Fairness and Petri nets

Answer sketches for Home Assignment 2

To hand in before or on February 17, 2013.

	28	29	30	31			
February					1	2	3
	4	5	6	7	8	9	10
	11	12	13	14	15	16	17

Electronic versions (PDF only) can be sent by email to ⟨schmitz@lsv.ens-cachan.fr⟩, paper versions should be handed in on the 17th or put in my mailbox at LSV, ENS Cachan.

The following formula will be applied to late assignments: $g - 3d + 1$, where $0 \leq g \leq 20$ is the original grade and $d > 0$ is the number of days of delay, rounded up onward from the 13th at midnight.

This assignment is concerned with *fairness* in Petri nets. Fairness properties are employed to rule out behaviours where a process might wait indefinitely before being activated.

The numbers in the margins next to exercises are indications of time and difficulty. Judging from your answers, I did a pretty bad job of evaluating this difficulty, sorry about that.

1 The Fairness Fragment

We restrict ourselves to a fragment TL(AP, GF) where the only temporal modality is GF (sometimes also written $\overset{\infty}{\mathsf{F}}$), i.e. with syntax

$$\varphi ::= p \mid \top \mid \neg\varphi \mid \varphi \wedge \varphi \mid \mathsf{GF}\,\varphi$$

where p ranges over AP.

Let us consider the case of state-based LTL model-checking for Petri nets. In this framework, we are checking infinite sequences of markings of a Petri net $\mathcal{N} = \langle P, T, W, m_0 \rangle$, i.e. infinite sequences $m_0 m_1 \cdots m_i \cdots$ in $(\mathbb{N}^P)^\omega$ such that, for all i in $\mathbb{N}$, $m_i \to_{\mathcal{N}} m_{i+1}$ is a transition step of $\mathcal{N}$ according to some t in T, thus verifying for all p in P that $m_i(p) \geq W(p,t)$ and $m_{i+1}(p) = m_i(p) - W(p,t) + W(t,p)$. More generally, the *effect* $\Delta(u)$ of a transition sequence u in T^* is defined by $\Delta(\varepsilon) = 0^P$ and $\Delta(ut) = \Delta(u) - W(P,t) + W(t,P)$.

The atomic propositions in $\mathrm{AP} \subseteq P$ represent places, and are interpreted over such a sequence by $m_i \models p$ iff $m_i(p) > 0$. Put differently, a Petri net $\mathcal{N}$ gives rise to an infinite Kripke structure $\mathfrak{M}_{\mathcal{N}} = \langle \mathbb{N}^P, \to_{\mathcal{N}}, \{m_0\}, \mathrm{AP}, \ell \rangle$ where $\ell(m) = \{p \in \mathrm{AP} \mid m(p) > 0\}$.

Recall (from TD 6, Exercise 3) that the model-checking problem for state-based TL(AP, X, U) is undecidable in general and PSPACE-complete for safe nets. In the following we restrict ourselves to the model-checking problem for *ordinary* Petri nets, which verify $W(t,p) \leq 1$ and $W(p,t) \leq 1$ for all p in P and t in T (note that this does *not* imply that the net is safe).

[4] **Exercise 1** (Fairness in Petri Nets)**.** Let $\mathcal{N}$ be an ordinary Petri net, m a marking in $\{0,1\}^P$, and t be a transition in T. Reduce the following problems to $\text{MC}^\exists(\mathsf{AP}, \mathsf{GF})$ model-checking instances on an ordinary Petri net $\mathcal{N}'$:

repeated coverability ($\text{RC}^\exists$): there exists an infinite execution where m is covered infinitely often, i.e. such that $m_i \geq m$ for an infinite number of indices i,

weak fairness ($\text{WF}^\forall$): every infinite execution either fires t infinitely often, or t is infinitely often not firable,

strong fairness ($\text{SF}^\forall$): in every infinite execution, if t is firable infinitely often, then it is actually fired infinitely often.

One can also consider the existential questions $\text{WF}^\exists$ or $\text{SF}^\exists$, which ask whether there *exist* some fair infinite execution. We will see in sections 2 and 3 that $\text{RC}^\exists$ and $\text{WF}^\exists$ are decidable.

It turns out that $\text{SF}^\exists$ is not decidable. Because the proof of this result [1] is a bit involved, we rather look at a decidable case:

Exercise 2 (Strong Fairness in Safe Nets)**.** We consider the $\text{SF}^\exists$ problem when the Petri net $\mathcal{N}$ is known to be *safe*, i.e. verifying $m(p) \leq 1$ for all reachable m and all places p in P.

[4] Show that $\text{SF}^\exists$ is PSPACE-complete when $\mathcal{N}$ is safe. *Hint: For hardness, reduce from reachability in safe nets.*

The upper bound was immediate, using Exercise 1 and the fact that $\text{MC}^\exists(\text{AP}, \mathsf{X}, \mathsf{U})$ is in PSPACE for 1-safe Petri nets (from TD 6, Exercise 3.1).

2 Repeated Coverability

We show in this section that $\text{RC}^\exists$ is decidable, relying for this on the properties of the coverability graph seen during TD 7.

Exercise 3 (Decidability)**.** We prove in this exercise that repeated coverability is decidable. Let $\mathcal{N}$ be Petri net, G be its coverability graph and m some marking in $\mathbb{N}^P$.

