

Some Petri Net Problems with Counterintuitive Solutions

Piotr Chrzastowski-Wachtel^{1,†}, Piotr Ulanowski^{1,†}

¹*Institute of Informatics, University of Warsaw, Banacha 2, PL-02-097 Warszawa, Poland*

Abstract

The first author has been giving a one-semester monograph lecture for decades at the Warsaw University. During the lectures he has offered solving Petri net puzzles, as one of the ways to pass the course. Some of them are pretty hard and counterintuitive. Despite the excellency of students at the Warsaw University, some of the puzzles were solved by only a few of them. This paper is a joint paper with one of the current students (the second author), who has solved the hardest puzzles in a convincing way.

1. Introduction

Petri nets are a very elegant framework to expose the problems of concurrent systems. They allow to define problems in a comprehensible way and to focus on important issues like liveness or boundedness of a given state. Some of the problems concerning liveness and boundedness are counterintuitive. They offer deep didactic value to look for something that appears impossible at the first glance. Solving such problems gives the students much satisfaction. The questionnaires filled anonymously after each course give no doubt about it. But it also demonstrates good understanding of the subject, so such homework can replace the final exam, if the originality of the solution leaves no doubt about the authorship. We will demonstrate solutions for four problems:

1. Find a net and two markings: M and M' such that
 - M is live
 - $M' \geq M$
 - M' is not live
2. Find a net and two markings: M and M' such that
 - M is live and bounded
 - $M' \geq M$
 - M' is live and not bounded
3. Find a net and initial marking M such that
 - M is live and bounded
 - the reachability graph of M has two finals (two maximal strongly connected components).
4. Find a net and initial marking M such that
 - M is live
 - there exist two places p_1, p_2 , such that they are both unbounded, but not simultaneously.

2. Definitions

Petri nets are triples $\langle P, T, F \rangle$, where P and T are finite, nonempty and disjoint sets of *places* and *transitions* respectively, while $F : P \times T \cup T \times P \rightarrow \mathbf{N}$ is a function defining the structure of a Petri net graph. We assume that if $F(x, y) = 0$, then there is no edge in the graph between nodes x and y . A positive value means that the edge connecting x and y exists and denotes the capacity of the edge,

✉ pch@mimuw.edu.pl (P. Chrzastowski-Wachtel); pu429640@students.mimuw.edu.pl (P. Ulanowski)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

so the number of tokens that will be passing along this arc and the direction of token flow. If for all $x, y : F(x, y) \leq 1$, we call such a net simple.

A net is a *P-system* if the numbers of incoming and outgoing edges of any transition do not exceed one. A net is a *T-system* if the numbers of incoming and outgoing edges of any place do not exceed one.

Let's fix a net $\mathcal{N} = \langle P, T, F \rangle$. A marking is a mapping from P to \mathbf{N} . A transition $t \in T$ is enabled at marking M , if for all $p \in P : M(p) \geq F(p, t)$. An enabled transition at a given marking M can *fire* changing M into M' according to the rule: for all $p \in P : M'(p) = M(p) - F(p, t) + F(t, p)$. Firing a transition t at M changing it to M' is denoted by $M[t]M'$. The reachability set of M is the smallest set of markings $[M]$ satisfying two conditions:

- $M \in [M]$
- if $M' \in [M]$ and $\exists t \in T$ and M'' such that $M'[t]M''$ then $M'' \in [M]$.

The *reachability graph* $G(M)$ of a marking M is a graph $\langle V, E \rangle$, where $V = [M]$ and $E = \{(M', M'') : \exists t \in T : M'[t]M''\}$; in this case we presume that the edge between M' and M'' is labeled by t . A subgraph (V', E') is called a *final* of (V, E) iff $V' \subseteq V, E' = E \cap (V' \times V')$ is a maximal strongly connected component, so any two nodes of V' are connected by a path in $G(M)$. If the set $[M]$ is finite, then finals represent possible ultimate behaviors of a net after making possibly irreversible decisions, during resolving conflicts with transition firings.

A place p at marking M is *bounded*, if there exists $k \in \mathbf{N}$ such that $\forall M' \in [M] : M'(p) \leq k$. Otherwise we call such a place *unbounded* at M . A marking M is bounded if all places are bounded at M (or, equivalently, the set $[M]$ is finite).

Marking M is *live* if $\forall t \in T : \forall M' \in [M] : \exists M'' \in [M'] : t$ is enabled at M'' . A marking M is *dead* if no transition is enabled at M .

