

**A TAXONOMY OF FAIRNESS AND TEMPORAL
LOGIC PROBLEMS FOR PETRI NETS**

Rodney R. Howell, Louis E. Rosier, and
Hsu-Chun Yen

Department of Computer Sciences
The University of Texas at Austin
Austin, Texas 78712-1188

TR-88-03

January 1988

A Taxonomy of Fairness and Temporal Logic Problems for Petri Nets*

Rodney R. Howell, Louis E. Rosier
Dept. of Computer Sciences
The University of Texas at Austin
Austin, TX 78712

Hsu-Chun Yen
Dept of Computer Science
Iowa State University
Ames, IA 50011

January 12, 1988

Abstract

In this paper, we define a temporal logic for reasoning about Petri nets. We show the model checking problem for this logic to be PTIME equivalent to the Petri net reachability problem. Using this logic and two refinements, we show the fair nontermination problem to be PTIME equivalent to reachability for several definitions of fairness. For other versions of fairness, this problem is shown to be either PTIME equivalent to the boundedness problem or highly undecidable. In all, 24 versions of fairness are examined.

1 Introduction

In the specification and analysis of concurrent systems, some notion of fairness is often necessary to exclude from consideration certain computations which prevent particular events from occurring, even though these events may be possible infinitely often. Many definitions of fairness have been proposed (see, e.g., [5,6,7,8,12,19,20,28]), and each has its merit in particular applications. Several versions of fairness have been defined (or adapted) for Petri nets [5,6,7,8,28], a powerful formalism often employed to model concurrent systems [26]. Decidability issues concerning fairness in Petri nets were considered in [6,7,34]. One problem examined in these papers is the fair nontermination problem; i.e., for a certain definition of fairness, does there exist an infinite fair computation? Though the decidability of a number of these problems has been determined, only a few rough complexity bounds have been given for those known to be decidable [7]. In [15], we examined the complexity of the fair nontermination problem for conflict-free Petri nets with respect to several definitions of fairness. In this paper, we extend this study to general Petri nets, examining the 24 versions of fairness presented in [5,6,7,8,12,19,20,28]. The results of this study are summarized in a table in Section 5.

In the study of Petri nets, a scarcity of knowledge concerning computational complexity is by no means unique to the fair nontermination problem. For example, the precise complexity of the reachability problem has remained elusive for many years. For this problem, the most efficient algorithm is not primitive recursive [18,23], whereas the best known lower bound is exponential space [21]. We show in this paper that there is a close relationship between the reachability problem and several of the fair nontermination problems. In particular, we show several of these problems to be equivalent to reachability; i.e., their complexities are equivalent to that of reachability with respect to PTIME many-one reductions. (Throughout this paper, we

*This work was supported in part by U. S. Office of Naval Research Grant No. N00014-86-K-0763.

will use the word “equivalent” in this sense when referring to decision problems.) Thus, to determine the precise complexity of any of these problems (modulo PTIME reductions), it is sufficient to consider only the reachability problem. For examples of other (more classical) problems known to be equivalent to reachability, see [26].

Clearly, since some versions of the fair nontermination problem have been shown to be undecidable [6,7], not all versions are equivalent to reachability. Rather, one of the main points of this paper is that most versions are either highly undecidable (in particular, complete for Σ_1^1 — the first level of the analytical hierarchy), equivalent to reachability, or equivalent to boundedness (i.e., exponential space complete with respect to PTIME many-one reductions). In some sense, it is easier to show problems to be highly undecidable or equivalent to boundedness than it is to show equivalence to reachability. The reason for this is that highly undecidable problems and exponential space complete problems as a whole are fairly well understood. The reachability problem, however is not well understood, as is evidenced by the lack of knowledge concerning its complexity. Hence, as a way to overcome this difficulty, we develop in this paper a framework based upon temporal logic for reducing fair nontermination problems to the reachability problem.

For some time, temporal logic has been considered an appropriate formalism for reasoning about systems of concurrent programs [22,27]. A typical problem involving temporal logic is the model checking problem; i.e., determining whether a given structure defines a model of a correctness specification expressed in the temporal logic. The problem can be formally stated in a variety of ways. For the purposes of this paper, model checking is the problem of deciding whether in a given Petri net there is an infinite firing sequence satisfying a given temporal logic formula. This version of model checking (with respect to finite-state structures) was referred to as “determination of truth in a structure” in [31]. What makes the model checking problem useful for us is that most fairness specifications can be stated in some temporal logic (see, e.g., [10]). Thus, a fair nontermination problem can be reduced to a model checking problem. Now in order for this reduction to be useful (for our purposes), the model checking problem must be no harder than reachability. Unfortunately, we were able to show in [15] that for a fairly simple temporal logic, the model checking problem is undecidable, even for conflict-free Petri nets (see also [33]). On the other hand, we were able to show a subset of the logic to be NP-complete for conflict-free nets. This logic utilizes the predicates $ge(p,c)$ (place p is greater than or equal to c), $en(t)$ (transition t is enabled) and $fi(t)$ (transition t is the next to fire). Call this set of predicates Q' . The operators used are F (sometime), X (next time), \wedge (and), \vee (or), and \neg (not), where \neg is used only on predicates. The problem with this logic (which we will call $\tilde{\mathcal{L}}(Q',F,X)$) is that it can only make assertions about finite portions of (possibly infinite) firing sequences. Hence, it is only marginally useful for expressing fairness constraints.

With respect to general Petri nets, it is not hard to see that $\tilde{\mathcal{L}}(Q',F,X)$ is powerful enough to express reachability. Thus, the model checking problem for this logic is as hard as reachability. In this paper, we extend this logic by adding new predicates which give it the power to specify certain loops. This strategy differs from our analysis in [15], where we kept the set of predicates fixed but restricted the use of the operators in order to develop various logics. With this new logic (which we will call $\tilde{\mathcal{L}}(Q,F,X)$) we associate a variation of the model checking problem which we call the finite model checking problem. This problem is to determine whether there exists a *finite* firing sequence that satisfies a given formula. The reason we introduce the finite model checking problem is to allow us to use predicates asserting that a firing sequence produces a nonnegative (or zero) net change on a given place; this assertion would not make sense for arbitrary infinite

firing sequences. We are then able to show that the finite model checking problem for $\tilde{\mathcal{L}}(Q, F, X)$ is equivalent to reachability. In so doing, we develop a methodology for admitting new predicates to the logic without destroying its equivalence to reachability. Because $\tilde{\mathcal{L}}(Q, F, X)$ has the power to specify loops, it can specify that certain types of events occur infinitely often. In particular, suppose we wish to determine whether a Petri net \mathcal{P} is a model for a formula $f \in \tilde{\mathcal{L}}(Q', F, X)$. We need only to find a finite firing sequence that satisfies f and ends in a loop. We can specify a path of this type in $\tilde{\mathcal{L}}(Q, F, X)$; hence, the model checking problem for $\tilde{\mathcal{L}}(Q', F, X)$ may be expressed as an instance of the finite model checking problem for $\tilde{\mathcal{L}}(Q, F, X)$. Furthermore, $\tilde{\mathcal{L}}(Q, F, X)$ is powerful enough to express formulas of the form “infinitely often q ,” where q is a Boolean combination of predicates from Q' . (Call this set of formulas $\mathcal{L}^\infty(Q')$.) Again, we can express the model checking problem for $\mathcal{L}^\infty(Q')$ as an instance of the finite model checking problem for $\tilde{\mathcal{L}}(Q, F, X)$. This answers a question left open in [15] as to whether the model checking problem for $\mathcal{L}^\infty(Q')$ is decidable for conflict-free Petri nets. Perhaps the most important property of $\tilde{\mathcal{L}}(Q, F, X)$ is its ability to succinctly express various definitions of fairness — namely, ∞ -fairness [5], fdp-T where T is a set of transitions [7], and five types of fairness introduced by Landweber [19] and Carstensen and Valk [8]. As a result, we are able to show the fair nontermination problems for all of these types of fairness to be equivalent to reachability. Thus, $\tilde{\mathcal{L}}(Q, F, X)$ provides a powerful mechanism for showing equivalence of certain fair nontermination problems to reachability. Finally, we might also mention that the existence of such a logic is somewhat surprising in view of the aforementioned negative results of [15].

Since so many of the decidable versions of the fair nontermination problem are equivalent to more classical Petri net problems, one might surmise that the undecidable versions are equivalent to containment and equivalence, which have been shown to be undecidable in [4] and [13], respectively. Since reachability is decidable, it is not hard to see that both of these problems are in Π_1 ; i.e., they are co-r.e. As Vidal-Naquet [6] and Carstensen [7] have demonstrated, the key step in showing that any version is undecidable is to show that the Petri net under the given fairness constraint can simulate zero-testing. Once zero-testing can be performed, the Petri net can simulate an arbitrary Turing machine. Carstensen [7] has also exhibited another interesting capability of Petri nets under certain fairness constraints: the ability to nondeterministically generate any natural number without running the risk of entering an infinite loop. Thus, unbounded nondeterminism can be simulated. We use this fact to show that these versions are complete for Σ_1^1 . Hence, these problems are highly undecidable and not equivalent to containment or equivalence. These results may be compared and contrasted with those of [1,2,3,9,14,25,32].

The remainder of the paper is organized as follows. In Section 2, we give the basic definitions of Petri nets and temporal logic. In Section 3, we develop the logics that are later used to show various types of fair nontermination problems to be equivalent to reachability. In Section 4, we examine the fair nontermination problem for the 24 types of fairness. In most cases, we are able to show the problems to be either highly undecidable, equivalent to reachability, or equivalent to boundedness. One exception is with respect to fairness (as defined by Lehman, Pnueli, and Stavi [20]) for bounded Petri nets, as is mentioned in [7]. Although this problem is clearly decidable, we are able to show it to be nonprimitive recursive. We conclude in Section 5 with a summary of our results and a discussion of open problems.

2 Definitions

A *Petri Net* (PN, for short) is a tuple (P, T, φ, μ_0) , where P is a finite set of *places*, T is a finite set of *transitions*, φ is a *flow function* $\varphi : (P \times T) \cup (T \times P) \rightarrow \mathbb{N}$, and μ_0 is the *initial marking* $\mu_0 : P \rightarrow \mathbb{N}$, where \mathbb{N} is the set of natural numbers. A *marking* is a mapping $\mu : P \rightarrow \mathbb{N}$. We often establish an order on the places, p_1, \dots, p_k , and designate a marking μ as a vector in \mathbb{N}^k , where the i th component represents $\mu(p_i)$. A transition $t \in T$ is *enabled* at a marking μ iff for every $p \in P$, $\varphi(p, t) \leq \mu(p)$. A transition t may *fire* at a marking μ if t is enabled at μ . We then write $\mu \xrightarrow{t} \mu'$, where $\mu'(p) = \mu(p) - \varphi(p, t) + \varphi(t, p)$ for all $p \in P$. A sequence of transitions $\sigma = t_1 \dots t_n$ is a *firing sequence* from μ_0 iff $\mu_0 \xrightarrow{t_1} \mu_1 \xrightarrow{t_2} \dots \xrightarrow{t_n} \mu_n$ for some sequence of markings μ_1, \dots, μ_n . We also write $\mu_0 \xrightarrow{\sigma} \mu_n$, and denote $t_1 \dots t_j$ by $\sigma[j]$ for $1 \leq j \leq n$. We extend these notions to infinite firing sequences in the obvious way.

For a PN $\mathcal{P} = (P, T, \varphi, \mu_0)$, the *reachability set* of \mathcal{P} is the set $R(\mathcal{P}) = \{\mu \mid \mu_0 \xrightarrow{\sigma} \mu \text{ for some } \sigma\}$. Given a marking μ of \mathcal{P} , the *reachability problem* (RP) is to determine whether $\mu \in R(\mathcal{P})$. The *boundedness problem* (BP) is to determine whether $R(\mathcal{P})$ is finite. Throughout this paper, we will define several fairness properties for firing sequences. Given a fairness property x , the *nontermination problem with respect to x* (NTP x) is to determine whether there is an infinite firing sequence σ in \mathcal{P} that satisfies x .

A *labeled Petri net* is a triple $\mathcal{P} = (\mathcal{P}_1, \Sigma, h)$, where $\mathcal{P}_1 = (P, T, \varphi, \mu_0)$ is a PN, Σ is a finite set of labels, and $h : T \rightarrow \Sigma \cup \{\epsilon\}$ is a labeling function. We also extend $h : T^* \rightarrow \Sigma^*$ by $h(\epsilon) = \epsilon$ and $h(\sigma t) = h(\sigma)h(t)$. Given a marking μ of \mathcal{P} , we define the *terminal language of \mathcal{P} with respect to μ* as $L^t(\mathcal{P}, \mu) = \{h(\sigma) \mid \mu_0 \xrightarrow{\sigma} \mu\}$.

Let \mathcal{N} denote the set of all PNs, T^* denote all finite firing sequences of nets in \mathcal{N} , T^ω denote all infinite firing sequences of nets in \mathcal{N} , and $T^\infty = T^* \cup T^\omega$. A *predicate* is a partial function $q : \mathcal{N} \times T^\infty \times \mathbb{N} \rightarrow \{\text{true}, \text{false}\}$. For $\mathcal{P} \in \mathcal{N}$, σ a (finite or infinite) firing sequence of \mathcal{P} , and $n \in \mathbb{N}$, we say $\langle \mathcal{P}, \sigma, n \rangle \models q$ iff $q(\mathcal{P}, \sigma, n) = \text{true}$. A *well-formed formula* (wff) is either a predicate or of the form $\neg f$, $f \wedge g$, $\mathbf{X}f$, or $f \mathbf{U} g$, where f and g are wffs. For a firing sequence σ of \mathcal{P} and wffs f and g , we say:

- $\langle \mathcal{P}, \sigma, n \rangle \models \neg f$ iff not $(\langle \mathcal{P}, \sigma, n \rangle \models f)$;
- $\langle \mathcal{P}, \sigma, n \rangle \models \mathbf{X}f$ iff $\langle \mathcal{P}, \sigma, n+1 \rangle \models f$;
- $\langle \mathcal{P}, \sigma, n \rangle \models f \mathbf{U} g$ iff $\exists r > n$ such that $\langle \mathcal{P}, \sigma, r \rangle \models g$ and $\forall s, n < s < r, \langle \mathcal{P}, \sigma, s \rangle \models f$;
- $\langle \mathcal{P}, \sigma, n \rangle \models f \wedge g$ iff $\langle \mathcal{P}, \sigma, n \rangle \models f$ and $\langle \mathcal{P}, \sigma, n \rangle \models g$.

