f\sigma \in [\![[t]\phi]\!]_{f\theta}$ and $\sigma_0 ; \sigma_1 \ldots \sigma_k = \theta(t)$. Then putting $\sigma' = f\sigma$ we can find l such that $f(\sigma_0 ; \sigma_1 \ldots \sigma_k) = \sigma'_0 ; \sigma'_1 ; \ldots \sigma'_l$

Example 1 *The following formulae are preserved:*

$E(t)$	*$\theta(t)$ is enabled*
$\exists x.E(x)$	*some $\theta(\alpha)$ is enabled*
$E(t) \vee E(t')$	*either $\theta(t)$ or $\theta(t')$ is enabled.*

The following formulae are reflected:

$\neg E(t)$	*$\theta(t)$ is not enabled*
$\Diamond\neg E(t)$	*eventually $\theta(t)$ is disabled*
$\Box\neg E(t)$	*$\theta(t)$ is never enabled*
$\forall x.\neg E(x)$	*no $\theta(\alpha)$ is enabled (relative deadlock).*
$\Box\Diamond\neg E(t)$	*a marking is always reachable in which $\theta(t)$ is disabled.*

The following formulae are minimally preserved:

$\Diamond E(t)$	*$\theta(t)$ is eventually enabled*
$\Diamond\neg E(t)$	*$\theta(t)$ is eventually disabled*
$\neg E(t)$	*$\theta(t)$ is not enabled*
$\Diamond\exists x.E(x)$	*eventually some $\theta(\alpha)$ is enabled*
$\forall x.\neg E(x)$	*no $\theta(\alpha)$ is enabled (relative deadlock)*
$\Box\exists x.E(x)$	*some $\theta(t)$ is always enabled.*

The following formulae are minimally reflected:

$E(t)$	*$\theta(t)$ is enabled*
$\exists x.E(x)$	*some $\theta(\alpha)$ is enabled*
$\Box\exists x.E(x)$	*some $\theta(t)$ is always enabled.*

There are many more examples of formulae whose properties we can deduce from the results presented above. A selection is given in Example 2.

A common situation is illustrated by the following lemma:

Lemma 21 *Let I index the set $\{t_i \mid f\theta(t_i) = f\theta(t)$. If $f\sigma \in [\![E(t)]\!]_{f\theta}$ then $\sigma \in \bigcup_{i\in I} [\![E(t_i)]\!]_\theta$ and whenever $\langle M', N'\rangle \models_{f\theta} E(t)$ it is the case that $\langle M, N\rangle \models_\theta \bigvee_I E(t_i)$.*

Proof: Straightforward □

Remark 22 *It is an immediate consequence of the previous lemma that if $f\theta t = f\theta t'$ implies that $\theta t = \theta t'$ (and in particular, if f is injective) then $E(t)$-computations are minimally reflected and so $E(t)$ is minimally reflected.*

It is not in general the case that $\Box\phi$ is preserved or that $\Box\phi$-computations are preserved, even by a minimal morphism. For example, returning to the net N illustrated at the start of Section 3, the identity morphism $\langle id, id\rangle$ maps $\langle b_0, N\rangle$ to $\langle 2b_0, N\rangle$ but $\langle b_0, N\rangle \models_\theta \Box\neg E(\alpha_0)$ and $\langle 2b_0, N\rangle \not\models_{\mathrm{id}\theta} \Box\neg E(\alpha_0)$. The following lemma establishes a special case in which we can infer properties of a formula $\Box\phi$ from properties of ϕ.

Lemma 23
$\Box\Diamond E(t)$-computations are preserved and $\Box\Diamond E(t)$ is minimally preserved. If $f\theta(t) = f\theta(t')$ implies that $\theta(t) = \theta(t')$ and $\langle f, F\rangle$ is minimal then $\langle f, F\rangle$ reflects $\Box\Diamond E(t)$-computations.

Proof: Suppose $\sigma \in [\![\Box\Diamond E(t)]\!]_\theta$. Then for every i there exists j such that $\overline{\sigma}_{i+j} \in [\![E(t)]\!]_\theta$. Suppose that $f\sigma \notin [\![\Box\Diamond E(t)]\!]_{f\theta}$. Then there exists some k such that for all l, $\overline{f\sigma}_{k+l} \notin [\![E(t)]\!]_{f\theta}$. It follows that there exists $m \geq k$ such that for all l, $f(\overline{\sigma}_{m+l}) \notin [\![E(t)]\!]_{f\theta}$. Since $E(t)$-computations are preserved, this would imply that we could find some m such that for all l, $\overline{\sigma}_{m+l} \notin [\![E(t)]\!]_\theta$, which contradicts our assumption that $\sigma \in [\![\Box\Diamond E(t)]\!]_\theta$. Hence $f\sigma \in [\![E(t)]\!]_{f\theta}$.

It follows that $\Box\Diamond E(t)$ is minimally preserved, by Proposition 13.