A net is *well formed*, if there exists a live and bounded marking in it. A net is *structurally bounded* if all markings are bounded. When a net is unbounded, we construct *coverability graph* following the idea of Karp and Miller [KM69]. Whenever during the construction of the coverability graph a marking bigger than some previously considered node on a path from M is created, we replace all the coordinates, which are greater, by a symbol ω to indicate the possibility of generating arbitrarily many tokens. Additionally we introduce ω_o indicating the possibility of generating any odd number of tokens, and ω_e indicating the possibility of generating any even number of tokens.

For two markings M, M' we say that $M \leq M'$ iff $\forall p \in P : M(p) \leq M'(p)$. We call a property *monotonic*, if it is preserved under increasing a marking. Otherwise it is called non-monotonic.

We present now four problems with (subjectively) increasing difficulty. All of them seem at the first glance to have no solution, but solutions indeed exist.

3. Non-monotonicity of liveness

Find a net and two markings: M and M' such that

- M is live
- $M' \geq M$
- M' is not live

This is an old problem, proposed by the first author in the [PNB13] long ago, solved by many, including [WR85]. Here we present the most compact solution consisting of a net with just two transitions and one place. The marking of Figure 1 is live: invariantly the number of tokens on place p is odd, so one can always increase the number of them by firing t_1 , but if we increase the number of tokens in the initial marking to 2, the marking is not live: we can fire t_2 making a marking dead.

This problem has some important moral: When something goes well, adding some extra goods can destroy wellness. Like it happens in real life, for instance with people who win a fortune on a lottery give up their good jobs, lose friends and end up in disaster. Or in economics, when too much money can trigger unnecessary investments, killing some well working parts of a system.




Figure 1: A live marking of a net solving the "Non-monotonicity of liveness" problem.

4. Non-monotonicity of boundedness for live nets

Find a net and two markings: M and M' such that

- M is live and bounded
- $M' \geq M$
- M' is live and not bounded

This problem is much harder. In fact the first author has been puzzled with this problem by Teruel [ET92], while working together on T-systems and P-systems in 1992. The hypothesis of Teruel was that well-formed P-systems are structurally bounded, which would exclude the example satisfying these requirements. Surprisingly, one of the solutions for this problem is a P-system, which makes the hypothesis false.

The first author proposes this problem to the students, as stated, without a requirement that the net should be a P-system. Maybe this makes the problem even harder. Anyway, the example presented here on Figure 2 with only two places and four transitions is easily verifiable, as can be seen on the coverability graphs. Notice that every transition can be rollbacked with an opposite transition. To be precise, after firing t_1 the net will restore its marking by firing t_4 , and after firing t_2 , the net will restore its marking by firing t_3 and vice versa. This means that every marking that can activate t_1 and t_2 or t_3 and t_4 is live. This is because, after applying transitions in some order from such marking, you can always return to the original marking by applying the corresponding opposite transitions in the reverse order. This way, the net will restore a marking, from which it was possible to activate all the transitions at some point. Notice that marking $(0, 5)$ on Figure 2 can activate both t_3 and t_4 , so the net with unbounded marking is live.

Mind that such an example cannot be found for free-choice nets (a consequence of the rank theorem) [DE94].

Let us present another net on Figure 3 and Figure 4, which was discovered by the second author and will serve as a basis of the solution of the next problem.




Figure 2: A bounded marking of a P-System net solving the "Non-monotonicity of boundedness for live nets" problem, coverability graphs for bounded and unbounded markings




Figure 3: A bounded marking of a net solving the "Non-monotonicity of boundedness for live nets" problem, proposed by second author.




Figure 4: Coverability graph for not bounded marking * - Transitions with asterisk will fire only when $\omega_o = 1$

5. Two finals

Find a net and initial marking M such that

- M is live and bounded
- the coverability graph of M has two finals.

To find a net, whose coverability graph has two finals is a very easy task. But adding the requirement of liveness can make the problem pretty hard.

Recently the first author communicated Devillers, who mentioned that the problem was first solved probably by Best and Voss in [BV84].

The solution presented here was proposed by the second author. Mind that the initial marking is live and bounded, and depending on the choice of the first transition to fire, we fall into two different finals, which of course do not share any single marking. The presented reachability graph leaves no doubt that the problem is solved indeed: after the choice of the first transition to fire, the net falls into disjoint cycles.