We also use the following abbreviations:

- $f \vee g \equiv \neg(\neg f \wedge \neg g)$;
- $f \supset g \equiv \neg f \vee g$;
- $\mathbf{F}f \equiv \text{true} \mathbf{U} f$;
- $\mathbf{G}f \equiv \neg \mathbf{F} \neg f$.

We say that \mathcal{P} is a (*finite*) *model* for f iff there is an infinite (finite, respectively) firing sequence σ in \mathcal{P} such that $\langle \mathcal{P}, \sigma, 0 \rangle \models f$. Let \mathcal{F} be a set of wffs. The (*finite*) *model checking problem with respect to \mathcal{F}* , denoted MCP(\mathcal{F}) (FMCP(\mathcal{F}), respectively), is to determine whether a given PN \mathcal{P} is a (finite) model for a given formula $f \in \mathcal{F}$. Let \mathcal{Q} be a set of predicates. We then define

- $\mathcal{L}(Q) = \{f \mid f \text{ is a wff using predicates from } Q\}$;
- $\tilde{\mathcal{L}}(Q, \mathbf{F}, \mathbf{X}) = \{f \mid f \text{ is a wff using predicates from } Q \text{ and the operators } \mathbf{F}, \mathbf{X}, \wedge, \vee, \text{ and } \neg, \text{ such that } \neg \text{ is used only on predicates}\}$; and
- $\mathcal{L}^\infty(Q) = \{\mathbf{G}\mathbf{F}f \mid f \text{ is a Boolean combination of predicates from } Q\}$.

For a PN $\mathcal{P}=(P, T, \varphi, \mu_0)$, a wff f , and a natural number n , we define the *model language of \mathcal{P} with respect to f and n* as $L^m(\mathcal{P}, f, n) = \{\sigma \mid \sigma \text{ is finite and } \langle \mathcal{P}, \sigma, n \rangle \models f\}$. A set \mathcal{F} of wffs is said to be *RP-decidable* iff for all $f \in \mathcal{F}, \mathcal{P} \in \mathcal{N}$, we can construct in PTIME a labeled PN $\mathcal{P}'=(\mathcal{P}'_1, \Sigma, h)$, $\mathcal{P}'_1=(P', T', \varphi', \mu'_0)$, $P'=p_1, \dots, p_k$, and a marking μ' on p_1, \dots, p_{k-1} such that for all $n \in \mathbb{N}$, $L^m(\mathcal{P}, f, n) = L^t(\mathcal{P}', (\mu', n))$. We will refer to p_k as the *marker*.

3 A temporal logic for Petri nets

In this section, we will present a temporal logic for reasoning about Petri nets such that if \mathcal{F} is the set of all wffs in the logic, then $\text{FMCP}(\mathcal{F}) \equiv_{\text{PTIME}} \text{RP}$. The logic will be $\tilde{\mathcal{L}}(Q, \mathbf{F}, \mathbf{X})$ for a set Q of predicates to be defined later. Even though $\text{FMCP}(\tilde{\mathcal{L}}(Q, \mathbf{F}, \mathbf{X}))$ is no harder than RP, Q will contain a sufficient variety of predicates to provide a powerful mechanism for showing fair nontermination problems to be equivalent to RP. Furthermore, certain restrictions of the logic provide interesting extensions to the results shown in [15]. The first extension we show is that $\text{MCP}(\tilde{\mathcal{L}}(Q', \mathbf{F}, \mathbf{X})) \equiv_{\text{PTIME}} \text{RP}$, where Q' is the set of predicates from the logic developed in [15]. The second extension is that $\text{MCP}(\mathcal{L}^\infty(Q')) \equiv_{\text{PTIME}} \text{RP}$. Both of these extensions may be considered refinements of the main result of this section — that $\text{FMCP}(\mathcal{L}(Q, \mathbf{F}, \mathbf{X})) \equiv_{\text{PTIME}} \text{RP}$. All of these logics will be used in the next section to show various fair nontermination problems to be equivalent to RP. The main result of this section may therefore be viewed as an umbrella under which a number of the subsequent results in this paper are derived.

In order to show that $\text{FMCP}(\tilde{\mathcal{L}}(Q, \mathbf{F}, \mathbf{X})) \leq_{\text{PTIME}} \text{RP}$, we will first show that the reduction holds for any RP-decidable set of predicates whose negations are also RP-decidable. We will then define Q and show that both Q and $\tilde{Q} = \{\neg q \mid q \in Q\}$ are RP-decidable. We first state the following lemma, which follows immediately from the definition of the terminal language of a labeled PN.

Lemma 3.1: Given a labeled PN $\mathcal{P}=(\mathcal{P}_1, \Sigma, h)$ and a marking μ , $L^t(\mathcal{P}, \mu) \neq \emptyset$ iff $\mu \in R(\mathcal{P}_1)$.

We can now give the following lemma, relating the FMCP to the RP.

Lemma 3.2: For a set of wffs \mathcal{F} , if \mathcal{F} is RP-decidable, then $\text{FMCP}(\mathcal{F}) \leq_{\text{PTIME}} \text{RP}$.

Proof. Given an RP-decidable set \mathcal{F} of wffs, let $f \in \mathcal{F}$, and let \mathcal{P} be an arbitrary PN. Since \mathcal{F} is RP-decidable, we can construct in PTIME a labeled PN $\mathcal{P}'=(\mathcal{P}'_1, \Sigma, h)$ and a marking μ such that $L^t(\mathcal{P}', \mu) = L^m(\mathcal{P}, f, 0)$. Then \mathcal{P} is a finite model for f iff $L^m(\mathcal{P}, f, 0) \neq \emptyset$ iff $L^t(\mathcal{P}', \mu) \neq \emptyset$ iff $\mu \in R(\mathcal{P}'_1)$. Therefore, $\text{FMCP}(\mathcal{F}) \leq_{\text{PTIME}} \text{RP}$. □

The following theorem now gives a framework for defining our set of predicates.

Theorem 3.1: If a set of predicates Q is RP-decidable and $\tilde{Q}=\{\neg q \mid q \in Q\}$ is RP-decidable, then $\text{FMCP}(\tilde{\mathcal{L}}(Q, \mathbf{F}, \mathbf{X})) \leq_{\text{PTIME}} \text{RP}$.

Proof. We will show by induction on the structure of $f \in \tilde{\mathcal{L}}(Q, \mathbf{F}, \mathbf{X})$ that for any PN \mathcal{P} , we can construct in

P_{TIME} a labeled PN \mathcal{P}' and a submarking μ such that for all $n \in \mathbb{N}$, $L^t(\mathcal{P}',(\mu,n)) = L^m(\mathcal{P},f,n)$. The theorem will then follow from Lemma 3.2. If $f \in Q \cup \tilde{Q}$, this is trivial. Therefore, assume we have $f_1, f_2 \in \tilde{\mathcal{L}}(Q, F, X)$ such that for any PN \mathcal{P} , we can construct in P_{TIME} labeled PNs $\mathcal{P}'_1, \mathcal{P}'_2$ and markings μ_1, μ_2 such that for all $n \in \mathbb{N}$, $L^t(\mathcal{P}'_1,(\mu_1,n)) = L^m(\mathcal{P},f_1,n)$ and $L^t(\mathcal{P}'_2,(\mu_2,n)) = L^m(\mathcal{P},f_2,n)$.

Case 1: $f = f_1 \wedge f_2$.

We construct \mathcal{P}' from \mathcal{P}'_1 and \mathcal{P}'_2 as follows. Let T_{1a} (T_{2a}) be the set of all transitions labeled a in \mathcal{P}'_1 (\mathcal{P}'_2 , respectively) for all $a \in \Sigma \cup \{\epsilon\}$, where $\Sigma = \{a_1, \dots, a_m\}$ is the alphabet of $L^m(\mathcal{P},f,n)$. We relabel all transitions in \mathcal{P}'_1 and \mathcal{P}'_2 as ϵ , and add $m+1$ new transitions and $2m+3$ new places, as is shown in Figure 3.1. We define μ as:

- $\mu(p_{ij}) = \mu_i(p_{ij})$ for $i = 1, 2, j = 1, \dots, k_i - 1$;
- $\mu(p_{ik_i}) = 0$ where p_{ik_i} is the marker for \mathcal{P}'_i , $i = 1, 2$;
- $\mu(\text{ack}_i) = 1$ for $i = 1, 2$;
- $\mu(\text{en}_{ij}) = 0$ for $i = 1, 2, j = 1, \dots, m$.

p_{mark} is the marker for \mathcal{P}' . It is not hard to see that for all $n \in \mathbb{N}$, $\sigma \in L^t(\mathcal{P}',(\mu,n))$ iff $\sigma \in L^t(\mathcal{P}'_1,(\mu_1,n)) \cap L^t(\mathcal{P}'_2,(\mu_2,n))$ iff $\sigma \in L^m(\mathcal{P},f_1,n) \cap L^m(\mathcal{P},f_2,n)$ iff $\langle \mathcal{P}, \sigma, n \rangle \models f_1 \wedge f_2$ iff $\sigma \in L^m(\mathcal{P},f,n)$.

Case 2: $f = f_1 \vee f_2$.

\mathcal{P}' nondeterministically chooses to simulate one of \mathcal{P}'_1 or \mathcal{P}'_2 ; the other, it brings immediately to its final marking with a zero marker value. \mathcal{P}' can also transfer all tokens from the markers of \mathcal{P}'_1 and \mathcal{P}'_2 to its marker. The details are left to the reader.

Case 3: $f = Xf_1$.

\mathcal{P}' simulates \mathcal{P}'_1 , but in order to reach its final marking, it must subtract 1 from the marker.

Case 4: $f = Ff_1$.

This is the same as Case 3 except that \mathcal{P}' can subtract more than 1 from the marker.

□

We are now ready to define our predicates. For a PN $\mathcal{P} = (P, T, \varphi, \mu_0)$, $p \in P$, $t \in T$, $c, n \in \mathbb{N}$, and a finite firing sequence σ , let:

- $\langle \mathcal{P}, \sigma, n \rangle \models \text{ge}(p, c)$ iff $\mu_0 \xrightarrow{\sigma[n]} \mu$ and $\mu(p) \geq c$;
- $\langle \mathcal{P}, \sigma, n \rangle \models \text{ft}(t)$ iff t is the $(n+1)$ st transition in σ ;
- $\langle \mathcal{P}, \sigma, n \rangle \models \text{lp}(p)$ iff $\mu_0 \xrightarrow{\sigma[n]} \mu_1, \mu_0 \xrightarrow{\sigma} \mu_2$, and $\mu_1(p) \leq \mu_2(p)$;
- $\langle \mathcal{P}, \sigma, n \rangle \models \text{zl}(p)$ iff $\mu_0 \xrightarrow{\sigma[n]} \mu_1, \mu_0 \xrightarrow{\sigma} \mu_2$, and $\mu_1(p) = \mu_2(p)$;
- $\langle \mathcal{P}, \sigma, n \rangle \models \text{co}(\mu)$ iff $\mu_0 \xrightarrow{\sigma[n]} \mu_1$ and there is a $\mu_2 \in R(P, T, \varphi, \mu_1)$ such that $\mu_2 \geq \mu$.

Let Q be the set of all of the above predicates for all places p and transitions t , and $\tilde{Q} = \{q \mid \neg q \in Q\}$. We wish to show that $\text{FMCP}(\tilde{\mathcal{L}}(Q, F, X)) \leq_{P\text{TIME}} \text{RP}$. From Theorem 3.1, we need only show that Q and \tilde{Q} are RP-decidable. For most elements of $Q \cup \tilde{Q}$, this is straightforward. The main difficulty lies with

$\neg\text{co}(\mu)$. In showing $\neg\text{co}(\mu)$ to be RP-decidable, we will construct a PN that will in some sense produce all markings from which no marking greater than or equal to μ can be reached. In order to construct such a PN, we first construct a modified Turing machine (TM) that accomplishes the same purpose. This modified TM will be such that using Lipton's construction [21], we can transform it to a PN. In order to define the precise function of the modified TM, we introduce the following notation. Given an unmarked PN $\mathcal{P} = (P, T, \varphi)$, a marking μ , and a positive c , let $S(\mathcal{P}, \mu) = \{\mu_0 \mid \forall \mu' \geq \mu, \mu' \notin R(P, T, \varphi, \mu_0)\}$, and $S'(\mathcal{P}, \mu, c) = \{\mu_0 \#^{2^{cn \lceil \log n \rceil}} \mid \mu_0 \in S(\mathcal{P}, \mu)\}$, where n is the number of bits needed to encode \mathcal{P} and μ . Also, let $\bar{S}(\mathcal{P}, \mu)$ and $\bar{S}'(\mathcal{P}, \mu, c)$ denote the complements of $S(\mathcal{P}, \mu)$ and $S'(\mathcal{P}, \mu, c)$, respectively. We now give two lemmas which define the function of the modified TM and give its construction.

Lemma 3.3: There is a positive constant c such that for any unmarked PN $\mathcal{P} = (P, T, \varphi)$ and marking μ , we can construct in PTIME a linearly bounded automaton (LBA) M that accepts the set $S'(\mathcal{P}, \mu, c)$.

Proof. For any fixed constant c , we can clearly verify deterministically in linear space whether the input is syntactically correct and the number of $\#$ s is correct. From [29], if $\mu_0 \xrightarrow{\sigma_1} \mu_1 \geq \mu$, then $\mu_0 \xrightarrow{\sigma_2} \mu_2 \geq \mu$ such that the length of σ_2 is no more than $2^{2^{dn \log n}}$ for some constant d . Thus, any marking produced in this firing sequence can be stored in $\text{size}(\mu_0) + 2^{cn \log n}$ bits for some constant c . $\bar{S}'(\mathcal{P}, \mu, c)$ is therefore accepted by an LBA M' that can be constructed in PTIME. From [17], there is an LBA M that accepts $S'(\mathcal{P}, \mu, c)$. An inspection of the proof in [17] reveals that M can be constructed from M' in PTIME. □

Lemma 3.4: Let \mathcal{P} , μ , and n be as in Lemma 3.3, and let k be the number of places in \mathcal{P} . We can construct in PTIME an $O(2^{an \log n})$ space bounded TM M with no input tape, augmented with k unbounded increment-only counters such that the set of final counter values in all accepting computations is $S(\mathcal{P}, \mu)$.