We now show that $\Box\Diamond E(t)$-computations are minimally reflected. Suppose $f\sigma \in [\![\Box\Diamond E(t)]\!]_{f\theta}$. Then for all i there exists j such that $\overline{f\sigma}_{i+j} \in [\![E(t)]\!]_{f\theta}$. It follows that for all i there exists $k \geq j$ such that $f(\overline{\sigma}_{i+k}) \in [\![E(t)]\!]_{f\theta}$. Since $\langle f, F\rangle$ is minimal, it follows from the proof of Lemma 21 that $\langle f, F\rangle$ reflects $E(t)$. Hence for all i there exists k such that $\overline{\sigma}_{i+k} \in [\![E(t)]\!]_\theta$.

□

Remark 24 *Observe that the proof above still goes through if we replace $E(t)$ by any formula ϕ which is preserved and minimally reflected. We can prove the usual dual results for formulae of the form $\Diamond\Box\phi$.*

If we extend T with arbitrary disjunctions then we can prove the following proposition:

Proposition 25 *If $\langle f, F\rangle$: $\langle M, N\rangle \to \langle M', gien\ \psi\psi Tf\psi\psi$*[illegible]

Proof: Suppose for example that $\langle M', N' \rangle \models_{f\theta} \Box E(t)$. We show that $\langle M, N \rangle \models_{\theta} \Box \bigvee_I E(t_i)$. In every computation of $\langle M', N' \rangle$ the computation $\theta(t)$ is continuously enabled. By minimality, in every computation of $\langle M, N \rangle$, there is always a computation enabled whose image under f equals $f(\theta t)$. Let I index the set $\{t_i \mid f\theta(t_i) = f\theta(t)\}$. Then $\langle M, N \rangle \models_{\theta} \Box \bigvee_I E(t_i)$. □

Note that, as in the case of Lemma 21, if $f\theta(t') = f\theta(t)$ implies that $t' = t$ and $\langle f, F \rangle : \langle M, N \rangle \longrightarrow \langle M', N' \rangle$ is minimal with $\langle M', N' \rangle \models_{f\theta} \Box E(t)$ then $\langle M, N \rangle \models_{\theta} \Box E(t)$.

Proposition 26 *If $\langle f, F \rangle : \langle M, N \rangle \rightarrow \langle M', N' \rangle$ is minimal and I indexes $\{t_i \mid f\theta(t_i) = f\theta(t)\}$, then*

if $\langle M, N \rangle \models_{\theta} \Box\Diamond \bigwedge_I \neg E(t_i)$ then $\langle M', N' \rangle \models_{f\theta} \Box\Diamond \neg E(t)$ and

if $\langle M, N \rangle \models_{\theta} \Diamond \bigwedge_I \neg E(t_i)$ then $\langle M', N' \rangle \models_{f\theta} \Diamond \neg E(t)$.

Proof: Analogous to that of Proposition 25 □

The results of this section together with the proof rules for temporal and modal logic determine a relatively large and expressive class of formulae which are either preserved or reflected by morphisms in $\mathbf{MNet}^{+}$. These formulae occur at all levels of Manna and Pnueli's hierarchy [7, 8].

Example 2 *The state formulae of $\mathcal{T}$ are those given by $\mathtt{tt} \mid E(t) \mid \phi \wedge \phi \mid \neg\phi$. If ϕ and ψ are state formulae then:*

- *$\Box\phi$ describes a safety property. Many such formulae, including mutual exclusion $\Box\neg E(t_0 + t_1))$, are reflected.*
- *$\Diamond\phi$ describes a termination property, guaranteeing a one-time goal. An example is $\Diamond E(\alpha)$, which is both minimally preserved and minimally reflected.*
- *$\Box\Diamond\phi$ describes a recurrence property or response property. An example is $\Box(E(t_0) \rightarrow \Diamond E(t_1))$, which is minimally preserved and minimally reflected.*
- *$\Diamond\Box\phi$ describes a persistence property. As an example, $\Diamond\Box E(t)$ is minimally reflected.*
- *$\Diamond\Box\phi \vee \Box\Diamond\psi$ describes a progress property. An example is $\Box(\Box\Diamond E(t_0) \rightarrow \Box\Diamond E(t_1))$ (strong fairness) which is minimally preserved, and furthermore is reflected by minimal morphisms $\langle f, F \rangle$ such that f is injective.*

5.1 Proving Properties of Nets

We now outline the formal proofs that the net $\langle M_1 + M_2 + S, N_1 \rangle$ of Section 3.2 preserves mutual exclusion and satisfies absence of starvation. These proofs follow our previous reasoning closely. For absence of starvation, we shall assume an invertible interpretation θ in $\langle m_1 + m_2 + s, N_3 \rangle$ with inverse η. The fact that s is marked infinitely often is expressed as $\langle m_1 + m_2 + s, N_3 \rangle \models_{\theta} \Box\Diamond E(\eta\, id_s)$. The fact that if q_1 is marked and cr_1 never occurs then q_1 remains marked is expressed as $\langle m_1 + m_2 + s, N_3 \rangle \models_{\theta} (E(\eta\, id_{q_1}) \wedge \neg\Diamond E(\eta\, cr_1)) \rightarrow \Box E(\eta\, id_{q_1}$