[4] Show that there exists an infinite computation s.t. $m \leq m_i$ for infinitely many indices i iff there exists an accessible loop $m' \xrightarrow{v}_G m'$ in G s.t. $m \leq m'$ and $\Delta(v) \geq 0^P$. Conclude. *Hint: Use Exercise 2 from TD 7.*

Assume $m_0 m_1 m_2 \cdots$ is an infinite computation of $\mathcal{N}$ with $m \leq m_i$ for infinitely many indices $i_0 < i_1 < i_2 < \cdots$. Because $(\mathbb{N}^{|P|}, \leq)$ is a wqo, we can extract an increasing subsequence $m_{j_1} \leq m_{j_2} \leq \cdots$ for $j_0 < j_1 < \cdots$ among $\{i_0, i_1, \dots\}$, i.e. $m_{j_k} \xrightarrow{v_k}_{\mathcal{N}} m_{j_{k+1}}$ verifies $\Delta(v_k) \geq 0^P$ for all k, and $m_0 \xrightarrow{u}_{\mathcal{N}} m_{j_0}$ for some u.

By Exercise 2.1 from TD 7, there exists an infinite computation $m_0 \xrightarrow{u}_G m'_{j_0} \xrightarrow{v_1} m'_{j_1} \cdots$ in G s.t. $m_{j_k} \leq m'_{j_k}$ for all k. Because for all k

- $m'_{j_k}(p) = m_{j_k}(p) \leq m_{j_{k+1}} = m'_{j_{k+1}}$ for all $p \in P \setminus \Omega(m'_{j_{k+1}})$,
- $\Omega(m'_{j_k}) \subseteq \Omega(m'_{j_{k+1}})$, i.e. $m'_{j_k}(p) = m'_{j_{k+1}}(p) = \omega$ for all $p \in \Omega(m'_{j_k})$, and
- $m'_{j_k}(p) < m'_{j_{k+1}}(p) = \omega$ for all $p \in \Omega(m'_{j_{k+1}}) \setminus \Omega(m'_{j_k})$,

we deduce that $m'_{j_k} \leq m'_{j_{k+1}}$. However, by definition of G, after $K \leq |P|$ of these steps, two successive values are equal: whenever $m_{j_k} < m_{j_{k+1}}$, at least one new ω value is introduced. Thus there exists a loop on m_{j_K} in G, with effect $\Delta(v_K) \geq 0^P$.

Conversely, if there exists an accessible loop $m_0 \xrightarrow{u}_G m' \xrightarrow{v}_G m'$ in G s.t. $m \leq m'$ and $\Delta(v) \geq 0^P$, then by Exercise 2.2 from TD 7 and using a partial marking $\max(\Theta(v)(p), m(p))$ for all p in $\Omega(m')$, there exists a run $m_0 \xrightarrow{u_0 w_1^{k_1} u_1 \cdots w_n^{k_n} u_{n+1}}_{\mathcal{N}} m_1$. Since $m_1(p) = m'(p) \geq m(p)$ on the places p in $P \setminus \Omega(m')$ and $m_1(p) \geq m(p)$ on the places p in $\Omega(m')$, $m_1 \geq m$. Furthermore, since $m' \xrightarrow{v}_G m'$, $m_1(p) = m'(p) \geq \Theta(v)(p)$ on the places p in $P \setminus \Omega(m')$, and $m_1(p) \geq \Theta(v)(p)$ on the places p in $\Omega(m')$, thus the sequence v can be fired in m_1: $m_1 \xrightarrow{v}_{\mathcal{N}} m_2$. Since $\Delta(v) \geq 0^P$, $m_2 \geq m_1$ thus $m_2 \geq m$ and $m_2 \geq \Theta(v)$, and we can construct an infinite computation that covers m infinitely often by repeatedly applying v.

Remark 3.1 (Non-negative Cycles in G). Beware that, in a cycle $m' \xrightarrow{v}_G m'$ in the coverability graph, the overall effect $\Delta(v)$ is not necessarily ≥ 0. Consider for instance the following Petri net and its coverability graph:

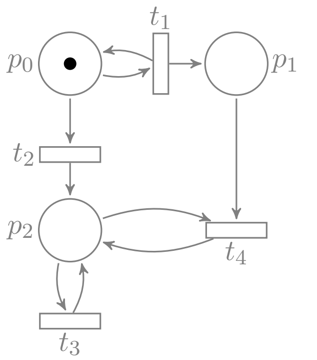

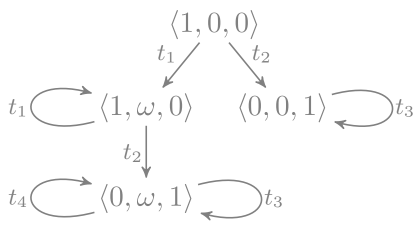

Its coverability graph has a *decreasing* cycle on $\langle 0, \omega, 1 \rangle$ when firing t_4.

Regarding the next question on action-based LTL model-checking, consider the labeling $\lambda(t_1) = \lambda(t_2) = \{a\}$, $\lambda(t_3) = \emptyset$, and $\lambda(t_4) = \{b\}$. Then $\mathcal{N} \models_\forall \mathsf{FG}\,\neg b$, because, in any infinite execution of $\mathcal{N}$, either t_2 is never fired and b never holds, or t_4 can only be fired as many times as t_1 was fired before firing t_2, i.e. finitely many times. However,

$G \models_\exists \mathsf{GF}\, b$, by going to the state $\langle 0, \omega, 1\rangle$ and looping on t_4. Thus one cannot simply check G with an LTL formula in order to deduce that the same formula holds in $\mathcal{N}$.

Remark 3.2 (Finding a Witness in G). I did not really expect you to give details on how a reachable cycle $m' \xrightarrow{v}_G m'$ with $m \leq m'$ and $\Delta(v) \geq 0^P$ could be found in G, so I was not disappointed not to see it in most answers. In fact, it is easily reduced to a linear programming problem. This problem is thus in PTime in the size of G. A reference on the complexity of such problems is [2] (beware that most of the paper is dedicated to the case where $|P|$ is fixed).