Proof. Let d be the constant from [29] mentioned in the proof of Lemma 3.3. We define a function g mapping markings to markings such that

$$\begin{aligned} g(\mu)(p) &= \mu(p) \text{ if } \mu(p) \leq 2^{2^{dn \lceil \log n \rceil}} \\ g(\mu)(p) &= 2^{2^{dn \lceil \log n \rceil}} \text{ otherwise.} \end{aligned}$$

It follows from [29] that $\mu_0 \in S(\mathcal{P}, \mu)$ iff $g(\mu_0) \in S(\mathcal{P}, \mu)$. We can therefore construct M to operate as follows. M first nondeterministically generates k nonnegative integers no larger than $2^{2^{dn \lceil \log n \rceil}}$ and stores them on its worktape and on the k counters. If any of these values are exactly $2^{2^{dn \lceil \log n \rceil}}$, the corresponding counters may be incremented arbitrarily many times. M then writes $\#^{2^{cn \lceil \log n \rceil}}$ on its worktape, where c is the constant from Lemma 3.3, and simulates the machine given by Lemma 3.3 on the contents of its worktape. The result follows from Lemma 3.3. □

We are now ready to show the main result of this section, that $\text{FMCP}(\tilde{\mathcal{L}}(Q, F, X))$ is equivalent to RP. From this result we will subsequently derive two refinements concerning logics developed in [15]; these refinements will be given in Theorems 3.3 and 3.4. We will then use Theorem 3.2 and its refinements in Section 4 to show seven fair nontermination problems to be equivalent to RP. The reason we can use a finite model checking problem to encode a fair nontermination problem is that $\tilde{\mathcal{L}}(Q, F, X)$ has the power to

express certain loops which may be iterated to produce an infinite “fair” path. Thus, Theorem 3.2 is an umbrella under which powerful machinery is developed for proving certain fair nontermination problems to be equivalent to RP.

Theorem 3.2: $\text{FMCP}(\tilde{\mathcal{L}}(\mathcal{Q}, \mathbf{F}, \mathbf{X})) \equiv_{PTIME} \text{RP}$.

Proof. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be an arbitrary PN, and let μ be an arbitrary marking of \mathcal{P} . Clearly, $\mu \in R(\mathcal{P})$ iff there is a finite firing sequence σ such that

$$\langle \mathcal{P}, \sigma, 0 \rangle \models [\bigwedge_{p \in P} (\text{ge}(p, \mu(p)) \wedge \neg \text{ge}(p, \mu(p)+1))] \vee \mathbf{F} [\bigwedge_{p \in P} (\text{ge}(p, \mu(p)) \wedge \neg \text{ge}(p, \mu(p)+1))].$$

Thus, $\text{RP} \leq_{PTIME} \text{FMCP}(\tilde{\mathcal{L}}(\mathcal{Q}, \mathbf{F}, \mathbf{X}))$.

In order to show that $\text{FMCP}(\tilde{\mathcal{L}}(\mathcal{Q}, \mathbf{F}, \mathbf{X})) \leq_{PTIME} \text{RP}$, we need only to show that \mathcal{Q} and $\tilde{\mathcal{Q}}$ are RP-decidable; the theorem then follows from Theorem 3.1. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be an arbitrary PN. We need to show that for any element q of $\mathcal{Q} \cup \tilde{\mathcal{Q}}$, we can construct in PTIME a labeled PN \mathcal{P}' and a marking μ' such that for any $n \in \mathbb{N}$, $L^m(\mathcal{P}, q, n) = L^t(\mathcal{P}', (\mu', n))$. Since most of the cases are straightforward, we will only show the cases $\text{ge}(p_i, c)$ and $\neg \text{co}(\mu)$.

Case 1: $q = \text{ge}(p_i, c)$.

We construct \mathcal{P}' as shown in Figure 3.2. Let $\mu'(en_2) = 1$, $\mu'(p) = 0$ if $p \neq en_2$ and $p \neq p_{\text{mark}}$, where p_{mark} is the marker. In order to reach (μ', n) , \mathcal{P}' must pass through two phases. In the first phase, \mathcal{P}' simulates n transitions of a firing sequence σ on two copies of \mathcal{P} . In the second phase, the remainder of σ is simulated on one of the copies. At any time, the transitions $t_{z_1}, \dots, t_{z_k}, t'_{z_1}, \dots, t'_{z_k}$ may fire. These transitions enable no new transitions, but allow all places in both copies of \mathcal{P} to be brought to 0 after the simulation of σ . Furthermore, t_{g_e} must fire exactly once, guaranteeing that p_i is at least c after $\sigma[n]$ is executed in \mathcal{P} . It can therefore be seen that $L^m(\mathcal{P}, q, n) = L^t(\mathcal{P}', (\mu', n))$.

Case 2: $q = \neg \text{co}(\mu)$.

Before we construct \mathcal{P}' , we will construct a PN \mathcal{P}'' with designated places p''_1, \dots, p''_k , where k is the number of places in \mathcal{P} , and a submarking μ'' on the remaining places in \mathcal{P}'' , such that for any marking μ_1 of \mathcal{P} , $(\mu_1, \mu'') \in R(\mathcal{P}'')$ iff $\mu_1 \in S((P, T, \varphi), \mu)$. Let M be the machine given by Lemma 3.4 to compute $S((P, T, \varphi), \mu)$. Since the worktape of M is bounded by $O(2^{an \log n})$ space, it can be simulated by four $O(2^{2^{an \log n}})$ bounded counters; furthermore, this construction can be done in PTIME. From [21], we can construct in PTIME a PN and a marking μ'' such that any firing sequence yielding μ'' simulates an accepting computation of M without its k unbounded counters. Since the k unbounded counters are increment-only, they can clearly be implemented by k places in the PN. We have therefore constructed \mathcal{P}'' and μ'' . The construction of \mathcal{P}' and μ' from \mathcal{P} , \mathcal{P}'' , and μ'' is now similar to Case 1 and is shown in Figure 3.3. It should be clear that $L^m(\mathcal{P}, q, n) = L^t(\mathcal{P}', (\mu', n))$. □

We now will examine two logics developed in [15]. We can show that the MCP for both problems can be expressed as restrictions of $\text{FMCP}(\tilde{\mathcal{L}}(\mathcal{Q}, \mathbf{F}, \mathbf{X}))$. We will use these two logics in Theorems 4.2 and 4.6 to give succinct reductions from various fair nontermination problems to RP.

Let \mathcal{Q}' be the set of predicates $\text{ge}(p, c)$ and $\text{fi}(t)$ extended to infinite firing sequences, and let $\tilde{\mathcal{Q}}' = \{\neg q \mid q \in \mathcal{Q}'\}$. $\text{MCP}(\tilde{\mathcal{L}}(\mathcal{Q}', \mathbf{F}, \mathbf{X}))$ was shown in [15] to be NP-complete for conflict-free PNs. (The logic in [15] also included predicates asserting that a transition t is enabled; this assertion and its negation can clearly

be encoded in $\tilde{\mathcal{L}}(Q', \mathbf{F}, \mathbf{X})$.) Although $\tilde{\mathcal{L}}(Q', \mathbf{F}, \mathbf{X})$ can only express loops in which the repeated markings are explicitly stated, this is sufficient to encode several of the types of fairness given by Landweber [19] and Carstensen and Valk [8]. We now show that $\text{MCP}(\tilde{\mathcal{L}}(Q', \mathbf{F}, \mathbf{X}))$ for general PNs is equivalent to reachability.

Theorem 3.3: $\text{MCP}(\tilde{\mathcal{L}}(Q', \mathbf{F}, \mathbf{X})) \equiv_{\text{PTIME}} \text{RP}$.

Proof. We will first show $\text{RP} \leq_{\text{PTIME}} \text{MCP}(\tilde{\mathcal{L}}(Q', \mathbf{F}, \mathbf{X}))$. Let \mathcal{P} be an arbitrary PN and μ be an arbitrary marking of \mathcal{P} . We construct a PN \mathcal{P}' from \mathcal{P} by adding a transition t that does nothing and is always enabled. Clearly, $\mu \in \text{R}(\mathcal{P}')$ iff $\mu \in \text{R}(\mathcal{P})$, and any finite firing sequence in \mathcal{P}' can be made infinite. Since we can express reachability in $\tilde{\mathcal{L}}(Q', \mathbf{F}, \mathbf{X})$ as in the proof of Theorem 3.2, $\text{RP} \leq_{\text{PTIME}} \text{MCP}(\tilde{\mathcal{L}}(Q', \mathbf{F}, \mathbf{X}))$.

We will now show $\text{MCP}(\tilde{\mathcal{L}}(Q', \mathbf{F}, \mathbf{X})) \leq_{\text{PTIME}} \text{RP}$. We will reduce the MCP to FMCP($\tilde{\mathcal{L}}(Q', \mathbf{F}, \mathbf{X})$); the result will then follow from Theorem 3.2. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be an arbitrary PN, and let f be an arbitrary wff in $\tilde{\mathcal{L}}(Q', \mathbf{F}, \mathbf{X})$. Let $f' \equiv f \wedge \mathbf{F} \bigwedge_{p \in P} \text{lp}(p)$. It is a straightforward matter to show that \mathcal{P} is a model for f iff \mathcal{P} is a finite model for f' . □

One question left open in [15] was whether $\text{MCP}(\mathcal{L}^\infty(Q'))$ is decidable for conflict-free PNs. In Theorem 3.4, we will give a positive answer to this question by showing the problem with respect to general PNs to be equivalent to RP. Again, Theorem 3.4 may be viewed as a refinement of Theorem 3.2. Clearly, $\mathcal{L}^\infty(Q')$ can be used to express fairness constraints for which a certain event must be repeated infinitely often. We now define the following terminology. A set of wffs \mathcal{F} is *Q-pumpable* iff there is a PTIME function $g : \mathcal{F} \rightarrow \tilde{\mathcal{L}}(Q, \mathbf{F}, \mathbf{X})$ such that

1. for any $f \in \mathcal{F}$, and any infinite firing sequence σ in \mathcal{P} , if $\langle \mathcal{P}, \sigma, 0 \rangle \models \text{Gff}$, then for any infinite set I of natural numbers such that if $i \in I$, then $\langle \mathcal{P}, \sigma, i \rangle \models f$, there is an infinite subset I' of I such that for any $i, j \in I'$, $i < j$, $\langle \mathcal{P}, \sigma[j], i \rangle \models g(f)$; and
2. if $\mathcal{P} = (P, T, \varphi, \mu_0)$ is such that $\langle \mathcal{P}, \sigma, n \rangle \models g(f)$ for some finite $\sigma = \sigma[n]\sigma'$, then $\mu_0 \xrightarrow{\sigma[n]} \mu_1 \xrightarrow{\sigma'} \mu_2 \xrightarrow{\sigma'} \dots$ such that $\langle (P, T, \varphi, \mu_j), \sigma'[1], 0 \rangle \models f$ for all $j \geq 1$.

Before showing Theorem 3.4, we give the following lemma.

Lemma 3.5: The set of all positive Boolean combinations of predicates in $Q' \cup \tilde{Q}'$ is Q-pumpable.

Proof. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be an arbitrary PN. We will show by induction on the structure of f that we can construct in PTIME a $g(f)$ satisfying conditions (1) and (2) above.

Base Case 1: $f = \text{ge}(p, c)$.

Let $g(f) = \text{ge}(p, c) \wedge \bigwedge_{p' \in P} \text{lp}(p')$. Let σ be any infinite firing sequence such that $\langle \mathcal{P}, \sigma, 0 \rangle \models \text{Gff}$, and let I be any infinite set of natural numbers such that if $i \in I$, then $\langle \mathcal{P}, \sigma, i \rangle \models f$. Let μ_i be such that $\mu_0 \xrightarrow{\sigma[i]} \mu_i$, and let θ be the sequence $\mu_{i_1}, \mu_{i_2}, \dots$ such that $\{i_1, i_2, \dots\} = I$ and $i_1 \leq i_2 \leq \dots$. Since θ is infinite, there must be an infinite subsequence of θ , $\theta' = \mu_{j_1}, \mu_{j_2}, \dots$ such that $\mu_{j_1} \leq \mu_{j_2} \leq \dots$. Clearly, for any natural numbers r and s such that $r < s$, $\langle \mathcal{P}, \sigma[j_s]_{j_r}, 0 \rangle \models g(f)$. Part (2) clearly holds.

Base Case 2: $f = \neg \text{ge}(p, c)$.

$g(f) = \neg \text{ge}(p, c) \wedge \text{zl}(p) \wedge \bigwedge_{p' \in P} \text{lp}(p')$ clearly works.

Base Case 3: $f = \text{fi}(t)$.

$g(f) = \text{fi}(t) \wedge \bigwedge_{p' \in P} \text{lp}(p')$ clearly works.

Base Case 4: $f = \neg \text{fi}(t)$.

$g(f) = \neg \text{fi}(t) \wedge \bigwedge_{p' \in P} \text{lp}(p')$ clearly works.

Now assume we have f_1 and f_2 for which conditions (1) and (2) hold.

Case 5: $f = f_1 \vee f_2$.

Let $g(f) = g(f_1) \vee g(f_2)$. Let σ be any infinite firing sequence such that $\langle \mathcal{P}, \sigma, 0 \rangle \models \text{GF}f$, and let I be any infinite set of natural numbers i such that $\langle \mathcal{P}, \sigma, i \rangle \models f$. There must be some infinite subset I' of I such that either $\forall i \in I' \langle \mathcal{P}, \sigma, i \rangle \models f_1$, or $\forall i \in I' \langle \mathcal{P}, \sigma, i \rangle \models f_2$. Assume without loss of generality that $\forall i \in I' \langle \mathcal{P}, \sigma, i \rangle \models f_1$. From the induction hypothesis, there is an infinite subset I'' of I' such that for any $i, j \in I''$, $i < j$, $\langle \mathcal{P}, \sigma[j], i \rangle \models g(f_1)$. Then $\langle \mathcal{P}, \sigma[j], i \rangle \models g(f)$.

