



On-the-Fly Synthesis for Strictly Alternating Games

Shyam Lal Karra, Kim Guldstrand Larsen, Marco Muñoz, and Jiří Srba^(✉)

Aalborg University, Aalborg, Denmark
{shyamla1, kgl, muniz, srba}@cs.aau.dk

Abstract. We study two-player zero-sum infinite reachability games with strictly alternating moves of the players allowing us to model a race between the two opponents. We develop an algorithm for deciding the winner of the game and suggest a notion of alternating simulation in order to speed up the computation of the winning strategy. The theory is applied to Petri net games, where the strictly alternating games are in general undecidable. We consider soft bounds on Petri net places in order to achieve decidability and implement the algorithms in our prototype tool. Finally, we compare the performance of our approach with an algorithm proposed in the seminal work by Liu and Smolka for calculating the minimum fixed points on dependency graphs. The results show that using alternating simulation almost always improves the performance in time and space and with exponential gain in some examples. Moreover, we show that there are Petri net games where our algorithm with alternating simulation terminates, whereas the algorithm without the alternating simulation loops for any possible search order.

1 Introduction

An embedded controller often has to interact continuously with an external environment and make decisions in order for the overall system to evolve in a safe manner. Such systems may be seen as (alternating) games, where two players—the controller and the environment—race against each other in order to achieve their individual objectives. The environment and controller alternate in making moves: the environment makes a move and gives the turn to the controller who can correct the behaviour of the system and give the control back to the environment and so on. We consider zero-sum turn-based games where the objective of the controller is to reach a set of goal states, while the objective of the environment is to avoid these states. Winning such a game requires to synthesize a strategy for the moves of the controller, so that any run under this strategy leads to a goal state no matter the moves of the environment. We talk about *synthesizing a winning controller strategy*.

Consider the game in Fig. 1. The game has six configurations $\{s_0, \dots, s_5\}$ and the solid edges indicate controller moves, while the dashed edges are environmental moves. The game starts at s_0 and the players alternate in taking turns,

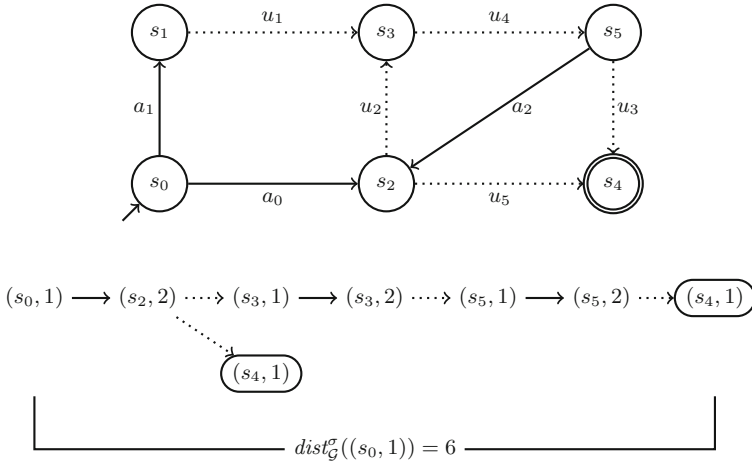


Fig. 1. A game graph and the distance of its winning strategy σ where $\sigma(s_0, 1) = a_0$, $\sigma(s_1, 1) = \sigma(s_2, 1) = \sigma(s_3, 1) = \sigma(s_4, 1) = \sigma(s_5, 1) = \epsilon_1$

assuming that the controller has to make the first move. It has three choices, i.e., move to s_1 or s_2 (shown by solid arrows) or stay in s_0 without moving. Assume that the controller chooses to move to s_2 and gives the control to environment. The environment now has the choice to go to s_3 or s_4 (which is a goal configuration that the environment tries to avoid). Observe that there can be states like s_5 from which both the environment and controller can make a move, based on whose turn it is. Also the game may have states like s_0 or s_3 where one of the players does not have any explicit move in which case they play an empty move and give the control to the other. The controller can play the empty move also in the situation where other controllable moves are enabled, whereas the environment is allowed to make the empty move only if no other environmental moves are possible (in order to guarantee progress in the game). The goal of the controller is to reach s_4 and it has a winning strategy to achieve this goal as shown below the game graph (here the second component in the pair denotes the player who has the turn).

Petri nets [9] is a standard formalism that models the behaviour of a concurrent system. Petri net games represent the reactive interaction between a controller and an environment by distinguishing controllable transitions from uncontrollable ones. In order to account for both the adversarial behaviour of the environment and concurrency, one can model reactive systems as alternating games (as mentioned before) played on these nets where each transition is designated to either environment or controller and the goal of the controller is to eventually reach a particular marking of the net.

Contribution: We define the notion of alternating simulation and prove that if a configuration c' simulates c then c' is a winning configuration whenever c is a winning configuration. We then provide an on-the-fly algorithm which uses

this alternating simulation relation and we prove that the algorithm is partially correct. The termination of the algorithm is in general not guaranteed. However, for finite game graphs we have total correctness. We apply this algorithm to games played on Petri nets and prove that these games are in general undecidable and therefore we consider a subclass of these games where the places in nets have a bounded capacity, resulting in a game over a finite graph. As an important contribution, we propose an efficiently computable (linear in the size of the net) alternating simulation relation for Petri nets. We also show an example where this specific simulation relation allows for a termination on an infinite Petri net game, while the adaptation of Liu-Smolka algorithm from [2] does not terminate for all possible search strategies that explore the game graph. Finally, we demonstrate the practical usability of our approach on three case studies.

Related Work: The notion of *dependency graphs* and their use to compute fixed points was originally proposed in [7]. In [2,3] an adaptation of these fixed point computations has been extended to two player games with reachability objectives, however, without the requirement on the alternation of the player moves as in our work. Two player coverability games on Petri nets with strictly alternating semantics were presented in [10] and [1], but these games are restricted in the sense that they assume that the moves in the game are monotonic with respect to a partial order which makes the coverability problem decidable for the particular subclass (B-Petri games). Our games are more general and hence the coverability problem for our games played on Petri nets, unlike the games in [10], are undecidable. Our work introduces an on-the-fly algorithm closely related to the work in [3] which is an adaptation of the classical Liu Smolka algorithm [7], where the set of losing configurations is stored, which along with the use of our alternating simulation makes our algorithm more efficient with respect to both space and time consumption. Our main novelty is the introduction of alternating simulation as a part of our algorithm in order to speed up the computation of the fixed point and we present its syntactic approximation for the class of Petri net games with encouraging experimental results.

2 Alternating Games

We shall now introduce the notion of an alternating game where the players strictly alternate in their moves such that the first player tries to enforce some given reachability objective and the other player's aim is to prevent this from happening.

Definition 1. A Game graph is a tuple $\mathcal{G} = (S, Act_1, Act_2, \longrightarrow_1, \longrightarrow_2, Goal)$ where:

- S is the set of states,
- Act_i is a finite set of Player- i actions where $i \in \{1, 2\}$ and $Act_1 \cap Act_2 = \emptyset$,
- $\longrightarrow_i \subseteq S \times Act_i \times S$, $i \in \{1, 2\}$ is the Player- i edge relation such that the relations \longrightarrow_i are deterministic i.e. if $(s, \alpha, s') \in \longrightarrow_i$ and $(s, \alpha, s'') \in \longrightarrow_i$ then $s' = s''$, and
- $Goal \subseteq S$ is a set of goal states.