To approach this problem, one can start with a net that is modeling the flow of two processes rivaling for an access to a critical section in a concurrent environment. One can add an additional restriction to ensure liveness and eliminate starvation by enforcing on the processes to enter the critical section in an interleaving sequence – one after the other. This scenario can be represented by the net in Figure 5 without the place a . Tokens in places i_1, i_2 mean accordingly that the process is idle. Place m is representing the binary semaphore. Tokens in cs_1 and cs_2 mean, that a specific process is in critical section. Token in e means that one of the processes just released the semaphore and wants to become idle.

With such a net, we can just add an asymmetry to the net with place a , which solves the given problem. The solution can be seen on Figure 5.

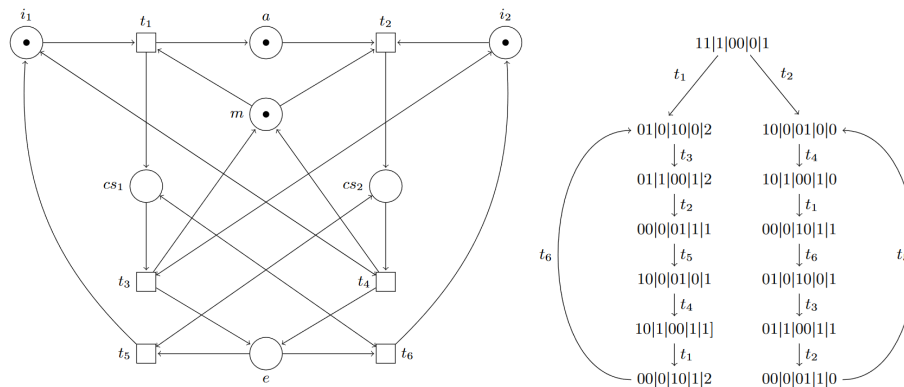


Figure 5: A marking of a net solving the "Two finals" problem and coverability graph of this net. Values are encoded accordingly $i_1 i_2 | m | cs_1 cs_2 | e | a$

6. Asynchronous unboundedness

Find a net and initial marking M such that

- M is live
- there exist two places p_1, p_2 , such that they are both unbounded, but not simultaneously. In other words: for every $k \in \mathbb{N}$ there exists a marking $M_1 \in [M]$ with $M_1(p_1) > k$, and there exists a marking $M_2 \in [M]$ with $M_2(p_2) > k$, but there exists a constant m such that for every $M' \in [M]$ either $M'(p_1) < m$ or $M'(p_2) < m$

The solution for this problem is presented on Figure 6. The net is composed of one net that was used for solving the "Two Finals" problem and two nets that were created by the second author for the "Non-monotonicity of boundedness for live nets" problem. By choosing non-deterministically the first transition t_1 or t_2 , the net will put two tokens either on a_1 or on a_2 accordingly, but it will be not possible to achieve two tokens on both these places simultaneously. As we showed earlier the upper nets with their a_i being marked with two tokens are unbounded causing the place p_i to achieve any number of tokens. All the three nets are live and taking/adding tokens from places a_i by other nets won't affect their liveness in a long run. As a result we've received a live net with places p_1 and p_2 that fit the requirements of the problem.




Figure 6: A marked net solving the "Asynchronous unboundedness" problem. Places p_1 and p_2 can gain any number of tokens, but not at the same time. The marking is live.

7. Conclusions

All the presented solutions expose some interesting aspects of net behaviors. Having learned such examples enriches understanding of different non-intuitive properties and interrelations between them. Here the solution of the last example was merely a combination of two previous solutions. Usually

students find the last problem being the most difficult.

It is also nice to see that some standard problems from the concurrency theory, like mutex, give an insight to non-trivial solutions of the presented Petri net problems.

References

- [PNB13] Piotr Chrzastowski-Wachtel, Puzzle No 3, Petri Nets and Related Models, Newsletter No 13, p22, 1983
- [DE94] J.Desel, J,Esparza, Free Choice Petri Nets, Cambridge University Press, 1995.
- [WR85] Wolfgang Reisig, Petri Nets, an Introduction, Springer Verlag 1985.
- [ET92] E. Teruel, private communication, 1992.
- [BV84] E.Best, K.Voss, Free Choice Systems Have Home States Acta Informatica 21, pp. 89–100, 1984.
- [KM69] R.Karp, R.Miller, Parallel program schemata. Computer and System Sciences, 1969