Now suppose $\mathcal{P} = (P, T, \varphi, \mu_0)$ such that $\langle \mathcal{P}, \sigma, n \rangle \models g(f)$ for some finite $\sigma = \sigma[n]\sigma'$. Then either $\langle \mathcal{P}, \sigma, n \rangle \models g(f_1)$ or $\langle \mathcal{P}, \sigma, n \rangle \models g(f_2)$. Without loss of generality, assume the former. From the induction hypothesis, $\mu_0 \xrightarrow{\sigma[n]} \mu_1 \xrightarrow{\sigma'} \mu_2 \xrightarrow{\sigma'} \dots$ such that $\langle (P, T, \varphi, \mu_j), \sigma[1], 0 \rangle \models f_1$ for all $j \geq 1$. Therefore, $\langle (P, T, \varphi, \mu_j), \sigma[1], 0 \rangle \models f_1 \vee f_2 = f$.

Case 6: $f = f_1 \wedge f_2$.

Let $g(f) = g(f_1) \wedge g(f_2)$. Let σ be any infinite firing sequence such that $\langle \mathcal{P}, \sigma, 0 \rangle \models \text{GF}f$, and let I be any infinite set of natural numbers i such that $\langle \mathcal{P}, \sigma, i \rangle \models f$. Then for any $i \in I$, $\langle \mathcal{P}, \sigma, i \rangle \models f_1$. From the induction hypothesis, there is an infinite subset I' of I such that for any $i, j \in I'$, $i < j$, $\langle \mathcal{P}, \sigma[j], i \rangle \models g(f_1)$. Since $I' \subseteq I$, for any $i \in I'$, $\langle \mathcal{P}, \sigma, i \rangle \models f_2$. From the induction hypothesis, there is an infinite subset I'' of I' such that for any $i, j \in I''$, $i < j$, $\langle \mathcal{P}, \sigma[j], i \rangle \models g(f_2)$. Consequently, $\langle \mathcal{P}, \sigma[j], i \rangle \models g(f)$.

Now suppose $\mathcal{P} = (P, T, \varphi, \mu_0)$ such that $\langle \mathcal{P}, \sigma, n \rangle \models g(f)$ for some finite $\sigma = \sigma[n]\sigma'$. Then $\langle \mathcal{P}, \sigma, n \rangle \models g(f_1)$ and $\langle \mathcal{P}, \sigma, n \rangle \models g(f_2)$. From the induction hypothesis, $\mu_0 \xrightarrow{\sigma[n]} \mu_1 \xrightarrow{\sigma'} \mu_2 \xrightarrow{\sigma'} \dots$ such that $\langle (P, T, \varphi, \mu_j), \sigma'[1], 0 \rangle \models f_1$ and $\langle (P, T, \varphi, \mu_j), \sigma'[1], 0 \rangle \models f_2$ for all $j \geq 1$. Therefore, $\langle (P, T, \varphi, \mu_j), \sigma'[1], 0 \rangle \models f_1 \wedge f_2 = f$.

□

Theorem 3.4: $\text{MCP}(\mathcal{L}^\infty(Q')) \equiv_{\text{PTIME}} \text{RP}$.

Proof. Let $\text{GF}f \in \mathcal{L}^\infty(Q')$. We can construct in PTIME an $f' \equiv f$ such that negations occur only on predicates in f' . Let g be the PTIME function constructed in Lemma 3.5. It now follows immediately from the properties of g that \mathcal{P} is a model for f iff \mathcal{P} is a finite model for $g(f')$. Thus, $\text{MCP}(\mathcal{F}) \leq_{\text{PTIME}} \text{RP}$. The construction given in Theorem 3.3 shows $\text{RP} \leq_{\text{PTIME}} \text{MCP}(\mathcal{L}^\infty(Q'))$.

□

4 The fair nontermination problem

In this section, we examine the complexities of 24 fair nontermination problems. We use the machinery developed in Section 3 to prove Theorems 4.2, 4.6, and 4.12, where a total of seven of these problems are shown to be equivalent to RP. Particularly in the latter two theorems, this machinery provides for very succinct reductions to RP, whereas “brute force” reductions are much longer and considerably more tedious. In the remainder of the theorems in this section, most of the problems we study are shown to be either equivalent to BP or Σ_1^1 -complete. The results of this section are summarized in a table in Section 5.

The first notions of fairness we consider were defined in Landweber [19], and Carstensen and Valk [8]. These definitions of fairness are such that a fair firing sequence must visit certain *predefined* markings or transitions infinitely often. It is worth mentioning that the notion of “enabledness” does not play any role in these definitions (other than the fact that the definitions deal with firing sequences). Given an infinite firing sequence $\sigma = t_1 t_2 \dots$, we define $\text{inf}^M(\sigma)$ ($\text{inf}^T(\sigma)$) to be the set of markings (transitions) that occur infinitely often in σ (i.e., $\text{inf}^M(\sigma) = \{\mu \mid \text{there are infinitely many } i \text{ such that } \mu_0 \xrightarrow{\sigma^{[i]}} \mu\}$ and $\text{inf}^T(\sigma) = \{t_i \mid t_i \text{ occurs infinitely often in } \sigma\}$). Let \mathcal{A} be a finite set of finite nonempty sets of markings. An infinite firing sequence $\sigma = t_1 t_2 \dots$ is said to be

- M1-fair iff $\exists A \in \mathcal{A}, \exists i \in \mathbb{N} : \mu_0 \xrightarrow{\sigma^{[i]}} \mu \in A$ (i.e., some marking reached by σ is in A).
- M1'-fair iff $\exists A \in \mathcal{A}, \forall i \in \mathbb{N} : \mu_0 \xrightarrow{\sigma^{[i]}} \mu_i \in A$ (i.e., every marking reached by σ is in A).
- M2-fair iff $\exists A \in \mathcal{A}, \text{inf}^M(\sigma) \cap A \neq \emptyset$ (i.e., some marking reached infinitely often by σ is in A).
- M2'-fair iff $\exists A \in \mathcal{A}, \text{inf}^M(\sigma) \neq \emptyset$ and $\text{inf}^M(\sigma) \subseteq A$ (i.e., σ reaches some marking infinitely often and every marking reached infinitely often by σ is in A).
- M3-fair iff $\exists A \in \mathcal{A}, \text{inf}^M(\sigma) = A$ (i.e., the set of markings reached infinitely often by σ is an element of \mathcal{A}).
- M3'-fair iff $\exists A \in \mathcal{A}, A \subseteq \text{inf}^M(\sigma)$ (i.e., every marking in A is reached infinitely often by σ).

Similarly, let \mathcal{A} be a finite set of nonempty subsets of transitions. σ is said to be:

- T1-fair iff $\exists A \in \mathcal{A}, \exists i \in \mathbb{N}^+ : t_i \in A$.
- T1'-fair iff $\exists A \in \mathcal{A}, \forall i \in \mathbb{N}^+ : t_i \in A$.
- T2-fair iff $\exists A \in \mathcal{A}, \text{inf}^T(\sigma) \cap A \neq \emptyset$.
- T2'-fair iff $\exists A \in \mathcal{A}, \text{inf}^T(\sigma) \subseteq A$.
- T3-fair iff $\exists A \in \mathcal{A}, \text{inf}^T(\sigma) = A$.
- T3'-fair iff $\exists A \in \mathcal{A}, A \subseteq \text{inf}^T(\sigma)$.

We will now investigate the complexity of the nontermination problems with respect to the above fairness constraints. The first theorem gives an exception to our general classification scheme due to the fact that for M1'-fairness the entire allowable reachability set is given as input to the problem.

Theorem 4.1: $\text{NTP}^{M1'}$ is NLOGSPACE-complete.

Proof. The proof method here is similar to the proof used in [15] to solve a similar problem for conflict-free Petri nets. For the sake of completeness, we sketch the proof in what follows. To show the problem to be in NLOGSPACE, we first guess an A in \mathcal{A} . Next, we check to make sure that the initial marking is in A . During the entire procedure, we will maintain a pointer to keep track of the current marking. (Note here that the amount of space needed to maintain such a pointer is logarithmic in the size of the input.) Initially, the pointer points to the initial marking. We then repeatedly guess a transition t and a marking μ in A and verify that μ can be produced by firing t in the current marking. If so, μ becomes the current marking. If the

above procedure can be performed for more than $|A|$ times, then there is an $M1'$ -fair computation. Clearly, the above procedure is in NLOGSPACE. The problem was shown to be NLOGSPACE-hard for conflict-free PNs in [15]; thus, it must also be NLOGSPACE-hard for general PNs.

□

In showing the following theorem, we make use of the temporal logic results given in Theorems 3.3 and 3.4. This machinery makes the proof very succinct; however, this particular theorem is not exceedingly difficult to prove without using these results. The real power of our logic will be exploited in Theorems 4.6 and 4.12.

Theorem 4.2: $NTP^x \equiv_{PTIME} RP$, for $x \in \{M1, M2, M2', M3, M3'\}$.

Proof. We first show $RP \leq_{PTIME} NTP^x$. We use a technique from [15]. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be an arbitrary PN and μ be an arbitrary marking of \mathcal{P} . We now construct a Petri net \mathcal{P}' and a set \mathcal{A} such that $\mu \in R(\mathcal{P})$ iff there is an x -fair computation with respect to \mathcal{A} in \mathcal{P}' . The new Petri net \mathcal{P}' is identical to \mathcal{P} except that it has an additional transition that is always enabled and does nothing. Now let $\mathcal{A} = \{\{\mu\}\}$. Clearly, $\mu \in R(\mathcal{P})$ iff \mathcal{P}' has an x -fair computation with respect to \mathcal{A} , where $x \in \{M1, M2, M2', M3, M3'\}$. Furthermore, the reduction can be done in polynomial time.

We will now show $NTP^x \leq_{PTIME} RP$ for $x \in \{M1, M2, M2', M3, M3'\}$. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be an arbitrary PN, and let \mathcal{A} be a finite set of nonempty finite sets of markings. For each element A of \mathcal{A} , let $A = \{\mu_{A1}, \mu_{A2}, \dots, \mu_{Ak_A}\}$. We will give for each x a wff f_x such that \mathcal{P} has an infinite x -fair path iff \mathcal{P} is a model for f_x ; furthermore, f_x will be in either $\tilde{\mathcal{L}}(Q', F, X)$ or $\mathcal{L}^\infty(Q')$. The theorem will then follow from Theorems 3.3 and 3.4. We define:

- $eq(\mu) \equiv \bigwedge_{p \in P} [ge(p, \mu(p)) \wedge \neg ge(p, \mu(p)+1)]$;
- $in(A) \equiv \bigvee_{\mu \in A} eq(\mu)$;
- $f_{M1} \equiv \bigvee_{A \in \mathcal{A}} (in(A) \vee F in(A))$;
- $f_{M2} \equiv \bigvee_{A \in \mathcal{A}} G F in(A)$;
- $f_{M2'} \equiv \bigvee_{A \in \mathcal{A}} F(in(A) \wedge X(in(A) \wedge X(in(A) \dots X(in(A)) \dots)))$
(there are k_A X operators);
- $f_{M3'} \equiv \bigvee_{A \in \mathcal{A}} F(eq(\mu_{A1}) \wedge F(eq(\mu_{A2}) \wedge \dots \wedge F(eq(\mu_{Ak_A}) \wedge F eq(\mu_{A1})) \dots))$.

For $M3$, we first construct $\mathcal{A}' = \{A \mid A \in \mathcal{A} \text{ and there is an infinite } M1'\text{-fair firing sequence for } \{A\} \text{ in } (P, T, \varphi, \mu_{A1})\}$. From Theorem 4.1, \mathcal{A}' can be constructed in PTIME. We now define:

- $f_{M3} \equiv \bigvee_{A \in \mathcal{A}'} F eq(\mu_{A1})$.

It should be clear that for each x , there is an infinite x -fair firing sequence in \mathcal{P} iff \mathcal{P} is a model for f_x .

□

We now turn our attention to the 6 transition-related types of fairness. The NTP^x , for $x \in \{T1, T1', T2, T2', T3, T3'\}$, has been shown to be decidable in [34]. However, no complexity analysis was given there. In what follows, we show that these problems are equivalent to the BP. The proofs require

showing the fact that if an infinite x -fair computation exists, then there must be a short “witness” to this fact. The proof of such a fact, generally speaking, is based on the method that Rackoff used in the complexity analysis of the BP in [29] (see also [30]).

A finite firing sequence $\sigma\sigma'$ is *self-covering* iff $\mu_0 \xrightarrow{\sigma} \mu_i \xrightarrow{\sigma'} \mu_j$ and $\mu_j \geq \mu_i$. (Note that σ' can be executed infinitely many times.) We define $T(\sigma)=\{t \mid \text{transition } t \text{ occurs in } \sigma\}$. (Similarly, $T(\sigma')=\{t \mid \text{transition } t \text{ occurs in } \sigma'\}$.) Then we have the following easily shown lemma:

Lemma 4.1: Given a Petri net \mathcal{P} and a finite set \mathcal{A} of finite nonempty sets of transitions, there is an infinite x -fair firing sequence, where $x \in \{T1, T1', T2, T2', T3, T3'\}$, iff there is a self-covering firing sequence $\sigma\sigma'$ such that

1. (T1-fair) $\exists A \in \mathcal{A}, A \cap T(\sigma) \neq \emptyset$.
2. (T1'-fair) $\exists A \in \mathcal{A}, T(\sigma) \cup T(\sigma') \subseteq A$.
3. (T2-fair) $\exists A \in \mathcal{A}, A \cap T(\sigma') \neq \emptyset$.
4. (T2'-fair) $\exists A \in \mathcal{A}, T(\sigma') \subseteq A$.
5. (T3-fair) $\exists A \in \mathcal{A}, T(\sigma') = A$.
6. (T3'-fair) $\exists A \in \mathcal{A}, A \subseteq T(\sigma')$.

We now give a lemma that allows us to derive an upper bound on the space complexity of NTP^x , where $x \in \{T1, T1', T2, T2', T3, T3'\}$. The proof closely parallels the proof of Rackoff [29]; hence, it is omitted.

Lemma 4.2: Given a transition t and a set of transitions B , the length of the shortest self-covering firing sequence $\sigma\sigma'$, where $t \in T(\sigma)$ and $T(\sigma')=B$, is bounded by $2^{2^{cn \log n}}$, for some constant c independent of t, B and n .

In order to derive lower bounds for NTP^x , $x \in \{T1, T1', T2, T2', T3, T3'\}$, we define the following problems, which are equivalent to BP with respect to PTIME many-one reductions. (Note that these equivalences are not known to hold for LOG/LIN many-one reductions.)