If $(s, a, s') \in \longrightarrow_i$ then we write $s \xrightarrow{a}_i s'$ where $i \in \{1, 2\}$. Also, we write $s \longrightarrow_i$ if there exists an $s' \in S$ and $a \in Act_i$ such that $s \xrightarrow{a}_i s'$; otherwise we write $s \not\longrightarrow_i$. In the rest of this paper, we restrict ourselves to game graphs that are *finitely branching*, meaning that for each action a (of either player) there are only finitely many a -successors.

Alternating Semantics for Game Graphs. Given a game graph $\mathcal{G} = (S, Act_1, Act_2, \longrightarrow_1, \longrightarrow_2, Goal)$, the set $C = (S \times \{1, 2\})$ is the set of configurations of the game graph and $C_1 = (S \times \{1\}), C_2 = (S \times \{2\})$ represent the set of configurations of *Player-1* and *Player-2* respectively, $\Longrightarrow \subseteq (C_1 \times (Act_1 \cup \{\epsilon_1\}) \times C_2) \cup (C_2 \times (Act_2 \cup \{\epsilon_2\}) \times C_1)$ is the set of transitions of the game such that

- if $(s, \alpha, s') \in \longrightarrow_i$ then $((s, i), \alpha, (s', 3 - i)) \in \Longrightarrow$
- for all configurations $(s, 1) \in C_1$ we have $((s, 1), \epsilon_1, (s, 2)) \in \Longrightarrow$
- for all configurations $(s, 2) \in C_2$ where $s \not\longrightarrow_2$ we have $((s, 2), \epsilon_2, (s, 1)) \in \Longrightarrow$

Apart from the actions available in Act_1 and Act_2 , the controller and the environment have two empty actions ϵ_1 and ϵ_2 , respectively. Given a state $s \in S$, when it is its turn, the controller or the environment can choose to take any action in Act_1 or Act_2 enabled at s respectively. While, the controller can always play ϵ_1 at a given state s , the environment can play ϵ_2 only when there are no actions in Act_2 enabled at s so that the environment can not delay or deadlock the game.

We write $(s, i) \xrightarrow{\alpha}_i (s', 3 - i)$ to denote $((s, i), \alpha, (s', 3 - i)) \in \Longrightarrow$ and $(s, i) \xrightarrow{\alpha}_i$ if there exists $s' \in S$ such that $(s, i) \xrightarrow{\alpha}_i (s', 3 - i)$. A configuration $c_g = (s_g, i)$ is called *goal configuration* if $s_g \in Goal$. By abuse of notation, for configuration c we use $c \in Goal$ to denote that c is a goal configuration. A run in the game \mathcal{G} starting at a configuration c_0 is an infinite sequence of configurations and actions $c_0, \alpha_0, c_1, \alpha_1, c_2, \alpha_2 \dots$ where for every $j \in \mathbb{N}^0$, $c_j \xrightarrow{\alpha_j} c_{j+1}$. The set $Runs(c_0, \mathcal{G})$ denotes the set of all runs in the game \mathcal{G} starting from configuration c_0 . A run $c_0, \alpha_0, c_1, \alpha_1, c_2, \alpha_2 \dots$ is *winning* if there exists j such that c_j is a goal configuration. The set of all winning runs starting from c_0 is denoted by $WinRuns(c_0, \mathcal{G})$. A *strategy* σ for *Player-1* where $\sigma : C_1 \rightarrow Act_1 \cup \{\epsilon_1\}$ is a function such that if $\sigma(c) = \alpha$ then $c \xrightarrow{\alpha}_1$. For a given strategy σ , and a configuration $c_0 = (s_0, i_0)$, we define $OutcomeRuns_{\mathcal{G}}^{\sigma}(c_0) = \{c_0, \alpha_0, c_1, \alpha_1 \dots \in Runs(c_0, \mathcal{G}) \mid \text{for every } k \in \mathbb{N}^0, \alpha_{2k+i_0-1} = \sigma(c_{2k+i_0-1})\}$. A strategy σ is *winning* at a configuration c if $OutcomeRuns_{\mathcal{G}}^{\sigma}(c) \subseteq WinRuns(c, \mathcal{G})$. A configuration c is a *winning configuration* if there is a winning strategy at c . A winning strategy for our running example is given in the caption of Fig. 1.

Given a configuration $c = (s, i)$ and one of its *winning* strategies σ , we define the quantity $dist_{\mathcal{G}}^{\sigma}(c) := \max(\{n \mid c_0, \alpha_0 \dots c_n \dots \in OutcomeRuns_{\mathcal{G}}^{\sigma}(c) \text{ with } c = c_0 \text{ and } c_n \in Goal \text{ while } c_i \notin Goal \text{ for all } i < n\})$. The quantity represents the *distance* to the goal state, meaning that during any play by the given strategy there is a guarantee that the goal state is reached within that many steps. Due

to our assumption on finite branching of our game graph, we get that distance is well defined.

Lemma 1. *Let σ be a winning strategy for a configuration c . The distance function $\text{dist}_{\mathcal{G}}^{\sigma}(c)$ is well defined.*

Proof. Should $\text{dist}_{\mathcal{G}}^{\sigma}(c)$ not be well defined, meaning that the maximum does not exist, then necessarily the set $T = \{c_0, \alpha_0 \dots c_n \mid c_0, \alpha_0 \dots c_n \dots \in \text{OutcomeRuns}_{\mathcal{G}}^{\sigma}(c) \text{ where } c = c_0 \text{ and } c_n \in \text{Goal} \text{ with } c_i \notin \text{Goal} \text{ for all } i < n\}$ induces an infinite tree. By definition the game graph \mathcal{G} is a deterministic transition system with finite number of actions and hence the branching factor of \mathcal{G} and consequently of T is finite. By Kőning's Lemma the infinite tree T must contain an infinite path without any goal configuration, contradicting the fact that σ is a winning strategy. \square

Consider again the game graph shown in Fig. 1 where solid arrows indicate the transitions of *Player-1* while the dotted arrows indicate those of *Player-2*. A possible tree of all runs under a (winning) strategy σ is depicted in the figure and its distance is 6. We can also notice that the distance strictly decreases once *Player-1* performs a move according to the winning strategy σ , as well as by any possible move of *Player-2*. We can now observe a simple fact about alternating games.

Lemma 2. *If $(s, 2)$ is a winning configuration in the game \mathcal{G} , then $(s, 1)$ is also a winning configuration.*

Proof. Since $(s_0, 1) \xrightarrow{e_1} (s_0, 2)$ and there is a winning strategy from $(s_0, 2)$, we can conclude $(s_0, 1)$ is also a winning configuration. \square

We shall be interested in finding an efficient algorithm for deciding the following problem.