Remark 3.3 ($\mathrm{RC}^\forall$ is Undecidable). Because at least one of you claimed that the same ideas would lead to a decidability proof for $\mathrm{RC}^\exists$, I feel compelled to spell out an undecidability proof (I couldn't find any reference for this result). The proof uses arguments somewhat similar to the ones used in the undecidability proof of $\mathrm{SF}^\exists$ by Carstensen [1].

2-Counter Machines. The reduction is from the halting problem in 2-counters Minsky machines. Formally, such a machine is a tuple $\mathcal{M} = \langle Q, \delta, q_0, q_h\rangle$ where Q is a finite set of states containing an *initial state* q_0 and a *halting state* q_h, and δ maps a state in $Q \setminus \{q_h\}$ to a tuple

1. (i, q) standing for the action "$\texttt{c}_i\texttt{++; goto } q$", or
2. (i, q, q') standing for the action "$\texttt{if c}_i\texttt{==0 then goto } q\texttt{; else c}_i\texttt{--; goto } q'$"

for some i in $\{0, 1\}$ and q, q' in Q. A *configuration* of $\mathcal{M}$ is a triple $C = (q, n_0, n_1)$ in $Q \times \mathbb{N} \times \mathbb{N}$ holding the current state and the current values in the two counters $\texttt{c}_0$ and $\texttt{c}_1$. The initial configuration is $C_0 = (q_0, 0, 0)$, and in a configuration (q, n_0, n_1) we either halt if $q = q_h$, or apply the action $\delta(q)$. Because $\mathcal{M}$ is deterministic, there is a single possible execution starting from C_0. The question whether $\mathcal{M}$ halts, i.e. whether it reaches a configuration (q_h, n_0, n_1), regardless of the values n_0 and n_1 of the counters, is undecidable.

Never Covered. As a first step, let us prove that the following variant of $\mathrm{RC}^\forall$ is undecidable: given $\langle \mathcal{N}, m\rangle$, does there exist an infinite execution where m is *never* covered?

Given an instance $\langle \mathcal{M}\rangle$ of the 2-counter halting problem, we construct a Petri net $\mathcal{N} = \langle P, T, W, m_0\rangle$ that will simulate the counter machine: we set

$$\bar{Q} = \{\bar{q} \mid \exists i, q', q''.\delta(q) = (i, q', q'')\} \qquad\qquad P = Q \uplus \bar{Q} \uplus \{z, c_0, c_1\}\ ,$$

i.e. each state and each counter of $\mathcal{M}$ is associated to a place, and states with actions $\delta(q)$ of type 2 have an additional state $\bar{q}$; the place z is used to express that we are currently testing a counter for zero. The transitions of $\mathcal{N}$ depend on the actions of $\mathcal{M}$ and are best defined in pictures:

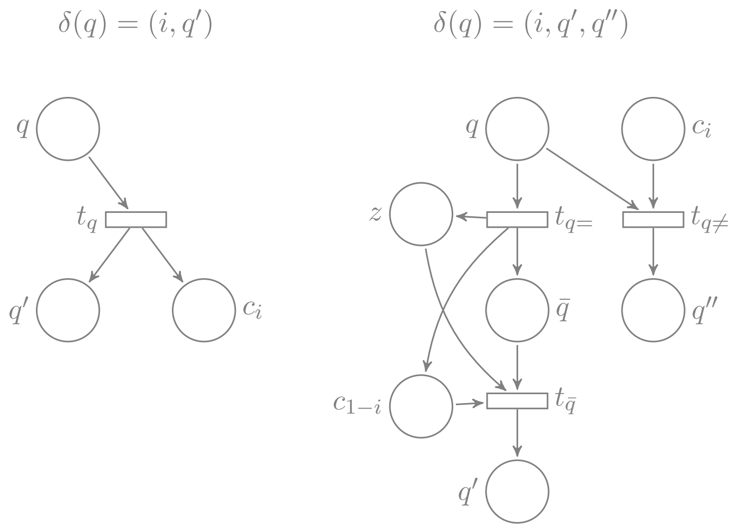

The initial marking is $m_0 = q_0$. We claim that $\mathcal{M}$ does *not* halt if and only if there exists an infinite computation where $m = z \oplus c_0 \oplus c_1$ is never covered (here $\oplus$ is the multiset union operator, i.e. $m(z) = m(c_0) = m(c_0) = 1$ and $m(p) = 0$ for all $p \in P \setminus \{m, c_0, c_1\}$).

Let us start with a few easy observations: the place z is marked whenever some $\bar{q}$ is, and the places in Q and $\bar{Q}$ contain together exactly one token at all times. We can with a configuration $C = (q, n_0, n_1)$ of $\mathcal{M}$ a sequence $\mu(C)$ of one or two markings in $\mathbb{N}^P$ defined by $\mu(C) = q \oplus c_0^{n_0} \oplus c_1^{n_1} \cdot \mu'(C)$ where $\mu'(C) = z \oplus \bar{q} \oplus c_1^{n_1+1}$ if $\delta(q) = (0, q', q'')$ and $n_0 = 0$, $\mu'(C) = z \oplus \bar{q} \oplus c_0^{n_0+1}$ if $\delta(q) = (1, q', q'')$ and $n_1 = 0$, and $\mu'(C)$ is the empty sequence otherwise. We can extend μ to be a homomorphism from $(Q \times \mathbb{N} \times \mathbb{N})^*$ to $(\mathbb{N}^P)^*$.