- **INF:**

Instance: Given a Petri net \mathcal{P} ,

Question: Is there an infinite firing sequence?

- **CP:**

Instance: Given a Petri net \mathcal{P} and a marking μ ,

Question: Is there a $\mu' \in R(\mathcal{P})$ such that $\mu \leq \mu'$?

The following lemma follows from results of Rackoff [29] and Lipton [21].

Lemma 4.3: $INF \equiv_{PTIME} CP \equiv_{PTIME} BP$.

We can now show

Theorem 4.3: $NTP^x \equiv_{PTIME} BP$, for $x \in \{T1, T1', T2, T2', T3, T3'\}$.

Proof. To show the $NTP^x \leq_{PTIME} BP$, we note that by applying Lemmas 4.1 and 4.2, we have an

$\text{NSPACE}(2^{cn \log n})$ upper bound for each problem. Since the BP is hard for $\text{NSPACE}(2^{p^{O(y)}})$ with respect to PTIME many-one reductions [21], $\text{NTP}^x \leq_{\text{PTIME}} \text{BP}$.

We now show $\text{BP} \leq_{\text{PTIME}} \text{NTP}^x$ for each x . Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be an arbitrary PN, and let $\mathcal{A} = \{T\}$. Clearly, \mathcal{P} has an infinite x -fair firing sequence for $x \in \{T1, T1', T2, T2', T3'\}$ iff it has an infinite firing sequence. Thus, $\text{BP} \leq_{\text{PTIME}} \text{INF} \leq_{\text{PTIME}} \text{NTP}^x$ for $x \in \{T1, T1', T2, T2', T3'\}$.

To show $\text{BP} \leq_{\text{PTIME}} \text{NTP}^{T3}$, we will use the fact that $\text{CP} \equiv_{\text{PTIME}} \text{BP}$. Let \mathcal{P} be an arbitrary PN, and let μ be an arbitrary marking. We construct \mathcal{P}' by adding to \mathcal{P} a new transition t which does nothing but which is only enabled at markings $\mu' \geq \mu$. Let $\mathcal{A} = \{\{t\}\}$. Clearly, \mathcal{P}' has an infinite T3-fair firing sequence iff there is a $\mu' \in R(\mathcal{P})$ such that $\mu \leq \mu'$. Thus, $\text{BP} \leq_{\text{PTIME}} \text{CP} \leq_{\text{PTIME}} \text{NTP}^{T3}$.

□

We now examine the NTP with respect to several notions of fairness in which the constraints are imposed in an implicit fashion, instead of by an explicit listing the markings and/or transitions that a “fair” firing sequence must visit. We will first examine three types introduced by Lehman, Pnueli, and Stavi [20] and two extensions given by Carstensen [7]. (See also Carstensen and Valk [8].) Given a Petri net \mathcal{P} and a set of subsets of transitions \mathcal{T} , an infinite firing sequence σ is said to be:

- *impartial* iff every transition in \mathcal{P} occurs infinitely often in σ .
- *just* iff every transition that is enabled almost everywhere in σ occurs infinitely often in σ .
- *fair* iff every transition that is enabled infinitely often in σ occurs infinitely often in σ .
- *fdp with respect to \mathcal{T}* (fdp- \mathcal{T}) iff for every $T \in \mathcal{T}$, if almost everywhere in σ some t in T is enabled, then some t' in T occurs infinitely often in σ . (Here, fdp stands for finite delay property.)
- *fair with respect to \mathcal{T}* (fair- \mathcal{T}) iff for every $T \in \mathcal{T}$, if some t in T is enabled infinitely often in σ , then some t' in T occurs infinitely often in σ .

Theorem 4.4: $\text{NTP}^{\text{imp}} \equiv_{\text{PTIME}} \text{BP}$.

Proof. It is not hard to see that NTP^{imp} is a special case of NTP^{T3} (by letting \mathcal{A} contain only the set of all transitions); hence, from Theorem 4.3, $\text{NTP}^{\text{imp}} \leq_{\text{PTIME}} \text{NTP}^{T3} \leq_{\text{PTIME}} \text{BP}$. In order to show that $\text{BP} \leq_{\text{PTIME}} \text{NTP}^{\text{imp}}$, we will use the fact that $\text{CP} \equiv_{\text{PTIME}} \text{BP}$. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be an arbitrary PN, and let μ be an arbitrary marking. We construct $\mathcal{P}' = (P, T', \varphi', \mu_0)$ by adding to \mathcal{P} a new transition t' such that:

- $\varphi(p, t') = \mu(p)$ for all $p \in P$; and
- $\varphi(t', p) = \mu(p) + \sum_{t \in T} \varphi(p, t)$ for all $p \in P$.

Clearly, there is an infinite impartial firing sequence in \mathcal{P}' iff there is a $\mu' \in R(\mathcal{P})$ such that $\mu' \geq \mu$.

□

In order to show the next theorem, we define the single-place zero-reachability problem, which is equivalent to RP.

- **RP-S0:**

Instance: A PN \mathcal{P} with a designated place p ,

Question: Is there a $\mu \in R(\mathcal{P})$ such that $\mu(p) = 0$?

The following lemma is well known; see, e.g., [26].

Lemma 4.4: $\text{RP-S0} \equiv_{PTIME} \text{RP}$.

At this time, we do not know whether NTP^{just} or $\text{NTP}^{fdp-\mathcal{T}}$ is decidable. The main difficulty, we feel, is due to the fact that these fairness properties are nonmonotonic in the sense that the existence of a just (fdp- \mathcal{T} , respectively) firing sequence starting in μ by no means guarantees a just (fdp- \mathcal{T}) firing sequence starting in any μ' , where $\mu' \geq \mu$. At the same time, we are unable to enforce zero-testing using either of these properties. However, we are able to show the following three theorems.

Theorem 4.5: $\text{RP} \leq_{PTIME} \text{NTP}^{fdp-\mathcal{T}}$ even if $\mathcal{T} = \{\{t\}\}$ for some transition t .

Proof. We will use the fact that $\text{RP-S0} \equiv_{PTIME} \text{RP}$. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be a PN with a designated place p_i . We will construct \mathcal{P}' as shown in Figure 4.1. There is clearly no infinite firing sequence in \mathcal{P}' in which t fires. Thus, in any infinite firing sequence that is fdp with respect to $\{\{t\}\}$, t must be disabled infinitely often; i.e., a marking μ must be reached such that $\mu(p_i) = 0$. On the other hand, if $\mu_0 \xrightarrow{\sigma} \mu, \mu(p_i) = 0$, then $\sigma t' t' t' \dots$ is fdp with respect to $\{\{t\}\}$. Therefore, $\text{RP} \leq_{PTIME} \text{RP-S0} \leq_{PTIME} \text{NTP}^{fdp-\mathcal{T}}$.

□

In [7], Carstensen showed that $\text{NTP}^{fdp-\mathcal{T}}$ is decidable if $|\mathcal{T}| = 1$. We are able to improve upon this result by showing the problem to be equivalent to RP. In this proof, as opposed to Theorem 4.2, the real power of the machinery developed in Section 3 is exploited. Without the umbrella of Theorem 3.2 and its subsequent refinement in Theorem 3.4, this proof would have been much longer and considerably more tedious. As it is, we are able to give a very succinct proof of an unobvious theorem.

Theorem 4.6: $\text{NTP}^{fdp-\mathcal{T}} \equiv_{PTIME} \text{RP}$ if $|\mathcal{T}| = 1$.

Proof. From Theorem 4.5, we need only show that $\text{NTP}^{fdp-\mathcal{T}} \leq_{PTIME} \text{RP}$ if $|\mathcal{T}| = 1$. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be an arbitrary PN, $T' \subseteq T$, and $\mathcal{T} = \{T'\}$. An infinite firing sequence σ is fdp with respect to \mathcal{T} iff

$$\langle \mathcal{P}, \sigma, 0 \rangle \models \text{GF}[(\bigwedge_{t \in T'} \neg (\bigwedge_{p \in P} \text{ge}(p, \varphi(p, t)))) \vee (\bigvee_{t \in T'} \text{fi}(t))].$$

Thus, from Theorem 3.4, $\text{NTP}^{fdp-\mathcal{T}} \leq_{PTIME} \text{RP}$ if $|\mathcal{T}| = 1$.

□

The following theorem gives a lower bound for NTP^{just} .

Theorem 4.7: $\text{RP} \leq_{PTIME} \text{NTP}^{just}$.

Proof. We will use the fact that $\text{RP} \leq_{PTIME} \text{NTP}^{fdp-\{\{t_i\}\}}$. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be an arbitrary PN, and let t_i be an arbitrary transition in T . We construct \mathcal{P}' as shown in Figure 4.2. Suppose σ is an infinite path in \mathcal{P} that is fdp with respect to $\{\{t_i\}\}$. We construct σ' by inserting infinitely often in σ $t'_j t''_j$ for all $j \neq i$, $1 \leq j \leq m$; since en is incremented infinitely often, this can clearly be done. Clearly, σ' is an infinite firing sequence in \mathcal{P}' that is fdp with respect to $\{\{t_i\}\}$. Furthermore, for all $j \neq i$, $1 \leq j \leq m$, t'_j and t''_j fire infinitely often, and t_i is disabled infinitely often. Therefore, σ' is just.

Now suppose σ is an infinite just firing sequence in \mathcal{P}' . In order for σ to be infinite, there clearly must be infinitely many occurrences of transitions from T . Therefore, by deleting all occurrences of transitions not

in T , we construct an infinite firing sequence σ' that is fdp with respect to $\{\{t_i\}\}$. Therefore, $RP \leq_{PTIME} NTP^{fdp-\{\{t\}\}} \leq_{PTIME} NTP^{just}$.

□

In [7], Carstensen showed the NTP^{fair} and the $NTP^{fair-\mathcal{T}}$ to be undecidable. In what follows, we improve this result by showing both to be complete for Σ_1^1 — the first level of the analytical hierarchy. We will later use these results to show a third version to be Σ_1^1 -complete. Our proofs involve certain generalized counter machines. We define an *infinite-branching counter machine* (ICM) to be a counter machine with the added capability to nondeterministically add to a counter any natural number in one step. The following lemma follows from Chandra [9].

Lemma 4.5: The set of all ICMs that contain an infinite computation on ϵ is Σ_1^1 -complete.

We are now ready to show that NTP^{fair} and $NTP^{fair-\mathcal{T}}$ are Σ_1^1 -complete.

Theorem 4.8: $NTP^{fair-\mathcal{T}}$ is Σ_1^1 -hard even if \mathcal{T} contains only singleton sets.

Proof. Let M be an arbitrary ICM. We will construct a PN $\mathcal{P} = (P, T, \varphi, \mu_0)$ and a set \mathcal{T} of singleton subsets of T such that \mathcal{P} has an infinite firing sequence that is fair with respect to \mathcal{T} iff M has an infinite computation on ϵ . \mathcal{P} must be able to simulate infinite-branching and zero-testing. The way we construct \mathcal{P} , zero-testing will be guaranteed to be correctly simulated only after some finite (but unknown) amount of time. Hence, we will first construct an ICM M' to simulate M such that for any configuration I of M' (whether reachable or not), any infinite computation of M' from I simulates an infinite computation of M . We will then construct \mathcal{P} to simulate M' .

We construct M' to repeatedly simulate finite subcomputations of M as follows. M' will have one set of counters to simulate M , one set of counters to store the final counter values from the previous simulation of M , two counters to simulate a clock (one for time remaining, the other for the initial value), and a fixed number of scratch counters. The finite-state control of M' will store the state of M in the current simulation and the final state from the previous simulation. Initially, the previous final configuration is the initial configuration of M , and the clock value is zero. M' iterates the following loop:

1. Let n be the value of the clock.
2. Simulate M for n steps; halt if M does.
3. If the final configuration of M differs from the previous final configuration, then halt.
4. Simulate M for one more step.
5. Save configuration of M .
6. Reset clock to $n + 1$.
7. Restore M to initial configuration.

Suppose M' starts in some arbitrary configuration I , and suppose there is an infinite computation σ from I . Clearly, each of the seven steps above must terminate, so each step occurs infinitely often in σ . Eventually, σ must reach Step 7. After this, when it reaches Step 2, it simulates a computation σ' of M for n steps,

where n is some natural number. Since σ is infinite, it will not halt in Step 3. Each subsequent execution of Step 4 extends σ' by one move; hence, σ simulates an infinite computation σ' of M .

We now construct \mathcal{P} to simulate M' using places to simulate states and counters in the obvious way, and transitions to simulate all moves except those of the form “ q_1 : if $c_i = 0$ then go to q_2 ” and “ q_1 : add some natural number to c_i and go to q_2 .” Carstensen [7] has shown how these moves can be implemented; the constructions are shown in Figures 4.3 and 4.4, respectively. Regarding Figure 4.3, since $\{t_2\} \in \mathcal{T}$, t_1 can be fired while $c_i > 0$ only finitely many times in any firing sequence that is fair with respect to \mathcal{T} . Likewise, for Figure 4.4, if t_2 is enabled, it must eventually fire in any infinite firing sequence that is fair with respect to \mathcal{T} . Clearly, if M has an infinite computation on ϵ , then \mathcal{P} has an infinite firing sequence that is fair with respect to \mathcal{T} . Suppose, on the other hand, that \mathcal{P} has an infinite firing sequence σ that is fair for \mathcal{T} . After some point in σ , σ correctly simulates M' from some initial configuration I . Since this simulation is infinite, M must have an infinite computation. □

Theorem 4.9: NTP^{fair} and $NTP^{fair-\mathcal{T}}$ are Σ_1^1 -complete.

Proof. $NTP^{fair-\mathcal{T}}$ was shown to be Σ_1^1 -hard in Theorem 4.8. Carstensen [7] gives a reduction from $NTP^{fair-\{\{t\}\}}$ to NTP^{fair} ; this reduction easily generalizes to $NTP^{fair-\mathcal{T}} \leq_{PTIME} NTP^{fair}$ if \mathcal{T} contains only singleton sets. Thus, from Theorem 4.8, NTP^{fair} is also Σ_1^1 -hard. See the proof of Theorem 4.10 below for an even more general version of this reduction.