Definition 2 (Reachability Control Problem). *For a given game \mathcal{G} and a configuration (s, i) , the reachability control problem is to decide if (s, i) is a winning configuration.*

3 Alternating Simulation and On-the-Fly Algorithm

We shall now present the notion of alternating simulation that will be used as the main component of our on-the-fly algorithm for determining the winner in the alternating game. Let $\mathcal{G} = (S, \text{Act}_1, \text{Act}_2, \longrightarrow_1, \longrightarrow_2, \text{Goal})$ be a game graph and let us adopt the notation used in the previous section.

Definition 3. *A reflexive binary relation $\preceq \subseteq C_1 \times C_1 \cup C_2 \times C_2$ is an alternating simulation relation iff whenever $(s_1, i) \preceq (s_2, i)$ then*

- if $s_1 \in \text{Goal}$ then $s_2 \in \text{Goal}$,
- if $(s_1, 1) \xrightarrow{a}_1 (s'_1, 2)$ then $(s_2, 1) \xrightarrow{a}_1 (s'_2, 2)$ such that $(s'_1, 2) \preceq (s'_2, 2)$, and
- if $(s_2, 2) \xrightarrow{u}_2 (s'_2, 1)$ then $(s_1, 2) \xrightarrow{u}_2 (s'_1, 1)$ such that $(s'_1, 1) \preceq (s'_2, 1)$.

An important property of alternating simulation is that it preserves winning strategies as stated in the following theorem.

Theorem 1. *Let (s, i) be a winning configuration and $(s, i) \preceq (s', i)$ then (s', i) is also a winning configuration.*

Proof. By induction on k we shall prove the following claim: if $(s, i) \preceq (s', i)$ and (s, i) is a winning configuration with $\text{dist}_{\mathcal{G}}^{\sigma}((s, i)) \leq k$ then (s', i) is also a winning configuration.

Base case ($k = 0$): Necessarily (s, i) is a goal state and by the definition of alternating simulation (s', i) is also a goal configuration and hence the claim trivially holds.

Induction step ($k > 0$): Case $i = 1$. Let σ be a winning strategy for $(s, 1)$ such that $(s, 1) \xrightarrow{\sigma((s,1))}_1 (s_1, 2)$. We define a winning strategy σ' for $(s', 1)$ by $\sigma'((s', 1)) = \sigma((s, 1))$. By the property of alternating simulation we get that $(s', 1) \xrightarrow{\sigma((s,1))}_1 (s'_1, 2)$ such that $(s_1, 2) \preceq (s'_1, 2)$ and $\text{dist}_{\mathcal{G}}^{\sigma}((s_1, 2)) < k$. Hence we can apply induction hypothesis and claim that $(s'_1, 2)$ has a winning strategy which implies $(s', 1)$ is a winning configuration as well. Case $i = 2$. If $(s', 2) \xrightarrow{u}_2 (s'_1, 1)$ for some $u \in \text{Act}_2$ then by the property of alternating simulation also $(s, 2) \xrightarrow{u}_2 (s_1, 1)$ such that $(s_1, 1) \preceq (s'_1, 1)$ and $\text{dist}_{\mathcal{G}}^{\sigma}((s_1, 2)) < k$. By the induction hypothesis $(s'_1, 1)$ is a winning configuration and so is $(s', 2)$. \square

As a direct corollary of this result, we also know that if $(s, i) \preceq (s', i)$ and (s', i) is not a winning configuration then (s, i) cannot be a winning configuration either. Before we present our on-the-fly algorithm, we need to settle some notation. Given a configuration $c = (s, i)$ we define

- $\text{Succ}(c) = \{c' \mid c \Longrightarrow c'\}$,
- $\text{MaxSucc}(c) \subseteq \text{Succ}(c)$ is a set of states such that for all $c' \in \text{Succ}(c)$ there exists a $c'' \in \text{MaxSucc}(c)$ such that $c' \preceq c''$, and
- $\text{MinSucc}(c) \subseteq \text{Succ}(c)$ is a set of states such that for all $c' \in \text{Succ}(c)$ there exists a $c'' \in \text{MinSucc}(c)$ such that $c'' \preceq c'$.

Remark 1. There can be many candidate sets of states that satisfy the definition of $\text{MaxSucc}(c)$ (and $\text{MinSucc}(c)$). In the rest of the paper, we assume that there is a way to fix one among these candidates and the theory proposed here works for any such candidate.

Given a game graph \mathcal{G} , we define a subgraph of \mathcal{G} denoted by \mathcal{G}' , where for all *Player-1* successors we only keep the maximum ones and for every *Player-2* successors we preserve only the minimum ones. In the following lemma we show that in order to solve the reachability analysis problem on \mathcal{G} , it is sufficient to solve it on \mathcal{G}' .

Definition 4. Given a game graph $\mathcal{G} = (S, Act_1, Act_2, \longrightarrow_1, \longrightarrow_2, Goal)$, we define a pruned game graph $\mathcal{G}' = (S, Act_1, Act_2, \longrightarrow'_1, \longrightarrow'_2, Goal)$ such that $\longrightarrow'_1 = \{(s_1, 1, s_2) \mid (s_1, 1, s_2) \in \longrightarrow_1 \text{ where } (s_2, 2) \in MaxSucc(s_1, 1)\}$ and $\longrightarrow'_2 = \{(s_2, 2, s_1) \mid (s_2, 2, s_1) \in \longrightarrow_2 \text{ where } (s_1, 1) \in MinSucc(s_2, 2)\}$.

Lemma 3. Given a game graph \mathcal{G} , a configuration (s, i) is winning in \mathcal{G} iff (s, i) is winning in \mathcal{G}' .

Proof. We prove the case for $i = 1$ (the argument for $i = 2$ is similar).

“ \implies ”: Let $(s, 1)$ be winning in \mathcal{G} under a winning strategy σ . We define a new strategy σ' in \mathcal{G}' such that for every $(s_1, 1)$ where $\sigma((s_1, 1)) = (s'_1, 2)$ we define $\sigma'((s_1, 1)) = (s''_1, 2)$ for some $(s''_1, 2) \in MaxSucc(s_1, 1)$ such that $(s'_1, 2) \preceq (s''_1, 2)$. By induction on $dist_{\mathcal{G}}^{\sigma}((s, 1))$ we prove that if $(s, 1)$ is winning in \mathcal{G} then it is winning also in \mathcal{G}' . Base case ($dist_{\mathcal{G}}^{\sigma}((s, 1)) = 0$) clearly holds as $(s, 1)$ is a goal configuration. Let $dist_{\mathcal{G}}^{\sigma}((s, 1)) = k$ where $k > 0$ and let $\sigma((s, 1)) = (s', 2)$ and $\sigma'((s, 1)) = (s'', 2)$. Clearly, $dist_{\mathcal{G}}^{\sigma}((s', 2)) < k$ and by induction hypothesis $(s', 2)$ is winning also in \mathcal{G}' . Because $(s', 2) \preceq (s'', 2)$ we get by Theorem 1 that $(s'', 2)$ is also winning and $dist_{\mathcal{G}}^{\sigma}((s'', 2)) < k$. Since $s \xrightarrow{\sigma'(s, 1)}_1 s''$ we get that σ' is a winning strategy for $(s, 1)$ in \mathcal{G}' .