In the case of $\delta(q) = (i, q', q'')$, note that $\mathcal{N}$ simulates $\mathcal{M}$ with a non-deterministic choice instead of the deterministic action: the transition $t_{q=}$ can be fired even if c_i is not empty. It is rather immediate that every execution $C_0 C_1 \cdots$ of $\mathcal{M}$ can be simulated by $\mathcal{N}$ (by the execution $\mu(C_0 C_1 \cdots)$), but that some executions of $\mathcal{N}$ do not have counterparts in $\mathcal{M}$, i.e. those executions where, in at least one occasion, a transition $t_{q=}$ is fired athough the corresponding counter place c_i is not empty. In such a case, the sequence of markings does not have a pre-image by μ.

If we reach some marking covering z, then we are necessarily covering some $\bar{q}$, i.e. we are simulating a zero-test for some $\delta(q) = (i, q, q')$; thus, if we are covering m, then our simulation of this test is incorrect: the associated configuration has $c_i > 0$, although in $\bar{q}$ we are already commited to the case `c`$_i$`==0` of our action.

Now, if $\mathcal{M}$ does not halt, then there is an infinite execution $(q_0, 0, 0)(q_1, c_{0,1}, c_{1,1}) \cdots$ of $\mathcal{M}$. Then we can simulate this infinite execution in $\mathcal{N}$ starting from m_0. Conversely, if there exists an infinite execution $m_0 m_1 \cdots$ of $\mathcal{N}$ that never covers m, then we can associate to this execution an infinite sequence $\mu^{-1}(m_0 m_1 \cdots)$ of configurations of $\mathcal{M}$: because we are never covering m, we know that our simulation only allows to choose a zero-test when the corresponding counter value is 0, and the sequence $m_0 m_1 \cdots$ is in the range of μ. Because this execution in $\mathcal{M}$ is infinite, $\mathcal{M}$ does not halt.

Eventually Never Covered. Let us now turn to $\mathrm{RC}^\forall$: we actually consider its complement, which asks whether there exists an infinite execution verifying $\mathsf{FG}\bigvee_{p\in m}\neg p$, i.e. where m is *eventually never* covered.

We are again going to reuse the previous construction of a Petri net $\mathcal{N}$ from a Minsky machine $\mathcal{M}$. Observe that, if $\mathcal{N}$ eventually never covers m, then it means that it eventually simulates $\mathcal{M}$ faithfully. The only issue is that the configuration of $\mathcal{M}$ at the point where the simulations starts being faithful can be arbitrary, and different from the initial configuration $C_0 = (q_0, 0, 0)$. Here we proceed differently from the "time budget" technique of Carstensen [1] and rather reduce from the immortality problem:

Immortality Problem. Given a 2-counter Minsky machine $\mathcal{M}$, the *immortality problem* (IP) asks whether there exists some configuration $C = (q, n_0, n_1)$ such that $\mathcal{M}$ does not halt when started in C. Hooper [4] proves this problem to be undecidable.

Let us reduce from IP: we modify the Petri net $\mathcal{N}$ constructed from $\mathcal{M}$: we construct a net $\mathcal{N}' = \langle P', T', W', m'_0\rangle$ and a marking m such that $\mathcal{M}$ has an immortal configuration if and only if $\mathcal{N}'$ has an infinite execution that eventually never covers m. Our net $\mathcal{N}'$ is based on $\mathcal{N}$ but features an initialization component tasked with generating some initial configuration for $\mathcal{M}$ to start with: we add a place i (for "initializing") to P and the following transitions:

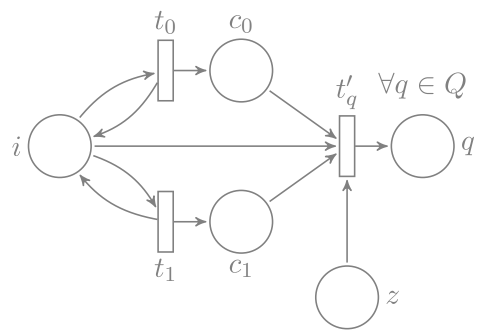

Let $m'_0 = i \oplus z \oplus c_0 \oplus c_1$; then $\mathcal{N}'$ initially puts some values $n_0 + 1$ and $n_1 + 1$ in c_0 and c_1 by repeatedly firing t_0 and t_1, before firing some "start" transition t'_q and simulating $\mathcal{N}$ with q as initial state, n_0 in c_0 and n_1 in c_1.

Let again $m = z \oplus c_0 \oplus c_1$: we claim that $\mathcal{M}$ has an immortal configuration if and only if $\langle \mathcal{N}', m\rangle$ is *not* an instance of $\mathrm{RC}^\forall$, i.e. if and only if there is an infinite execution in $\mathcal{N}'$ that eventually never covers m.

First suppose there exists an infinite execution $m'_0 m'_1 \cdots$ of $\mathcal{N}'$ that eventually never covers m, i.e. there exists an index j such that for all $k \geq j$, $m_k \not\geq m$. Because during the initialization phase of $\mathcal{N}'$ all the markings cover m, in this infinite execution, some "start" transition t'_q must be fired before this index j. Thus the infinite suffix $m'_j m'_{j+1} \cdots$ of this execution is in "running mode", where $\mathcal{N}'$ simulates $\mathcal{N}$, which in turn simulates $\mathcal{M}$. Because we never cover m in this suffix, the execution in $\mathcal{M}$ starting in $C = \mu^{-1}(m'_j)$ (or $C = \mu^{-1}(m'_j m'_{j+1})$ or $C = \mu^{-1}(m'_{j-1} m'_j)$ if places in $\bar{Q}$ are involved) is infinite, i.e. the configuration C is immortal in $\mathcal{M}$. Conversely, if there exists an immortal configuration in $\mathcal{M}$, say $C = (q, n_0, n_1)$, then the initialization phase of $\mathcal{N}'$ can build the marking $q \oplus n_0 \oplus n_1$ by firing n_0 times the transition t_0, n_1 times the transition t_1, and the "start" transition t'_q, and $\mathcal{N}'$ can then simulate $\mathcal{M}$ without errors, i.e. without ever covering m after the initialization.