Since NTP^{fair} is a restriction of $NTP^{fair-\mathcal{T}}$, we only need to show that $NTP^{fair-\mathcal{T}} \in \Sigma_1^1$. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be an arbitrary PN, and let \mathcal{T} be a set of subsets of T . Clearly an infinite firing sequence σ is fair with respect to \mathcal{T} iff it can be divided into a sequence of finite firing sequences $\sigma_0 \sigma_1 \dots$ such that for some subset \mathcal{T}' of \mathcal{T} , every set in \mathcal{T}' has at least one element in each σ_i , $i \geq 1$, and no transition in any set in $\mathcal{T} \setminus \mathcal{T}'$ is enabled at any time in σ after σ_0 . We therefore construct an ICM M to nondeterministically generate σ as follows. M first guesses \mathcal{T}' . It then guesses (in one step) n_0 , the length of σ_0 . Next, M nondeterministically generates a firing sequence σ_0 of length n_0 . For $i = 1, 2, \dots$, M then guesses n_i , the length of σ_i , and generates σ_i , verifying that it satisfies the conditions outlined above. Clearly, M has an infinite computation iff \mathcal{P} has an infinite firing sequence that is fair with respect to \mathcal{T} . Therefore, NTP^{fair} and $NTP^{fair-\mathcal{T}}$ are Σ_1^1 -complete. □

In [7], Carstensen also considered fairness for bounded PNs. Given a PN \mathcal{P} , we say an infinite firing sequence σ is

- *bd-fair* if \mathcal{P} is bounded and σ is fair.

Although $NTP^{bd-fair}$ is clearly decidable and as hard as BP, no tighter bounds have been given for the problem. We now show the problem to be nonprimitive recursive.

Theorem 4.10: $NTP^{bd-fair}$ is not primitive recursive.

Proof. We first define the following functions:

$$\begin{aligned} f_1(n) &= 2n, \\ f_i(n) &= f_{i-1}^{(n)}(1) \text{ for } i > 1, \end{aligned}$$

$$f(n) = f_n(n)$$

where $f_{i-1}^{(n)}$ denotes the n th-fold composition of f_{i-1} . It can be shown that f is not primitive recursive. Let M be an arbitrary $f(n)$ time bounded TM, and let x be an arbitrary input for M . Clearly, we can construct in PTIME a 2-counter machine M' that simulates M in $f(cn)$ time for some constant c . Carstensen [7] gave a PTIME construction of an unbounded PN \mathcal{P} with a designated transition t such that \mathcal{P} had an infinite firing sequence that was fair with respect to t iff M' halted on x . (In his construction, \mathcal{P} simulated an arbitrary 2-counter machine.) \mathcal{P} operated by first guessing m , the number of steps executed by M , then repeatedly simulating M on x . If M did not halt in m steps, \mathcal{P} would halt. \mathcal{P} contained three potentially unbounded places: the one containing m and two containing counter values. Clearly, since M' is $f(cn)$ time bounded, these places need not exceed $f(cn)$ to simulate M' . In [16], we gave a PTIME construction of a bounded PN which could produce a value of $f(cn)$ in a designated place (see also [24]). We can clearly use this construction to bound the three places in question by some number at least $f(cn)$. The details are left to the reader.

We will now show that given a bounded PN \mathcal{P} and a set \mathcal{T} of singleton subsets of transitions, we can construct in PTIME a bounded PN \mathcal{P}' that has an infinite bd-fair firing sequence iff \mathcal{P} has an infinite firing sequence that is fair with respect to \mathcal{T} . This construction is a generalization of one given in [7], and is shown in Figure 4.5. Note that the added places, $p'_1, \dots, p'_m, p''_1, p''_2$, must always contain a total of three tokens; hence \mathcal{P}' remains bounded. Suppose \mathcal{P} has an infinite path σ that is fair with respect to \mathcal{T} . Prior to each occurrence of t_i in σ , $1 \leq i \leq j$, we can insert the sequence t'_i, t''_i for some i' , and prior to each occurrence of t_i , $j+1 \leq i \leq m$, we can insert the sequence t'_i, t'_i, t''_i for some i' . It is easy to see that this new firing sequence is bd-fair. On the other hand, suppose there is an infinite firing sequence σ in \mathcal{P}' that is bd-fair. Since p''_2 must be positive before any transition in \mathcal{P} can fire, by removing all occurrences of transitions not in \mathcal{P} , we clearly get an infinite firing sequence that is fair with respect to \mathcal{T} .

□

In [5], Best extended the definition of fairness using the notions of i (∞)-*enabledness*. A transition t is said to be i -enabled (or ∞ -enabled if $i = \infty$) at a marking μ if there is a firing sequence σ no longer than i transitions such that $\mu \xrightarrow{\sigma} \mu'$, and t is enabled at μ' . For $1 \leq i \leq \infty$, an infinite computation σ is said to be:

- *i-fair* iff for every transition t , if t is i -enabled in infinitely many markings in σ , then t occurs infinitely often.

(Note that “0-fairness” coincides with “fairness” as defined in [20]. Also, an equivalent definition for ∞ -fairness was given in Queille and Sifakis [28].)

In what follows, we show that NTP^{i-fair} , $0 \leq i < \infty$, is Σ_1^1 -complete, but that $NTP^{\infty-fair}$ is equivalent to RP.

Theorem 4.11: For every i , $0 \leq i < \infty$, NTP^{i-fair} is Σ_1^1 -complete.

Proof. To show the lower bound, we will use a reduction from NTP^{i-fair} . Let \mathcal{P} be an arbitrary PN. We construct \mathcal{P}' as shown in Figure 4.6. Clearly, \mathcal{P} has an infinite fair firing sequence iff \mathcal{P}' has an infinite i -fair firing sequence. To show the upper bound we can clearly employ a similar strategy to that of the upper bound in Theorem 4.9, since i -enabledness is clearly a decidable property.

□

The following theorem is the last in which we use the machinery developed in Section 3. In this theorem, we use Theorem 3.2 directly to give a succinct reduction to RP. As in Theorem 4.6, a direct reduction to RP would have been much more tedious.

Theorem 4.12: $\text{NTP}^{\infty\text{-fair}} \equiv_{\text{PTIME}} \text{RP}$.

Proof. We first show $\text{NTP}^{\infty\text{-fair}} \leq_{\text{PTIME}} \text{RP}$. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be an arbitrary PN. For each $t \in T$, let μ_t be the minimum marking in which t is enabled. Consider the wff

$$f = \mathbf{F}[\bigwedge_{t \in T} (\neg \text{co}(\mu_t) \vee \mathbf{F} \text{fi}(t)) \wedge \bigwedge_{p \in P} \text{lp}(p)]$$

Clearly there is an infinite ∞ -fair firing sequence in \mathcal{P} iff \mathcal{P} is a finite model for f . Thus, from Theorem 3.2, $\text{NTP}^{\infty\text{-fair}} \leq_{\text{PTIME}} \text{RP}$.

We will now show $\text{RP} \leq_{\text{PTIME}} \text{NTP}^{\infty\text{-fair}}$. We will use the fact that $\text{RP-S0} \equiv_{\text{PTIME}} \text{RP}$. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be an arbitrary PN with a designated place p_i . We construct \mathcal{P}' as shown in Figure 4.7. Suppose \mathcal{P}' has an infinite ∞ -fair firing sequence. Since t'_1 is enabled until it fires, it must eventually fire. Clearly, t'_3 cannot fire in any infinite ∞ -fair firing sequence. Therefore, there must be a $\mu \in R(\mathcal{P})$ such that $\mu(p_i) = 0$. Suppose, on the other hand, that there is a marking $\mu \in R(\mathcal{P})$ such that $\mu(p_i) = 0$. Let σ be such that $\mu_0 \xrightarrow{\sigma} \mu$. Then the firing sequence $\sigma t'_1 t'_2 t'_2 t'_2 \dots$ is clearly ∞ -fair. □

Queille and Sifakis [28] have extended fairness in two other ways, namely, *fair choice from states*, and *fair reachability of predicates*. Applying these notions of fairness to Petri nets, we have that for a PN $\mathcal{P} = (P, T, \varphi, \mu_0)$, an infinite firing sequence σ satisfies:

- *fair choice from states* (state-fair, for short) iff for any marking μ reached infinitely often by σ , every transition enabled at μ is executed infinitely often from μ in σ .
- *fair reachability of predicates* (pred-fair, for short) iff for any (finite or infinite) set of markings M , if there are infinitely many i such that $\mu_0 \xrightarrow{\sigma[i]} \mu_i \xrightarrow{\sigma'_i} \mu'_i \in M$ for some σ'_i , then there must be infinitely many j such that $\mu_0 \xrightarrow{\sigma[j]} \mu_j \in M$.

At this time, we are unable to establish tight bounds for either $\text{NTP}^{\text{state-fair}}$ or $\text{NTP}^{\text{pred-fair}}$. $\text{NTP}^{\text{state-fair}}$ seems to be related to the problem of finding a *home state*, that is, a marking that is reachable from any reachable marking (see [11], where a decision procedure was given to determine whether a given marking is a home state). Because the definition of pred-fair is quantified over all sets of markings, it appears to be a very difficult problem. The next three theorems give the bounds we are able to derive for these two problems.

Theorem 4.13: $\text{NTP}^{\text{state-fair}}$ is decidable.

Proof. Let \mathcal{P} be an arbitrary PN. We first determine whether \mathcal{P} is bounded. If \mathcal{P} is unbounded, there is an infinite firing sequence σ which reaches each marking at most once. σ is clearly state-fair. On the other hand, if σ is bounded, we can construct the reachability graph. Then there is an infinite state-fair firing sequence iff the reachability graph contains a strongly connected component from which there is no exit. □

We now define the following problems, which will be used in giving our lower bound for $\text{NTP}^{\text{state-fair}}$.

- **BRP:**

Instance: Given a bounded Petri net \mathcal{P} and a marking μ ,

Question: Is $\mu \in R(\mathcal{P})$?

- **BRP-S0:**

Instance: Given a bounded Petri net \mathcal{P} with a designated place p ,

Question: Is there a $\mu \in R(\mathcal{P})$ such that $\mu(p) = 0$?

The following lemma can be shown in a similar manner as Lemma 4.4.

Lemma 4.6: $\text{BRP-S0} \equiv_{PTIME} \text{BRP}$.

Although it follows from [29,21] that $\text{BP} \leq_{PTIME} \text{BRP} \leq_{PTIME} \text{RP}$, we do not know whether either of these reductions can be strengthened to equivalences.

Theorem 4.14: $\text{BRP} \leq_{PTIME} \text{NTP}^{\text{state-fair}}$.

Proof. We will use the fact that $\text{BRP-S0} \equiv_{PTIME} \text{BRP}$. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be an arbitrary bounded PN with a designated place p_i . We construct \mathcal{P}' as shown in Figure 4.7. Suppose there is a $\mu \in R(\mathcal{P})$ such that $\mu(p_i) = 0$. Let σ be such that $\mu_0 \xrightarrow{\sigma} \mu$. Then the sequence $\sigma t'_1 t'_2 t'_2 t'_2 \dots$ in \mathcal{P}' is clearly state-fair. Now suppose there is an infinite state-fair path σ in \mathcal{P}' . Since \mathcal{P} is bounded, \mathcal{P}' is clearly bounded. Since t'_1 is enabled until it is fired, it must be fired in σ ; otherwise, there would be some marking reached infinitely often by σ in which t'_1 is enabled. Since t'_1 fires, t'_3 cannot fire; hence, there must be a marking $\mu \in R(\mathcal{P})$ such that $\mu(p_i) = 0$. Therefore, $\text{BRP} \leq_{PTIME} \text{BRP-S0} \leq_{PTIME} \text{NTP}^{\text{state-fair}}$. □

Theorem 4.15: $\text{RP} \leq_{PTIME} \text{NTP}^{\text{pred-fair}}$.

Proof. We will use the fact that $\text{RP-S0} \equiv_{PTIME} \text{RP}$. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be an arbitrary PN with a designated place p_i . We again construct \mathcal{P}' as shown in Figure 4.7. Suppose there is a $\mu \in R(\mathcal{P})$ such that $\mu(p_i) = 0$. Let σ be such that $\mu_0 \xrightarrow{\sigma} \mu$. We claim that $\sigma' = \sigma t'_1 t'_2 t'_2 t'_2 \dots$ is pred-fair. To see this, first note that after t'_1 fires only one marking is reachable. Since this marking is reached infinitely often by σ' , σ' must be pred-fair. Now suppose that there is an infinite pred-fair firing sequence σ in \mathcal{P}' . Let $M = \{\mu \mid \mu(p'_1) = 0 \text{ and } \mu(p'_2) = 0\}$. Since no marking in M can be reached in an infinite firing sequence, there must be some $i \in \mathbb{N}$ such that $\mu_0 \xrightarrow{\sigma[i]} \mu'$, such that no marking in M is in $R(P, T, \varphi, \mu')$. Clearly, t'_3 cannot occur in $\sigma[i]$, nor can it be enabled at μ' . Therefore, there must be some $\mu \in R(\mathcal{P})$ such that $\mu(p_i) = 0$. Thus, $\text{RP-S0} \leq_{PTIME} \text{RP} \leq_{PTIME} \text{NTP}^{\text{pred-fair}}$. □

Finally, we examine *equifairness* as defined by Francez [12]. Given a Petri net \mathcal{P} and an infinite firing sequence σ , σ is said to be

- *equifair* iff there exist infinitely many i 's such that all transitions occur the same number of times in $\sigma[i]$.

We will show this problem to be equivalent to the BP. In order to show this, we must introduce some terminology from [29] (see also [30]). Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be a PN. A *generalized marking* is a mapping $\mu : P \rightarrow \mathbb{Z}$, where \mathbb{Z} is the set of integers. A *generalized firing sequence* from μ_0 is any sequence of transitions. We

then extend the notation $\mu \xrightarrow{t} \mu'$ to generalized markings. An *i-loop* is a nonempty sequence σ of transitions $\mu_0 \xrightarrow{\sigma} \mu_0$ such that there are i places p_1, \dots, p_i such that for any $j \leq n$, where n is the length of σ , $\mu_0 \xrightarrow{\sigma[j]} \mu_j$ and $\mu_j(p_{i'}) \geq 0$ for $1 \leq i' \leq i$. Using a strategy similar to that of Rackoff [29], we can show the following lemma.

Lemma 4.7: Given a PN $\mathcal{P} = (P, T, \varphi, \mu_0)$ and an $i \leq |P|$, we can decide in space $2^{cn \log n}$ for some constant c whether there is an i -loop in \mathcal{P} .

We can now show

Theorem 4.16: $\text{NTP}^{\text{equifair}} \equiv_{\text{PTIME}} \text{BP}$.