“ \impliedby ”: Let $(s, 1)$ be winning in \mathcal{G}' by strategy σ' . We show that $(s, 1)$ is also winning in \mathcal{G} under the same strategy σ' . We prove this fact by induction on $dist_{\mathcal{G}'}^{\sigma'}((s, 1))$. If $dist_{\mathcal{G}'}^{\sigma'}((s, 1)) = 0$ then $(s, 1)$ is a goal configuration and the claim follows. Let $dist_{\mathcal{G}'}^{\sigma'}((s, 1)) = k$ where $k > 0$ and let $\sigma'(s, 1) = (s'', 2)$. Clearly $dist_{\mathcal{G}'}^{\sigma'}((s'', 2)) < k$. So by induction hypothesis, $(s'', 2)$ is a winning configuration in \mathcal{G} . Since $s \xrightarrow{\sigma'(s, 1)}_1 s''$ we get that σ' is a winning strategy for $(s, 1)$ in \mathcal{G} . \square

We can now present an algorithm for the reachability problem in alternating games which takes some alternating simulation relation as a parameter, along with the game graph and the initial configuration. In particular, if we initialize \preceq to the identity relation, the algorithm results essentially (modulo the additional *Lose* set construction) as an adaption of Liu-Smolka fixed point computation as given in [3]. Our aim is though to employ some more interesting alternating simulation that is fast to compute (preferably in syntax-driven manner as we demonstrate in the next section for the case of Petri net games). Our algorithm uses this alternating simulation at multiple instances which improves the efficiency of deciding the winning and losing status of configurations without having to explore unnecessary state space.

The algorithm uses the following data structures:

- W is the set of edges (in the alternating semantics of game graph) that are *waiting* to be processed,
- $Disc$ is the set of already *discovered* configurations,
- Win and $Lose$ are the sets of currently known *winning* resp. *losing* configurations, and
- D is a *dependency* function that to each configuration assigns the set of edges to be reinserted to the waiting set W whenever the configuration is moved to the set Win or $Lose$ by the help functions $AddToWin$ resp. $AddToLose$.

Algorithm 1. Game graph algorithm

Input Game graph $\mathcal{G} = (S, Act_1, Act_2, \rightarrow_1, \rightarrow_2, G)$, alternating simulation \preceq and the initial configuration (s_0, i_0)

Output true if (s_0, i_0) is a winning configuration, false otherwise

```

1:  $W = \emptyset$ ;  $Disc := \{(s_0, i_0)\}$ ;  $Lose := \emptyset$ ;  $Win := \emptyset$ ;  $D((s_0, i_0)) = \emptyset$ ;
2: If  $s_0 \in G$  then ADDTOWIN( $(s_0, i_0)$ );
3: ADDSUCCESSORS( $W, s_0, i_0$ );
4: while (
5:    $\nexists (s'_0, i_0). (s'_0, i_0) \preceq (s_0, i_0) \wedge (s'_0, i_0) \in Win$  and
6:    $\nexists (s'_0, i_0). (s_0, i_0) \preceq (s'_0, i_0) \wedge (s'_0, i_0) \in Lose$  and
7:    $W \neq \emptyset$ 
8: ) do
9:   Pick  $(s_1, i, s_2) \in W$ ;  $W := W \setminus \{(s_1, i, s_2)\}$ ;
10:  if  $(s_1, i) \in Win \cup Lose$  then Continue at line 4;
11:  if (
12:     $\exists s'_1. (s_1, i) \preceq (s'_1, i) \wedge (s'_1, i) \in Lose$  or
13:     $i = 1 \wedge \forall s_3. s_1 \xrightarrow{1} s_3$  implies  $\exists s'_3. (s_3, 2) \preceq (s'_3, 2) \wedge (s'_3, 2) \in Lose$  or
14:     $i = 2 \wedge \exists s_3. s'_3. (s_3, 1) \preceq (s'_3, 1) \wedge s_1 \xrightarrow{2} s_3 \wedge (s'_3, 1) \in Lose$  or
15:     $i = 2 \wedge \exists s'_1. (s_1, 1) \preceq (s'_1, 1) \wedge (s'_1, 1) \in Lose$  or
16:     $s_1 \not\xrightarrow{1} \wedge s_1 \not\xrightarrow{2}$ 
17:  ) then
18:    ADDTOLOSE( $W, s_1, i$ ); Continue at line 4;
19:  end if
20:  if (
21:     $\exists s'_1. (s'_1, i) \preceq (s_1, i) \wedge (s'_1, i) \in Win$  or
22:     $i = 1 \wedge \exists s_3. s'_3. (s_1, 1) \xrightarrow{1} (s_3, 2) \wedge (s'_3, 2) \in Win$  or
23:     $i = 2 \wedge \forall s_3. (s_1, 2) \xrightarrow{2} (s_3, 1)$  implies  $\exists s'_3. (s'_3, 1) \preceq (s_3, 1) \wedge (s'_3, 1) \in Win$ 
24:  ) then
25:    ADDTOWIN( $W, s_1, i$ ); Continue at line 4;
26:  end if
27:  if  $(s_2, 3 - i) \notin Win \cup Lose$  then
28:    if  $(s_2, 3 - i) \in Disc$  then
29:       $D((s_2, 3 - i)) := D((s_2, 3 - i)) \cup \{(s_1, i, s_2)\}$ ;
30:    else
31:       $Disc := Disc \cup \{(s_2, 3 - i)\}$ ;
32:       $D((s_2, 3 - i)) := \{(s_1, i, s_2)\}$ ;
33:      if  $s_2 \in G$  then
34:        ADDTOWIN( $W, s_2, 3 - i$ );
35:      else
36:        ADDSUCCESSORS( $W, s_2, 3 - i$ );
37:      end if
38:    end if
39:  end if
40: end while
41: if  $\exists s'_0. s'_0 \preceq s_0 \wedge (s'_0, i_0) \in Win$  then return true
42: else return false
43: end if

```

Algorithm 2. Helper procedures for Algorithm 1

```

1: procedure ADDSUCCESSORS( $W, s, i$ )
2:   if ( $i = 1$ ) then
3:      $L := \{(s, i, s') \mid \text{such that } (s', 3 - i) \in \text{MaxSucc}(s, i)\};$ 
4:   end if
5:   if ( $i = 2$ ) then
6:      $L := \{(s, i, s') \mid \text{such that } (s', 3 - i) \in \text{MinSucc}(s, i)\};$ 
7:   end if
8:    $W := W \cup L$ 
9: end procedure
10: procedure ADDTOWIN( $W, s, i$ )
11:    $\text{Win} := \text{Win} \cup \{(s, i)\};$ 
12:    $W := W \cup D(s, i) \setminus \{(s'', 3 - i, s) \mid (s'', 3 - i) \in (\text{Win} \cup \text{Lose})\};$ 
13: end procedure
14: procedure ADDTOLOSE( $W, s, i$ )
15:    $\text{Lose} := \text{Lose} \cup \{(s, i)\};$ 
16:    $W := W \cup D(s, i) \setminus \{(s'', 3 - i, s) \mid (s'', 3 - i) \in (\text{Win} \cup \text{Lose})\};$ 
17: end procedure

```

As long as the waiting set W is nonempty and no conclusion about the initial configuration can be drawn, we remove an edge from the waiting set and check whether the source configuration of the edge can be added to the losing or winning set. After this, the target configuration of the edge is explored. If it is already discovered, we only update the dependencies. Otherwise, we also check if it is a goal configuration (and call `AddToWin` if this is the case), or we add the outgoing edges from the configuration to the waiting set.