Undecidability of SF$^\exists$. Finally, let us see how the previous construction can be used to show the undecidability of SF$^\exists$. It suffices to add a transition t with incoming arcs from z, c_0 and c_1 and no outgoing arcs: then firing t deadlocks the Petri net (because the $t_{\bar{q}}$ transitions are then inhibited), thus no infinite execution of the net can ever fire t, and any strongly fair infinite execution needs to eventually never cover $m = W(P, t)$.

Exercise 4 (Action-Based LTL)**.** Recall from Exercise 3 of TD 6 that *action-based* LTL considers *labeled* Petri nets $\langle \mathcal{N}, \lambda \rangle$ where λ is a labelling from T to 2^{AP}. The model-checking problem then considers the infinite sequences $\lambda(t_1)\lambda(t_2)\cdots$ of transition labels along an execution $m_0 \xrightarrow{t_1}_{\mathcal{N}} m_1 \xrightarrow{t_2}_{\mathcal{N}} m_2 \cdots$; alternatively, we can consider the Kripke structure associated with $\mathcal{N}$ to be $\mathfrak{M}^\lambda_{\mathcal{N}} = \langle \mathbb{N}^P \times T, T', I, \ell \rangle$ where $T' = \{((m,t),(m',t')) \mid m \xrightarrow{t}_{\mathcal{N}} m' \wedge m' \geq W(P,t')\}$, $I = \{m_0\} \times \{t \in T \mid m_0 \geq W(P,t)\}$, and $\ell(m,t) = \lambda(t)$.

[3] Show that action-based LTL model-checking is decidable for labeled Petri nets.

Given an action-based LTL formula φ, construct its Büchi automaton $\mathcal{B}_{\neg\varphi} = \langle Q, \Sigma, \delta, I, F \rangle$; wlog. let $I = \{q_0\}$. From the labeled Petri net $\mathcal{N} = \langle P, T, W, m_0, \lambda \rangle$ to be checked against φ and $\mathcal{B}_{\neg\varphi}$, we construct a Petri net $\mathcal{N}_{\neg\varphi} = \langle P \uplus Q, T', W', m'_0 \rangle$ s.t. $\mathcal{N}_{\neg\varphi}$ has an infinite computation that covers some place in F infinitely often iff $\mathcal{N} \models \neg\varphi$.

Define for this $T' = \{(t,(q,a,q')) \in T \times \delta \mid \lambda(t) = a\}$ and

$$\begin{aligned} W'(p,(t,(q,a,q'))) &= W(p,t) & W'((t,(q,a,q')),p) &= W(t,p)\\ W'(q'',(t,(q,a,q'))) &= (q'' = q) & W'(t,(q,a,q')),q'') &= (q'' = q')\\ m'_0(p) &= m_0(p) & m'_0(q'') &= (q'' = q_0) \end{aligned}$$

for all p in P, q'' in Q, t in T, and (q,a,q') in δ with $\lambda(t) = a$.

Thus there is a computation $m'_0 \xrightarrow{t_1,(q_0,a_1,q_1)} m'_1 \xrightarrow{t_2,(q_1,a_2,q_2)} m'_2 \cdots$ in $\mathcal{N}_{\neg\varphi}$ iff there is one $m_0 \xrightarrow{t_1} m_1 \xrightarrow{t_2} m_2 \cdots$ in $\mathcal{N}$ and one $q_0 \xrightarrow{a_1} q_1 \xrightarrow{a_2} q_2 \cdots$ with $\lambda(t_i) = a_i$ and $m'_i = m_i \oplus q_i$, which yields the result.

Remark 4.1 (Complexity). Although the algorithms in exercises 3 and 4 are of Ackermannian complexity (due to their relying on the construction of the coverability graph), the problems are actually ExpSpace-complete, a result due to Habermehl [3].

3 Weak Fairness

We examine in this section the relationship between the existential weak fairness problem WF$^\exists$ and the *reachability problem* (RP) in Petri nets: given a Petri net $\mathcal{N}$ and a marking

m in $\mathbb{N}^P$, does there exist a finite run $m_0 \to_{\mathcal{N}}^* m$? This problem is known to be decidable [9, 7, 8], although its exact complexity is still open.

Exercise 5 (Lower Bounds)**.** We want to exhibit a reduction from RP to $\mathrm{WF}^\exists$.

[2] 1. Show that RP can be reduced to the following problem: given a Petri net $\mathcal{N}$ and a place p, does there exist a reachable marking m (i.e. verifying $m_0 \to_{\mathcal{N}}^* m$) such that $m(p) = 0$?

Let $\langle \mathcal{N}, m\rangle$ be an instance of RP with $\mathcal{N} = \langle P, T, W, m_0\rangle$; we construct an instance $\langle \mathcal{N}', s\rangle$ of the empty place problem s.t. $m_0 \to_{\mathcal{N}}^* m$ iff $\exists m'.m'(s) = 0 \wedge m'_0 \to_{\mathcal{N}'}^* m'$. The idea behind this reduction is that the place s in $\mathcal{N}'$ will hold the *sum* plus one of the numbers of tokens in all the places in $\mathcal{N}$, and a new transition t_m of $\mathcal{N}'$ will empty this place if and only if m was reachable in $\mathcal{N}$.