Proof. We first show $\text{NTP}^{\text{equifair}} \leq_{\text{PTIME}} \text{BP}$. Let $\mathcal{P} = (P, T, \varphi, \mu_0)$ be an arbitrary PN. We construct \mathcal{P}' as shown in Figure 4.8. Clearly, \mathcal{P}' has a $|P|$ -loop iff \mathcal{P} has an infinite equifair firing sequence. Therefore, from Lemma 4.7 and [21], $\text{NTP}^{\text{equifair}} \leq_{\text{PTIME}} \text{BP}$.

We now show $\text{BP} \leq_{\text{PTIME}} \text{NTP}^{\text{equifair}}$. We will use the fact that $\text{CP} \equiv_{\text{PTIME}} \text{BP}$. Let \mathcal{P} be an arbitrary PN, and let μ be an arbitrary marking. We construct \mathcal{P}' as in Theorem 4.4. Clearly, if \mathcal{P}' has an infinite equifair path, there is a $\mu' \in R(\mathcal{P})$ such that $\mu' \geq \mu$. Suppose, conversely, that $\mu_0 \xrightarrow{\sigma} \mu' \geq \mu$. (Here, σ is a firing sequence, as opposed to a generalized firing sequence.) Let i be the maximum number of occurrences of any transition in σ . Clearly, σ followed by i occurrences of t' is a firing sequence in \mathcal{P}' . This firing sequence can then be followed by enough occurrences of each transition so that each transition occurs exactly i times. This can then be repeated infinitely many times. Therefore, $\text{BP} \leq_{\text{PTIME}} \text{CP} \leq_{\text{PTIME}} \text{NTP}^{\text{equifair}}$.

□

5 Conclusion

We have exhibited a temporal logic powerful enough to express certain fairness constraints, yet whose FMCP is equivalent to the reachability problem for PNs. This logic was instrumental in showing seven fair nontermination problems to be equivalent to RP. In developing this logic, we were able to answer a question left open in [15], namely, is there a decision procedure for $\mathcal{L}^\infty(Q')$? We were able to give a positive answer to this question and show that the problem is equivalent to reachability. One question that remains open is whether there is a decision procedure for $\tilde{\mathcal{L}}(Q', \text{GF})$; i.e., wffs including predicates from Q' and the operators GF , \wedge , \vee , and \neg , where \neg is allowed only on predicates.

The following table summarizes the fairness results of this paper.

NTP	complexity	NTP	complexity
M1	RP	imp	BP
M1'	NL	just	\geq RP
M2	RP	fair	Σ_1^1
M2'	RP	bd-fair	NPR
M3	RP	fdp- \mathcal{T}	\geq RP
M3'	RP	fdp-T	RP
T1	BP	fair- \mathcal{T}	Σ_1^1
T1'	BP	i-fair	Σ_1^1
T2	BP	∞ -fair	RP
T2'	BP	state-fair	D, \geq BRP
T3	BP	pred-fair	\geq RP
T3'	BP	equifair	BP

- \geq BRP: as hard as the BRP,
- \geq RP: as hard as the RP,
- NPR: decidable, but not primitive recursive,
- D: decidable,
- fdp- \mathcal{T} : fdp- \mathcal{T} with $|\mathcal{T}| = 1$.

Most of the problems examined have been shown to be either equivalent to boundedness, equivalent to reachability, or Σ_1^1 -complete. Two exceptions to this general rule are $NTP^{M1'}$ and $NTP^{bd-fair}$. $NTP^{M1'}$ can be decided in nondeterministic logspace because the entire allowable reachability set is explicitly given as input. Using the fact that bounded PNs can generate very large numbers, we were able to show that $NTP^{bd-fair}$ is not primitive recursive. Aside from the fact that the precise complexity of RP is still unknown, the precise complexities of four of the fair nontermination problems we have examined remain open. Two of these, NTP^{just} and $NTP^{fdp-\mathcal{T}}$ are particularly interesting because they are related to the open temporal logic question mentioned above; i.e., they are both expressible in $\tilde{\mathcal{L}}(\mathcal{Q}', \mathbf{GF})$.

Acknowledgment: We would like to thank Prof. H. Carstensen for his suggestions and encouraging comments. In particular, we thank him for suggesting that model checking with respect to general PNs might be decidable for $\tilde{\mathcal{L}}(\mathcal{Q}', \mathbf{F}, \mathbf{X})$.

References

- [1] K. Apt and D. Kozen. Limits for automatic verification of finite-state concurrent systems. *Information Processing Letters*, 22:307–310, 1986.
- [2] K. Apt and E. Olderog. Proof rules and transformations dealing with fairness. *Sci. Comput. Prog.*, 3:65–100, 1983.
- [3] K. Apt, A. Pnueli, and J. Stavi. Fair termination revisited with delay. In *Proceedings of the 2nd Conference on Foundations of Software Technology and Theoretical Computer Science (FST-TCS)*, pages 146–170, 1982.
- [4] H. Baker. Rabin’s proof of the undecidability of the reachability set inclusion problem of vector addition systems. MIT Project MAC, 1973. CSGM 79, Cambridge, MA, 1973.
- [5] E. Best. Fairness and conspiracies. *Information Processing Letters*, 18:215–220, 1984. Addendum, Vol. 19, page 162, 1984.

- [6] G. Brams. *Reseaux de Petri: Theorie et Pratique – Tome 1: Theorie et Analyse*. Masson, Paris, 1983.
- [7] H. Carstensen. Decidability questions for fairness in Petri nets. In *Proceedings of the 4th Symposium on Theoretical Aspects of Computer Science*, pages 396–407, 1987. LNCS 247.
- [8] H. Carstensen and R. Valk. Infinite behaviour and fairness in Petri nets. In *Advances in Petri Nets 1984*, pages 83–100, Springer-Verlag, 1985. LNCS 188.
- [9] A. Chandra. Computable nondeterministic functions. In *Proceedings of the 19th IEEE Symposium on the Foundations of Computer Science*, pages 127–131, 1978.
- [10] E. Emerson and C. Lei. Modalities for model checking: branching time logic strikes back. To appear in *Science of Computer Programming*. Some of these results were presented at the *18th Annual Hawaii International Conference on System Sciences* and at the *12th Annual ACM Symposium on Principles of Programming Languages*.
- [11] D. Escrig. A collection of algorithms to decide liveness and other related properties of a place transition system. Unpublished manuscript.
- [12] N. Francez. *Fairness*. Springer-Verlag, 1986.
- [13] M. Hack. The equality problem for vector addition systems is undecidable. *Theoret. Comp. Sci.*, 2:77–95, 1976.
- [14] D. Harel. Effective transformations on infinite trees, with applications to high undecidability, dominoes, and fairness. *JACM*, 33:224–248, 1986.
- [15] R. Howell and L. Rosier. On questions of fairness and temporal logic for conflict-free Petri nets. In *Proceedings of the 8th European Workshop on Applications and Theory of Petri Nets*, pages 197–214, 1987.
- [16] R. Howell, L. Rosier, D. Huynh, and H. Yen. Some complexity bounds for problems concerning finite and 2-dimensional vector addition systems with states. *Theoret. Comp. Sci.*, 46:107–140, 1986.
- [17] N. Immerman. Nondeterministic space is closed under complement. Unpublished manuscript.
- [18] R. Kosaraju. Decidability of reachability in vector addition systems. In *Proceedings of the 14th Annual ACM Symposium on Theory of Computing*, pages 267–280, 1982.
- [19] L. Landweber. Decision problems for ω -automata. *Math. Syst. Theory*, 3:376–384, 1969.
- [20] D. Lehman, A. Pnueli, and J. Stavi. Impartiality, justice, and fairness: the ethics of concurrent termination. In *Proceedings of the 8th International Colloquium on Automata, Languages, and Programming*, pages 264–277, 1981. LNCS 115.
- [21] R. Lipton. *The Reachability Problem Requires Exponential Space*. Technical Report 62, Yale University, Dept. of CS., Jan. 1976.

- [22] Z. Manna and A. Pnueli. The modal logic of programs. In *Proceedings of the 6th International Colloquium on Automata, Languages, and Programming*, pages 385–410, 1979. LNCS 71.
- [23] E. Mayr. An algorithm for the general Petri net reachability problem. *SIAM J. Comput.*, 13(3):441–460, 1984. A preliminary version of this paper was presented at the *19th Annual Symposium on Theory of Computing*, 1981.
- [24] H. Müller. Weak Petri net computers for Ackermann functions. *Elektronische Informationsverarbeitung und Kybernetik*, 21:236–244, 1985.
- [25] E. Olderog and K. Apt. Transformations realizing fairness assumptions for parallel programs. In *Proceedings of STACS '84*, pages 20–42, 1984. LNCS 166.
- [26] J. Peterson. *Petri Net Theory and the Modeling of Systems*. Prentice Hall, Englewood Cliffs, NJ, 1981.
- [27] A. Pnueli. The temporal logic of programs. In *Proceedings of the 19th Annual Symposium on Foundations of Computer Science*, 1977.
- [28] J. Queille and J. Sifakis. Fairness and related properties in transition systems — a temporal logic to deal with fairness. *Acta Informatica*, 19:195–220, 1983.
- [29] C. Rackoff. The covering and boundedness problems for vector addition systems. *Theoret. Comp. Sci.*, 6:223–231, 1978.
- [30] L. Rosier and H. Yen. A multiparameter analysis of the boundedness problem for vector addition systems. *J. of Computer and System Sciences*, 32(1):105–135, 1986.
- [31] A. Sistla and E. Clarke. The complexity of propositional linear temporal logic. *JACM*, 32:733–749, 1985.
- [32] A. Sistla and S. German. Reasoning with many processes. In *Proceedings of IEEE Symposium on Logic in Computer Science*, pages 138–152, Ithaca, New York, 1987.
- [33] I. Suzuki. Fundamental properties and applications of temporal Petri nets. In *Proceedings of the 19th Annual Conference on Information Sciences and Systems*, pages 641–646, 1985.
- [34] R. Valk and M. Jantzen. The residue of vector sets with applications to decidability problems in Petri nets. *Acta Informatica*, 21:643–674, 1985.

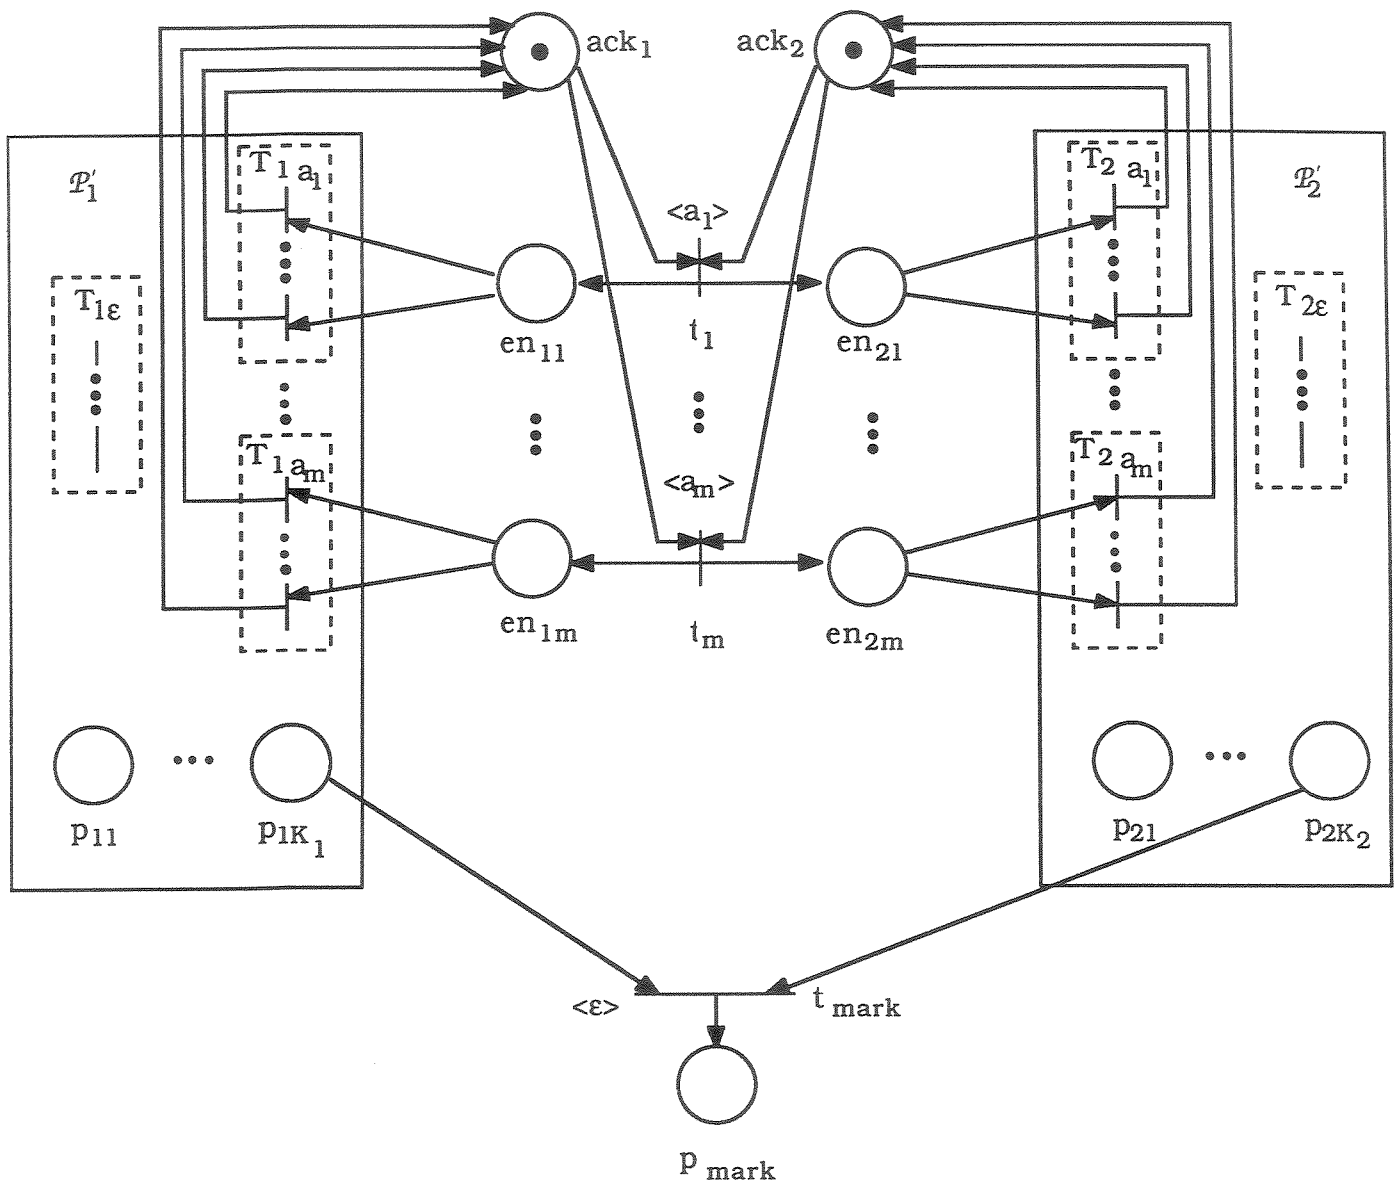


Figure 3.1: \mathcal{P}' for $f_1 \wedge f_2$.