In order to argue about the correctness of the algorithm, we introduce a number of loop invariants for the while loop in Algorithm 1 in order to argue that if the algorithm terminates then (s_0, turn) is winning if and only if $(s_0, \text{turn}) \in \text{Win}$.

Lemma 4. *Loop Invariant 1 for Algorithm 1: If $(s, i) \in \text{Win}$ and $(s, i) \preceq (s', i)$ then (s', i) is a winning configuration.*

Proof Sketch. Initially, $\text{Win} = \emptyset$ and the invariant holds trivially before the loop is entered. Let us assume that the invariant holds before we execute the body of the while loop and we want to argue that after the body is executed the invariant still holds. During the current iteration, a new element can be added to the set Win at lines 25 or 34. If line 25 is executed then at least one of the conditions at lines 21, 22, 23 must hold. We argue in each of these cases, along with the cases in which line 34 is executed, that the invariant continues to hold after executing the call to the function `AddToWin`. \square

Lemma 5. *Loop Invariant 2 for Algorithm 1: If $(s, i) \in \text{Lose}$ and $s' \preceq s$ then (s', i) is not a winning configuration.*

Proof Sketch. Initially $Lose = \emptyset$ and loop invariant 2 holds trivially before the while loop at line 4 is entered. Similarly to the previous lemma, we argue that in all the cases where the function $AddToLose$ is called, the invariant continues to hold after the call as well. \square

The next invariant is essential for the correctness proof.

Lemma 6. *Loop Invariant 3 for Algorithm 1: During the execution of the algorithm for any $(s, j) \in Disc \setminus (Win \cup Lose)$ invariantly holds*

(a) if $j = 1$ then for every $(s', 2) \in MaxSucc(s, 1)$

- I. $(s, 1, s') \in W$ or
- II. $(s, 1, s') \in D(s', 2)$ and $(s', 2) \in Disc \setminus Win$ or
- III. $(s', 2) \in Lose$

(b) if $j = 2$ then for every $(s', 1) \in MinSucc(s, 2)$

- I. $(s, 2, s') \in W$ or
- II. $(s, 2, s') \in D(s', 1)$ and $(s', 1) \in Disc \setminus Lose$ or
- III. $(s', 1) \in Win$.

This third invariant claims that for any discovered but yet undetermined configuration (s, j) and for any of its outgoing edges (maximum ones in case it is *Player-1* move and minimum ones for *Player-2* moves), the edge is either on the waiting list, or the target configuration is determined as winning or losing, or otherwise the edge is in the dependency set of the target configuration, so that in case the target configuration is later on determined as winning or losing, the edge gets reinserted to the waiting set and the information can possibly propagate to the parent configuration of the edge. This invariant is crucial in the proof of partial correctness of the algorithm which is given below.

Theorem 2. *If the algorithm terminates and returns true then the configuration (s_0, i) is winning and if the algorithm upon termination returns false then the configuration (s_0, i) is losing.*

Proof. Upon termination, if the algorithm returns true (this can happen only at the Line 41), then it means there exists a $s'_0 \in S$ such that $(s'_0, i) \in Win$ and $(s'_0, i) \preceq (s_0, i)$. From Lemma 4, we can deduce that (s_0, i) is winning. On the other hand if the algorithm returns false, then there are two cases.

Case 1: There exists a $s'_0 \in S$ such that $(s_0, i) \preceq (s'_0, i)$ and $(s'_0, i) \in Lose$. From Lemma 5, we can deduce that (s_0, i) does not have a winning strategy.

Case 2: $W = \emptyset$ and there exists no $s'_0 \in S$ such that $(s'_0, i) \in Win$ and $(s'_0, i) \preceq (s_0, i)$ and there exists no $s'_0 \in S$ such that $(s'_0, i) \in Lose$ and $(s_0, i) \preceq (s'_0, i)$. We prove that (s_0, i) is not winning for *Player-1* by contradiction. Let (s_0, i) be a winning configuration. Let σ be a *Player-1* winning strategy. Let $L = Disc \setminus (Win \cup Lose)$. Now we make the following two observations.

From Lemma 6, we can deduce that for all $(s, 1) \in L$, if $(s, 1) \implies_1 (s_1, 2)$ then there exists an s'_1 such that $(s_1, 2) \preceq (s'_1, 2)$ and $(s'_1, 2) \in L$ or $(s'_1, 2) \in Lose$.

Also, we can deduce that the case that for all $(s_2, 1) \in Succ(s, 2)$ we have $(s_2, 1) \in Win$ is not possible as follows. Among all $(s_2, 1) \in Succ(s, 2)$ consider the last $(s_2, 1)$ entered the set Win . During this process in the call to function `AddToWin`, the edge $(s, 2, s_2)$ is added to the waiting list. Since by the end of the algorithm, this edge is processed and it would have resulted in adding $(s, 2)$ to Win because of the condition at line 23. But this is a contradiction as $(s, 2) \in L$ and $L \cap Win = \emptyset$. Hence for all $(s, 2) \in L$, there exists an s_2 such that $(s, 2) \implies_2 (s_2, 1)$ and $(s_2, 1) \in L$.

Given these two observations, and the definition of \preceq , one can show that *Player-2* can play the game such that for the resulting run $\rho \in OutcomeRuns_{\mathcal{G}}^{\sigma}(s_0)$ where $\rho = \langle (s_0, i), a_0, (s_1, 3 - i), a_1 \dots \rangle$, there exists a $\rho' = \langle (s'_0, i), b_0, (s'_1, 3 - i), b_1 \dots \rangle \in Runs(s_0, \mathcal{G})$ such that $s'_0 = s_0$, for all $k \in \mathbb{N}$, $(s_k, 2 - ((i + k) \% 2)) \preceq (s'_k, 2 - ((i + k) \% 2))$ and $(s'_k, 2 - ((i + k) \% 2)) \in L$. In other words the configurations in the run ρ are all (one by one) simulated by the corresponding configurations in the run ρ' and the configurations from the run ρ' moreover all belong to the set L . Since σ is a winning strategy, there must exist an index $j \in \{0, 1, 2, \dots\}$ such that $s_j \in G$. Since the set of goal states are upward closed, it also means $s'_j \in G$. But $L \cap \{(s, 1), (s, 2) \mid s \in G\} = \emptyset$ because the line 34 of the algorithm adds a configuration in $\{(s, 1), (s, 2) \mid s \in G\}$ to Win when it enters the set $Disc$ for the first time. Hence our assumption that (s_0, i) has a winning strategy is contradicted. \square

The algorithm does not necessarily terminate on general game graphs, however, termination is guaranteed on finite game graphs.