Given a marking m in $\mathbb{N}^P$, we write $\sigma(m)$ for the sum of its components, i.e. for $\sum_{p\in P} m(p)$. We construct $\mathcal{N}' = \langle P', T \uplus \{t_m\}, W', m'_0\rangle$ where $P' = P \uplus \{s\}$, $m'_0 = m_0 \oplus s^{1+\sigma(m_0)}$, where the restriction of W' to P and T is W, but each transition t in T additionally verifies $W'(s,t) = \sigma(W(P,t))$ and $W'(t,s) = \sigma(W(t,P))$, and the transition t_m has $W'(P', t_m) = m \oplus s^{1+\sigma(m)}$ and $W'(t_m, P') = 0$. Note that the modified transitions in T preserve as an invariant that $m'(s) = 1 + \sigma(m'_{|P})$ for all reachable markings m' in $\mathcal{N}$.

Assume $m_0 \to_{\mathcal{N}}^* m$. Then, by firing the same sequence of transitions in $\mathcal{N}'$, $m'_0 = m_0 \oplus s^{1+\sigma(m_0)} \to_{\mathcal{N}'}^* m \oplus s^{1+\sigma(m)} \xrightarrow{t_m}_{\mathcal{N}'} 0$, i.e. a marking with 0 in s. Conversely, assume that $m'_0 \to_{\mathcal{N}'}^* m'$ with $m'(s) = 0$. Because $m'_0(s) > 0$ and due to the invariant on the transitions in T, necessarily t_m was fired last from the marking $m'' = m \oplus s^{1+\sigma(m)}$: hence the same sequence of transitions up to that marking m'' reaches m in $\mathcal{N}$.

[3] 2. Show that RP can be reduced to $\mathrm{WF}^\exists$.

We reduce from an instance $\langle \mathcal{N}, p\rangle$ of the previous problem. We construct for this a new Petri net $\mathcal{N}' = \langle P \uplus \{r\}, T \uplus \{t_p, t_e\}, W', m_0 \oplus r\rangle$, where W' extends W, where the place r is a "running place" and t_e is a dummy transition with $W'(P, t_e) = W'(t_e, P) = 0$ and $W'(r, t_e) = W'(r,t) = W'(t,r) = W'(t_e, r) = 1$ for all $t \in T$.

The transition t_p is used for the fairness condition: we let $W'(r, t_p) = 1$, $W'(P, t_p) = p$, and $W'(t_p, P \uplus \{r\}) = 0$, so that t_p can only be fired if p is marked, and if fired then it deadlocks $\mathcal{N}'$. Thus t_p cannot be fired infinitely often, and weak fairness can only be enforced if it is infinitely often not firable along some infinite execution.

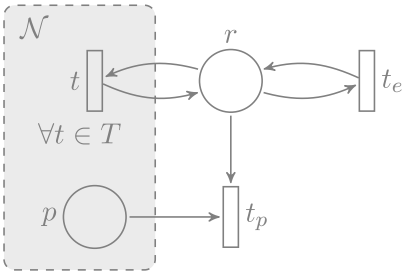

Assume that $m_0 \to^*_{\mathcal{N}} m$ with $m(p) = 0$. Then the same sequence of transitions can be fired in $\mathcal{N}'$, leading to $m_0 \oplus r \to^*_{\mathcal{N}'} m \oplus r$, from which t_e can be fired infinitely often: this constructs an infinite execution of $\mathcal{N}$ where t_p is infinitely often not firable, i.e. a weakly fair execution for t_p.

Conversely, if there exists an infinite execution $m'_0 m'_1 \cdots$ in $\mathcal{N}'$ where t_p is infinitely often not firable. Because any infinite execution leaves r marked at all times, i.e. $m'_i(r) = 1$ for all i, this means that p is empty infinitely often, i.e. $m'_i(p) = 0$ for infinitely many i. Pick the first such i: the execution $m_0 \oplus r \xrightarrow{w}_{\mathcal{N}'} m_i \oplus r$ with $w \in (T')^*$ can be mapped to an execution $m_0 \xrightarrow{\pi(w)}_{\mathcal{N}} m$ for the projection defined by $\pi(t_e) = \varepsilon$ and $\pi(t) = t$ for all $t \in T$. The marking $m'_i = m_i \oplus r$ verifies $m_i(p) = m'_i(p) = 0$.

Exercise 6 (Decidability)**.** We want to show that $\mathrm{WF}^\exists$ is decidable. Let us first concentrate on the subcase where there exists an infinite execution where t is infinitely often *not* firable.

[1] 1. Reduce this case to the question whether there exists a place p in P and an execution $m_0 \to^*_{\mathcal{N}} m \to^+_{\mathcal{N}} m'$ with $m \leq m'$ and $m(p) = m'(p) = 0$.

The question is badly written... The problem I had in mind was: *given* p, whether there exists an execution $m_0 \to^*_{\mathcal{N}} m \to^+_{\mathcal{N}} m'$ with $m \leq m'$ and $m(p) = m'(p) = 0$.

Because we assume our Petri nets to be ordinary, if t is not firable in a marking m, then there exists some place p such that $W(p,t) = 1$ and $m(p) = 0$.

We reduce the existence of an infinite execution where t is infinitely often not firable to a finite disjunction of instances of our new problem, one for each p such that $W(p,t) = 1$.

Assume first that there exists an infinite execution $m_0 \to^*_{\mathcal{N}} m \xrightarrow{w}_{\mathcal{N}} m'$ with w in T^+ verifying $\Delta(w) \geq 0$ and $m(p) = m'(p) = 0$ where $W(p,t) = 1$. Then $\Delta(w)(p) = 0$, and because $m' \geq m$, we can fire again the sequence w from m', leading to a marking $m'' \geq m' \geq m$ with $m''(p) = m'(p) = m(p) = 0$, and so on, thereby building an infinite execution where t is infinitely often not firable.