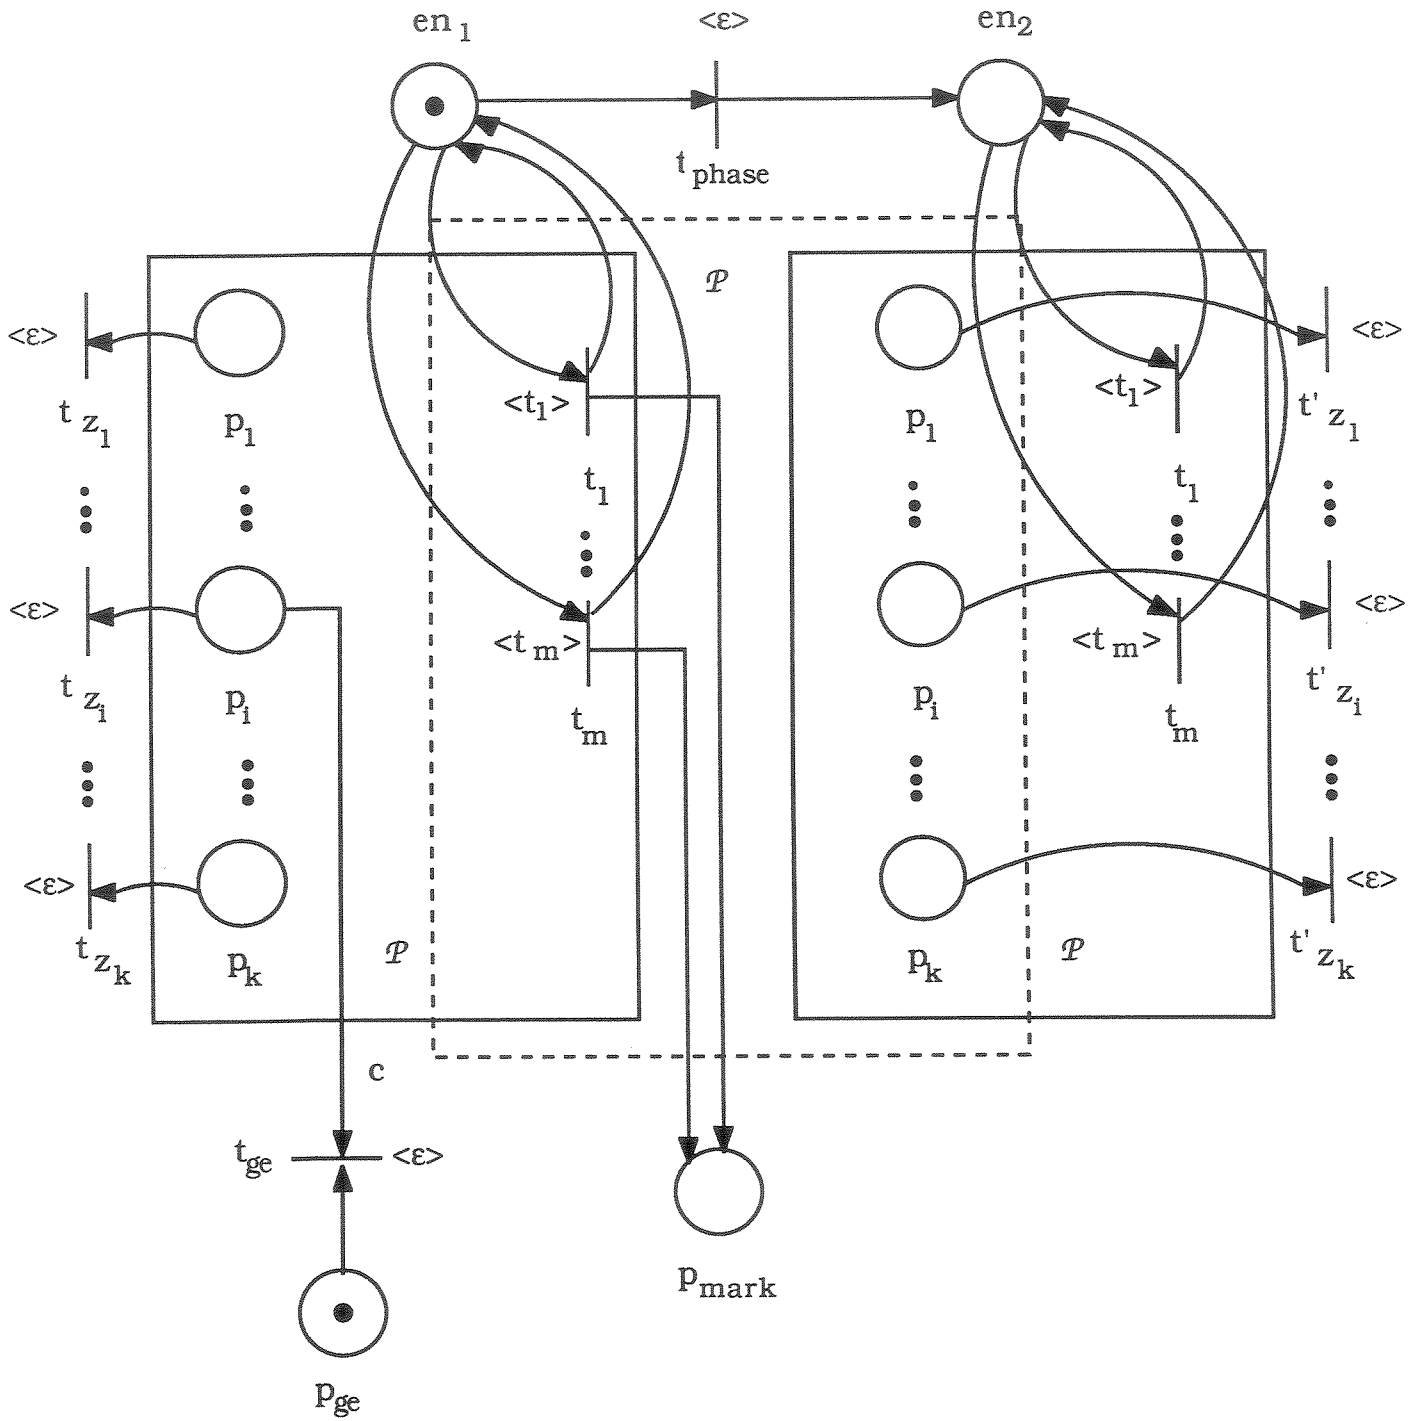


Figure 3.2: \mathcal{P}' for $ge(p_1, c)$.

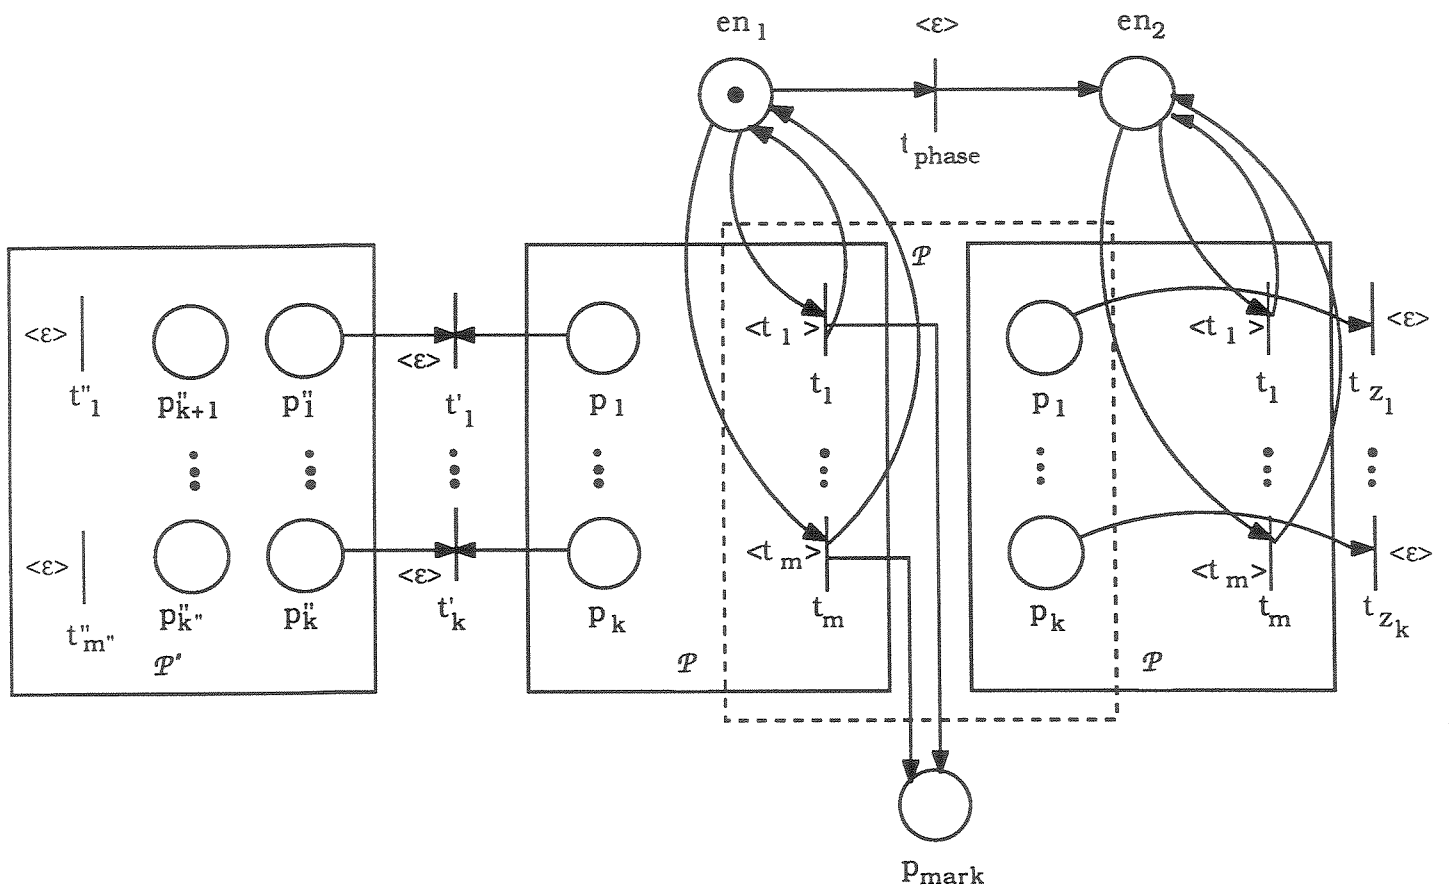


Figure 3.3: \mathcal{P}' for $\neg \text{co}(\mu)$.




Figure 4.1: \mathcal{P}' for $\text{fdp-}\mathcal{T}$, $\mathcal{T} = \{t\}$.




Figure 4.2: \mathcal{P}' for justice.




Figure 4.3: Implementation of zero-testing for fair- \mathcal{T} .




Figure 4.4: Implementation of infinite-branching for fair- \mathcal{T} .




Figure 4.5: \mathcal{P}' for bd-fairness.




Figure 4.6: \mathcal{P}' for i -fair.




Figure 4.7: \mathcal{P}' for ∞ -fair, state-fair, and pred-fair.

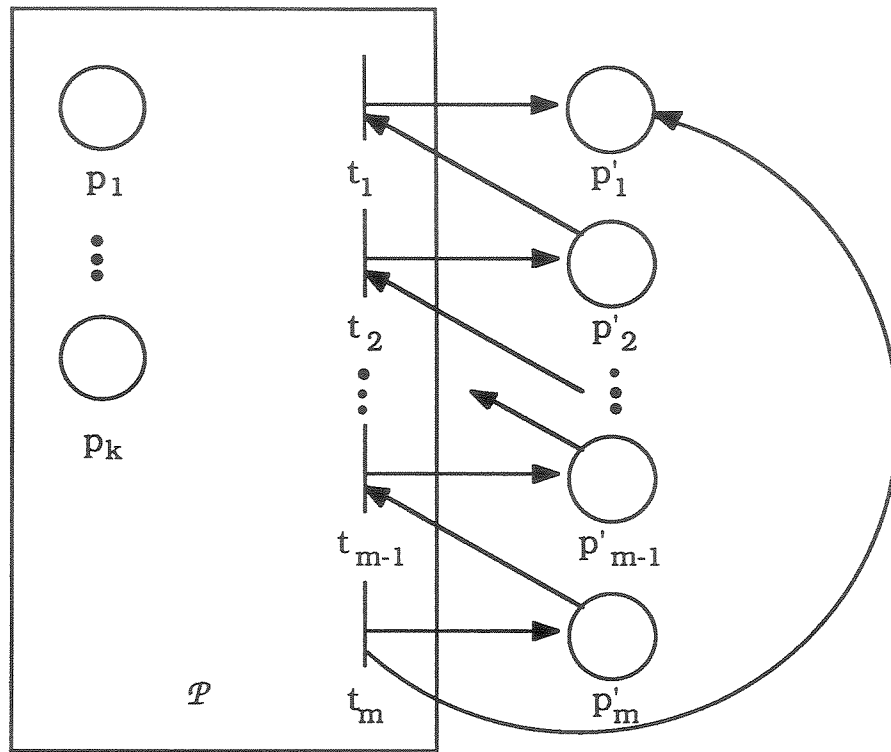


Figure 4.8: \mathcal{P}' for equifairness.