Theorem 3. *Algorithm 1 terminates on finite game graphs.*

Proof. In the following, we shall prove that any edge $e = (s_1, i, s_2)$ can be added to W at most twice and since there are only finitely many edges and every iteration of the while loop removes an edge from W (at line 9), W eventually becomes empty and the algorithm terminates on finite game graphs.

During the execution of while loop, an edge can only be added W through the call to functions `AddSuccessors` at line 36, `AddToWin` at lines 25, 34, or `AddToLose` at line 18. We shall show that these three functions can add an edge e at most twice to waiting list W .

Let $e = (s_1, i, s_2)$ be an edge added to W in iteration k through a call to `AddToWin` at line 34. This implies that, during iteration k , the condition in line 27 is true. Hence $(s_2, 3 - i) \notin Win \cup Lose$ before iteration k is executed and after line 34 is executed, $(s_2, 3 - i) \in Win \cup Lose$ (hence the condition in line 27 is now false). So the call to function `AddToWin` at line 34 can not add e to W after iteration k .

Let $e = (s_1, i, s_2)$ be an edge added to W in iteration k through a call to `AddToWin` at line 25. This implies that, during iteration k , the condition in line 10 is not true. Hence $(s_2, 3 - i) \notin Win \cup Lose$ before iteration k is executed and after line 25 is executed, $(s_2, 3 - i) \in Win \cup Lose$ (hence the condition in line 10 is false). So the call to function `AddToWin` at line 25 can not add e to W after iteration k .




Fig. 2. A Petri net game

Similar to previous cases, we can argue that the call to `AddToLose` at line 18 can add e to W atmost once. Also it is easy to observe that once $(s_2, 3 - i)$ is added to set $Win \cup Lose$ by any of the lines 25, 34, 18, the other two can not add e to W . So, all together, all these three lines can add e to W atmost once.

Now consider case when $e = (s_1, i, s_2)$ is added to W in iteration k through a call to `AddSuccessors` at line 36. This implies, during iteration k , the condition in the line 28 is not true. Hence $(s_1, i) \notin Disc$ before the iteration k is executed and after the line 31 is executed $(s_1, i) \in Disc$ by the end of iteration k and the condition in the line 28 is true . So the call to function `AddSuccessors` at line 36 can add e to W atmost once. In total, edge e can enter W atmost twice. \square

4 Application to Petri Games

We are now ready to instantiate our general framework to the case of Petri net games where the transitions are partitioned between the controller and environment transitions and we moreover consider a soft bound for the number of tokens in places (truncating the number of tokens that exceed the bound).

Definition 5 (Petri Net Games). A Petri Net Game $\mathcal{N} = (P, T_1, T_2, W, B, \varphi)$ is a tuple where

- P is a finite set of places,
- T_1 and T_2 are finite sets of transitions such that $T_1 \cap T_2 = \emptyset$,
- $W : (P \times (T_1 \cup T_2)) \cup ((T_1 \cup T_2) \times P) \rightarrow \mathbb{N}^0$ is the weight function,
- $B : P \rightarrow \mathbb{N}^0 \cup \{\infty\}$ is a function which assigns a (possible infinite) soft bound to every $p \in P$, and
- φ is a formula in coverability logic given by the grammar $\varphi ::= \varphi_1 \wedge \varphi_2 \mid p \geq c$ where $p \in P$ and $c \in \mathbb{N}^0$.

A marking is a function $M : P \rightarrow \mathbb{N}^0$ such that $M(p) \leq B(p)$ for all $p \in P$. The set $\mathcal{M}(\mathcal{N})$ is the set of all markings of \mathcal{N} .

In Fig. 2 we show an example of a Petri game where $P = \{p_0, \dots, p_3\}$ are the places, there are three controller transitions $T_1 = \{t_2, t_3, t_4\}$ and two environment transitions (indicated with a circle inside) $T_2 = \{t_0, t_1\}$. The function W assigns each arc a weight of 1, except the arcs from p_1 to t_2 and t_4 to p_0 that have weight 2. The bound function B assigns every place ∞ and $\varphi ::= p_3 \geq 1$ requires that in the goal marking the place p_3 must have at least one token. The initial marking of the net is $\langle 1, 0, 0, 0 \rangle$ where the place p_0 has one token and the places p_1 to p_3 (in this order) have no tokens.

For the rest of this section, let $\mathcal{N} = (P, T_1, T_2, W, B, \varphi)$ be a fixed Petri net game. The satisfaction relation for a coverability formula φ in marking M is defined as expected:

- $M \models p \geq n$ iff $M(p) \geq n$ where $n \in \mathbb{N}^0$, and
- $M \models \varphi_1 \wedge \varphi_2$ iff $M \models \varphi_1$ and $M \models \varphi_2$.

For a given formula φ , we now define the set of goal markings $\mathcal{M}_\varphi = \{M \in \mathcal{M}(\mathcal{N}) \mid M \models \varphi\}$.

Firing a transition $t \in T_1 \cup T_2$ from a marking M results in marking M' if for all $p \in P$ we have $M(p) \geq W((p, t))$ and

$$M'(p) = \begin{cases} M(p) - W((p, t)) + W((t, p)) & \text{if } M(p) - W((p, t)) + W((t, p)) \leq B(p) \\ B(p) & \text{otherwise.} \end{cases}$$

We denote this by $M \xrightarrow{t} M'$ and note that this is a standard Petri net firing transition in case the soft bound for each place is infinity, otherwise if the number of tokens in a place p of the resulting marking exceeds the integer bound $B(p)$ then it is truncated to $B(p)$.

Petri net game \mathcal{N} induces a game graph $(\mathcal{M}(\mathcal{N}), T_1, T_2, \longrightarrow_1, \longrightarrow_2, \mathcal{M}_\varphi)$ where $\longrightarrow_i = \{(M, t, M') \mid M \xrightarrow{t} M' \text{ and } t \in T_i\}$ for $i \in \{1, 2\}$.

For example, the (infinite) game graph induced by Petri game from Fig. 2 is shown in Fig. 3a. Next we show that reachability control problem for game graphs induced by Petri games is in general undecidable.

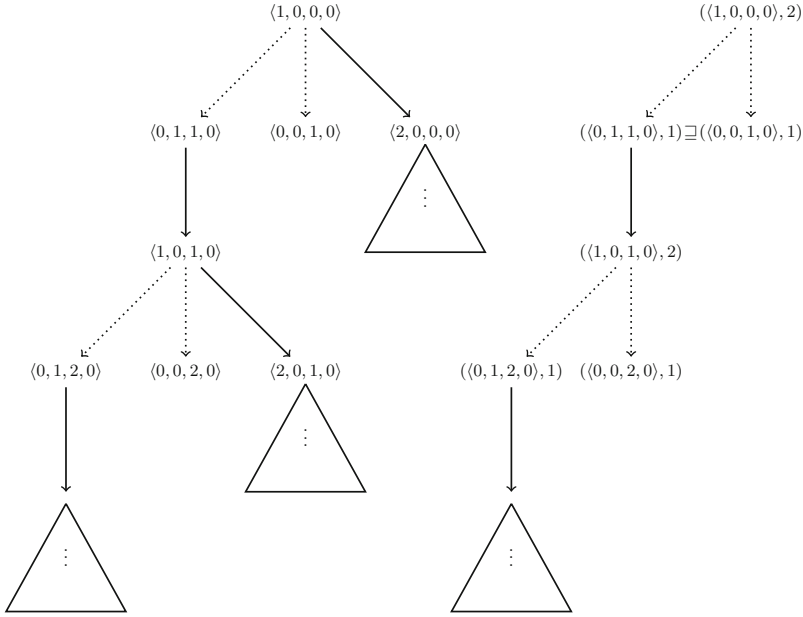
Theorem 4. *The reachability control problem for Petri games is undecidable.*

Proof. In order to show that reachability control problem for Petri games is undecidable, we reduce the undecidable problem of reachability for two counter Minsky machine [6] to the reachability control problem.