Conversely, suppose that there exists an infinite execution $m_0 m_1 \cdots$ in $\mathcal{N}$ such that t is infinitely often not firable. Color each of the infinitely many markings m_i that do not cover t in this sequence with c_p for some place p with $W(p,t) = 1$

but $m_i(p) = 0$. Since P is finite, there exists some color c_p that occurs infinitely often; this defines an infinite subsequence $m_{i_0} m_{i_1} \cdots$ of markings in $\mathbb{N}^P$ where $m_{i_j}(p) = 0$ for all j. By Dickson's Lemma, there exist two indices $j_1 < j_2$ such that $m_{i_{j_1}} \leq m_{i_{j_2}}$, thus exhibiting the desired witness $m_0 \to_{\mathcal{N}}^* m_{i_{j_1}} \to_{\mathcal{N}}^+ m_{i_{j_2}}$ with $m_{i_{j_1}}(p) = m_{i_{j_2}}(p) = 0$.

[2] 2. Reduce the previous question to an instance of the reachability problem.

An instance of the previous problem is a pair $\langle \mathcal{N}, p_0 \rangle$. We construct an instance $\langle \mathcal{N}', m \rangle$ of RP from it: the idea is that $\mathcal{N}'$ will work in three phases on a set of places $P \uplus \bar{P}$, such that during the first phase it simulates $m_0 \to_{\mathcal{N}}^* m$ on both the places in P and $\bar{P}$, and during the second phase it simulates $m \to_{\mathcal{N}}^+ m'$ on the places in P only. The third phase checks whether the execution is a witness for our initial problem: by decrementing the places p and $\bar{p}$ synchronously, $\mathcal{N}'$ can check that $m' \geq m$, and the final reachability test checks that $m(p_0) = m'(p_0) = 0$.

Let $P' = P \uplus \bar{P} \uplus \{p_1, p_2, p_3, p_r\}$ where $\bar{P}$ is a disjoint copy of P. The places p_1, p_2, p_3 indicate in which phase we are working. Let $T' = T_1 \uplus T_2 \uplus T_3 \uplus \{t_1, t_2, t_r\}$.

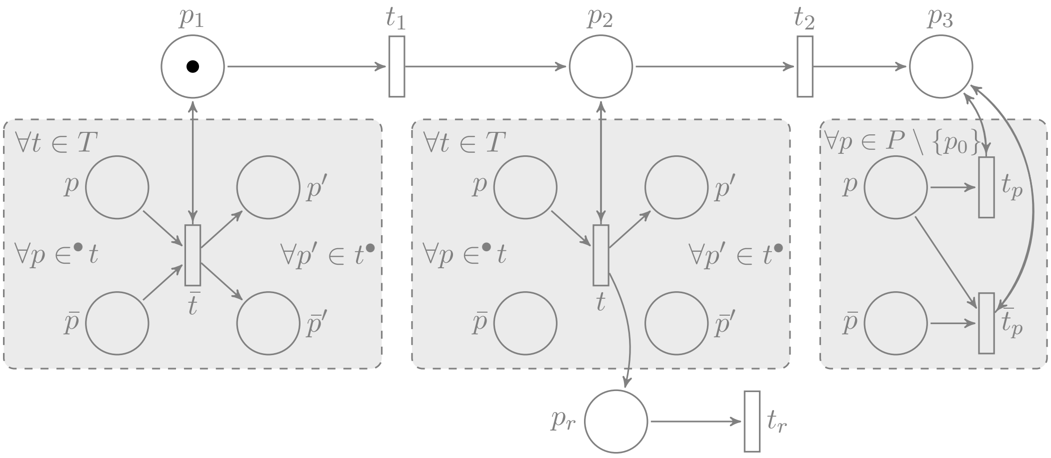

1. In the first phase, while there is a token in p_1, the transitions $\bar{t}$ in $T_1 = \{\bar{t} \mid t \in T\}$ verify $W'(p_1, \bar{t}) = W'(\bar{t}, p_1) = 1$, $W'(p, \bar{t}) = W'(\bar{t}, p) = 0$ for $p \in \{p_2, p_r, p_3\}$, $W'(P, \bar{t}) = W'(\bar{P}, \bar{t}) = W(P, t)$, and $W'(\bar{t}, P) = W'(\bar{t}, \bar{P}) = W(t, P)$; thus they simulate $\mathcal{N}$ on P and $\bar{P}$ simultaneously.
2. In the second phase, the transitions t in $T_2 = T$ verify $W'(p_2, \bar{t}) = W'(\bar{t}, p_2) = W'(t, p_r) = 1$, $W'(p_r, t) = 0$, $W'(p, t) = W'(t, p) = 0$ for $p \in \bar{P} \uplus \{p_1, p_3\}$, $W'(P, t) = W(P, t)$, and $W'(t, P) = W(t, P)$; thus they simulate $\mathcal{N}$ on P alone and increment p_r with each transition.
3. Finally, let $T_3 = \{t_p, \bar{t}_p \mid p \in P \setminus \{p_0\}\}$. The transitions $\bar{t}_p$ verify $W'(p, \bar{t}_p) = W'(\bar{p}, \bar{t}_p) = W'(p_3, \bar{t}_p) = W'(\bar{t}_p, p_3) = 1$ and have zero weights on all the other places: they decrement the places $p \neq p_0$ and $\bar{p}$ synchronously. The transitions t_p decrement the places $p \neq p_0$: $W'(p, t_p) = W'(p_3, t_p) = W'(t_p, p_3) = 1$ and zero everywhere else.

The transition t_r allows to decrement the place p_r: $W'(p_r, t_r) = 1$ and zero everywhere else. The transitions t_1 and t_2 move a token from p_i to p_{i+1} for $i \in \{1, 2\}$.