A Minsky Counter Machine with two non-negative counters c_1 and c_2 is a sequence of labelled instructions of the form:

$$1 : instr_1, 2 : instr_2, \dots n : HALT \tag{1}$$

where for all $i \in \{1, 2, \dots, n-1\}$ each $instr_i$ is of the form



(a) Game graph for PN from Figure 2 (b) Alternating semantics for Figure 3a

Fig. 3. Game graph and its alternating semantics

- $instr_i: c_r = c_r + 1; \text{ goto } j$ or
- $instr_i: \text{ if } c_r = 0 \text{ then goto } j \text{ else } c_r = c_r - 1 \text{ goto } k$

for $j, k \in \{1, 2, \dots, n\}$ and $r \in \{1, 2\}$.

A *computation* of a Minsky Counter Machine is a sequence of configurations (i, c_1, c_2) where $i \in \{1, 2 \dots n\}$ is the label of instruction to be performed and c_1, c_2 are values of the counters. The starting configuration is $(1, 0, 0)$. The computation sequence is deterministic and determined by the instructions in the obvious way. The *halting problem* of a Minsky Counter Machine as to decide whether the computation of the machine ever halts (reaches the instruction with label n).

Given a two counter machine, we construct a Petri game $\mathcal{N} = \langle P, T_1, T_2, W, B, \varphi \rangle$ where $P = \{p_1, \dots, p_n, p_{c_1}, p_{c_2}\} \cup \{p_{i:c_r=0}, p_{i:c_r \neq 0} \mid i \in \{1 \dots n\}\}$ is the set of places containing place for each of n instructions, a place for each counter and some helper places. The set of transitions is $T_1 = \{t_{i:1}, t_{i:2}, t_{i:3}, t_{i:4} \mid i \text{ is a decrement instruction}\} \cup \{t_i \mid i \text{ is an increment instruction}\}$ and $T_2 = \{t_{i:cheat} \mid i \text{ is a decrement transition}\}$. Figures 4a, 4b define the gadgets that are added for each increment and decrement instruction, respectively. The function B binds every place to 1, except for p_{c_1}, p_{c_2} each of which have a bound ∞ . The formula to be satisfied by the goal marking is $\varphi = p_n \geq 1$, stating that a token is put to the place p_n . The initial marking contains just one token in the place p_1 and all other places are empty. We now

show that the counter machine halts iff the controller has a winning strategy in the constructed game.

“ \implies ”: Assume that the Minsky machine halts. The controller’s strategy is to faithfully simulate the counter machine, meaning that if the controller is in place p_i for a decrement instruction labelled with i , then it selects the transition $t_{i:1}$ in case the counter c_r is not empty and transition $t_{i:2}$ in case the counter is empty. For increment instructions there is only one choice for the controller. As the controller is playing faithfully, the transition $t_{i:cheat}$ is never enabled and hence the place p_n is eventually reached and the controller has a winning strategy.

“ \impliedby ”: Assume that the Minsky machine loops. We want to argue that there is no winning controller’s strategy. For the contradiction assume that the controller can win the game also in this case. Clearly, playing faithfully as described in the previous direction will never place a token into the place p_n . Hence at some point the controller must cheat when executing the decrement instruction (no cheating is possible in increment instruction). There are two cheating scenarios. Either the place p_{c_r} is empty and the controller fires $t_{i:1}$ which leads to a deadlock marking and clearly cannot reach the goal marking that marks the place p_n . The other cheating move is to play $t_{i:2}$ in case the place p_{c_r} is not empty. However, now in the next round the environment has the transition $t_{i:cheat}$ enabled and can deadlock the net by firing it. In any case, it is impossible to mark the place p_n , which contradicts the existence of a winning strategy for the controller. \square

This means that running Algorithm 1 specialized to the case of Petri nets does not necessarily terminates on general unbounded Petri nets (nevertheless, if it terminates even for unbounded nets then it still provides a correct answer). In case a Petri net game is bounded, e.g. for each $p \in P$ we have $B(p) \neq \infty$, we can guarantee also termination of the algorithm.

Theorem 5. *The reachability control problem for bounded Petri net games is decidable.*

Proof. Theorem 2 guarantees the correctness of Algorithm 1 (where we assume that the alternating simulation relation is the identity relation). As the presence of soft bounds for each place makes the state space finite, we can apply Theorem 3 to conclude that Algorithm 1 terminates. Hence the total correctness of the algorithm is established. \square

Clearly the empty alternating simulation that was used in the proof of the above theorem guarantees correctness of the algorithm, however, it is not necessarily the most efficient one. Hence we shall now define a syntax-based and linear time computable alternating simulation for Petri net games. For a transition $t \in T_1 \cup T_2$, let $\bullet t = \{p \in P \mid W(p, t) > 0\}$. We partition the places P in *equality* places P_e and *remaining* places P_r such that $P_e = \bigcup_{t \in T_2} \bullet t$ and $P_r = P \setminus P_e$. Recall that a Petri net game \mathcal{N} induces a game graph with configurations $C_i = \mathcal{M}(\mathcal{N}) \times \{i\}$ for $i \in \{1, 2\}$.

Definition 6 (Alternating Simulation for Petri Net Games). *We define a relation $\sqsubseteq \subseteq C_1 \times C_1 \cup C_2 \times C_2$ on configurations such that $(M, i) \sqsubseteq (M', i)$ iff for all $p \in P_e$ we have $M(p) = M'(p)$ and for all $p \in P_r$ we have $M(p) \leq M'(p)$.*

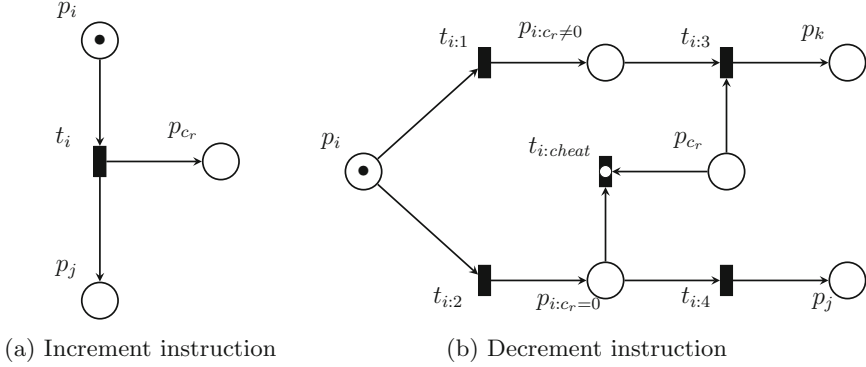


Fig. 4. Petri game gadgets for wwo counter machine instructions

Theorem 6. *The relation \sqsubseteq is an alternating simulation and it is computable in linear time.*

Proof. In order to prove that \sqsubseteq is alternating simulation, let us assume that $(M_1, i) \sqsubseteq (M'_1, i)$ and we need to prove the three conditions of the alternating simulation relation from Definition 3.