Given a marking m over $\mathbb{N}^P$, we write $\bar{m}$ for the marking over $\mathbb{N}^{\bar{P}}$ verifying $m(p) = \bar{m}(\bar{p})$ for all p in P. We then define the initial marking $m'_0 = p_1 \oplus m_0 \oplus \bar{m}_0$ and the target marking $m = p_3 \oplus p_r$.

Assume first $m_0 \xrightarrow{u}_{\mathcal{N}} m_1 \xrightarrow{v}_{\mathcal{N}} m_2$ with u in T^*, v in T^+, $m_1(p_0) = m_2(p_0) = 0$, and $m_1 \leq m_2$. Then we can simulate this execution by

$$
\begin{aligned}
m'_0 = m_0 \oplus \bar{m}_0 \oplus p_1 &\xrightarrow{\bar{u}}_{\mathcal{N}'} m_1 \oplus \bar{m}_1 \oplus p_1 \xrightarrow{t_1}_{\mathcal{N}'} m_1 \oplus \bar{m}_1 \oplus p_2 \\
&\xrightarrow{v}_{\mathcal{N}'} m_2 \oplus \bar{m}_1 \oplus p_r^{|v|} \oplus p_2 \xrightarrow{t_2}_{\mathcal{N}'} m_2 \oplus \bar{m}_1 \oplus p_r^{|v|} \oplus p_3 \\
&\xrightarrow{t_r^{|v|-1}}_{\mathcal{N}'} m_2 \oplus \bar{m}_1 \oplus p_r \oplus p_3 \\
&\xrightarrow{\prod_{p \in P \setminus \{p_0\}} \bar{t}_p^{m_1(p)}}_{\mathcal{N}'} (m_2 \ominus m_1) \oplus p_r \oplus p_3 \\
&\xrightarrow{\prod_{p \in P \setminus \{p_0\}} t_p^{m_2(p)-m_1(p)}}_{\mathcal{N}'} p_r \oplus p_3 \;.
\end{aligned}
$$

Conversely, assume that m is reachable in $\mathcal{N}'$. Then we have an execution rather similar to the above, except that the transitions t_p, $\bar{t}_p$, and t_r might occur in a different order (t_r might even occur during phase 2). By semi-commutativity arguments, namely that

- for all t in T_2, if $t_r t$ can be fired, then $t t_r$ can also be fired, and for all t in T_3, if $t t_r$ can be fired, then $t_r t$ can be fired,
- for all p in $P \setminus \{p_0\}$, if $t_p \bar{t}_p$ can be fired, then $\bar{t}_p t_p$ can be fired,
- for all $p \neq p'$ in $P \setminus \{p_0\}$, $t_p t_{p'}$ can be fired iff $t_{p'} t_p$ can be fired and $\bar{t}_p \bar{t}_{p'}$ can be fired iff $\bar{t}_{p'} \bar{t}_p$ can be fired,

we can actually find an execution of the desired form. Then, there exists an execution $m_0 \xrightarrow{u}_{\mathcal{N}} m_1 \xrightarrow{v}_{\mathcal{N}} m_2$ with $|v| > 0$ since p_r is marked, $m_1(p_0) = m_2(p_0) = 0$ since p_0 is not marked in m and is not modified by the transitions in phase 3, and $m_1 \leq m_2$ since we were able to execute the transitions $\bar{t}_p$ some (possibly null) amount of times to reach a marking $m_2 \ominus m_1$ over P.

[2] 3. Deduce that $\mathrm{WF}^\exists$ is decidable.

The other possibility for the existence of an infinite weakly fair execution is for t to be fired infinitely often, which easily reduces to an RC instance.

Remark 6.1. The interreducibility of $\mathrm{WF}^\exists$ and RP was first noted by Howell et al. [5]. The decidability of $\mathrm{WF}^\exists$ has been generalized by Jančar [6] to a fragment of state-based $\mathrm{MC}^\exists(\mathsf{GF})$ over Petri nets.

References

[1] H. Carstensen. Decidability questions for fairness in Petri nets. In F. J. Brandenburg, G. Vidal-Naquet, and M. Wirsing, editors, *STACS '87*, volume 247 of *LNCS*, pages 396–407. Springer, 1987. doi:10.1007/BFb0039622.

[2] E. Cohen and N. Megiddo. Strongly polynomial-time and NC algorithms for detecting cycles in periodic graphs. *J. ACM*, 40(4):791–830, 1993. doi:10.1145/153724.153727.

[3] P. Habermehl. On the complexity of the linear-time μ-calculus for Petri nets. In P. Azéma and G. Balbo, editors, *ICATPN '97*, volume 1248 of *LNCS*, pages 102–116. Springer, 1997. doi:10.1007/3-540-63139-9_32.

[4] P. K. Hooper. The undecidability of the Turing machine immortality problem. *J. Symb. Log.*, 31(2): 219–234, 1966. doi:10.2307/2269811.

[5] R. R. Howell, L. E. Rosier, and H.-C. Yen. A taxonomy of fairness and temporal logic problems for Petri nets. *Theor. Comput. Sci.*, 82(2):341–372, 1991. doi:10.1016/0304-3975(91)90228-T.

[6] P. Jančar. Decidability of a temporal logic problem for Petri nets. *Theor. Comput. Sci.*, 74(1):71–93, 1990. doi:10.1016/0304-3975(90)90006-4.

[7] S. R. Kosaraju. Decidability of reachability in vector addition systems. In *STOC '82*, pages 267–281. ACM Press, 1982. doi:10.1145/800070.802201.

[8] J. Leroux. Vector addition system reachability problem: a short self-contained proof. In *POPL 2011*, pages 307–316. ACM Press, 2011. doi:10.1145/1926385.1926421.

[9] E. W. Mayr. An algorithm for the general Petri net reachability problem. In *STOC '81*, pages 238–246. ACM Press, 1981. doi:10.1145/800076.802477.