- If (M_1, i) is a goal configuration, then by definition since the set of goal markings are upward closed, (M'_1, i) is also a goal configuration.
- If $(M_1, 1) \xrightarrow{t} (M_2, 2)$ where $t \in T_1$ then by definition of transition firing for every $p \in P$ holds that $M_1(p) \geq W(p, t)$. Since $(M_1, 1) \sqsubseteq (M'_1, 1)$ then for every $p \in P$ holds $M'_1(p) \geq M_1(p)$ implying that $M'_1(p) \geq W(p, t)$. Hence t can be fired at M'_1 and let $M'_1 \xrightarrow{t} M'_2$. It is easy to verify that $(M_2, 2) \sqsubseteq (M'_2, 2)$.
- The third condition claiming that if $(M'_1, 2) \xrightarrow{t} (M'_2, 1)$ where $t \in T_2$ then $(M_1, 2) \xrightarrow{t} (M_2, 1)$ such that $(M_2, 1) \sqsubseteq (M'_2, 1)$ follows straightforwardly as for all input places to the environment transition t we assume an equal number of tokens in both M_1 and M'_1 .

We can conclude with the fact that \sqsubseteq is an alternating simulation. The relation can be clearly decided in linear time as it requires only comparison of number of tokens for each place in the markings. \square

Remark 2. In Fig. 3b, at the initial configuration $c_0 = ((1, 0, 0, 0), 2)$ *Player-2* has two transitions enabled. The two successor configurations to c_0 are $c_1 = ((0, 1, 1, 0), 1)$ and $c_2 = ((0, 0, 1, 0), 1)$ which are reached by firing the transitions t_0 and t_1 respectively. Since $c_1 \sqsupseteq c_2$, Algorithm 1 explores c_2 and ignores the successor configuration c_1 . Hence our algorithm terminates on this game and correctly declares it as winning for *Player-2*.

We now point out the following facts about Algorithm 1 and how it is different from the classical Liu and Smolka's algorithm [7]. First of all, Liu and Smolka

Table 1. Fire alarm system

Fire alarm		Winning	Time (sec.)		States		Reduction in %	
Sensors	Channels		LS	ALT	LS	ALT	Time	States
2	2	True	1.5	0.2	116	19	87.5	83.6
3	2	True	30	0.7	434	56	97.6	87.1
4	2	False	351.7	1.1	934	60	99.7	93.6
5	2	False	2249.8	1.1	1663	63	99.9	96.2
6	2	False	8427.1	1.3	3121	64	99.9	98.0
7	2	False	T.O.	1.4	–	66	–	–
2	3	True	20.2	0.5	385	25	97.5	93.5
2	4	True	622.7	1.0	1233	31	99.8	97.5
2	5	True	10706.7	2.3	3564	37	99.9	98.9
2	6	True	T.O.	3.4	–	43	–	–

do not consider any alternating simulation and they do not use the set *Lose* in order to ensure early termination of the algorithm. As a result, Liu and Smolka’s algorithm does not terminate (for any search strategy) on the Petri net game from Fig. 2, whereas our algorithm always terminates (irrelevant of the search strategy) and provides a conclusive answer. This fact demonstrates that not only our approach is more efficient but it also terminates on a larger class of Petri net games than the previously studied approaches.

5 Implementation and Experimental Evaluation

We implemented both Algorithm 1 (referred to as ALT and using the alternating simulation on Petri nets given in Definition 6) and the classical Liu and Smolka algorithm (referred to as LS) in the prototype tool sAsET [8]. We present three use cases where each of the experiments is run 20 times and the average of the time/space requirements is presented in the summary tables. The columns in the tables show the scaling, a Boolean value indicating whether the initial configuration has a winning strategy or not, the time requirements, number of explored states (the size of the set *Disc*) and a relative reduction of the running time and state space in percentage. The experiments are executed on AMD Opteron 6376 processor running at 2.3 GHz with 10 GB memory limit.

5.1 Fire Alarm Use Case

The German company SeCa GmbH produces among other products fire alarm systems. In [4] the formal verification of a wireless communication protocol for a fire alarm system is presented. The fire alarm system consists of a central unit and a number of sensors that communicate using a wireless Time Division Multiple

Table 2. Student teacher use case

Model	Winning	Time (sec.)		States		Reduction in %	
		LS	ALT	LS	ALT	Time	States
$s_1a_6d_2w_2$	False	14.4	12.0	278	213	16.7	23.4
$s_1a_6d_2w_3$	False	220.9	101.0	734	527	54.3	28.2
$s_1a_6d_2w_4$	False	2107.7	620.1	1546	1027	70.6	33.6
$s_1a_6d_2w_5$	False	9968.1	2322.3	2642	1672	76.7	36.7
$s_1a_6d_2w_{10}$	True	7159.3	2710.0	2214	1744	62.1	21.2
$s_1a_6d_2w_{11}$	True	8242.0	2662.7	2226	1756	67.7	21.1
$s_1a_6d_2w_{12}$	True	7687.7	2867.2	2238	1768	62.7	21.0
$s_1a_6d_2w_{13}$	True	7819.7	2831.4	2250	1780	63.8	20.9
$s_1a_6d_2w_{14}$	True	7674.1	2960.3	2262	1792	61.4	20.8
$s_1a_6d_2w_{15}$	True	8113.8	2903.7	2274	1804	64.2	20.7
$s_1a_7d_2w_2$	False	35.3	25.8	382	286	26.9	25.1
$s_1a_7d_2w_3$	False	744.7	284.0	1114	783	61.9	29.7
$s_1a_7d_2w_4$	False	9400.1	2637.2	2604	1674	71.9	35.7
$s_1a_7d_2w_5$	False	T.O.	10477.3	T.O.	2961	–	–
$s_4a_6d_1w_2$	True	40.6	28.9	290	250	28.8	13.8
$s_4a_6d_1w_3$	True	60.7	45.3	326	286	25.4	12.3
$s_4a_6d_1w_4$	True	90.4	62.0	362	322	31.4	11.0
$s_4a_6d_1w_5$	True	122.9	93.9	398	358	23.6	10.1
$s_4a_6d_1w_6$	True	172.1	122.6	434	394	28.8	9.2
$s_4a_6d_1w_7$	True	197.7	163.6	470	430	17.2	8.5
$s_4a_6d_1w_8$	True	263.3	190.6	506	466	27.6	7.9
$s_2a_4d_1w_2$	True	1.7	1.9	90	82	–11.8	8.9
$s_2a_4d_1w_3$	True	2.9	3.1	110	102	–6.9	7.3
$s_2a_4d_1w_4$	True	4.3	4.7	130	122	–9.3	6.2
$s_2a_4d_1w_5$	True	6.3	6.4	150	142	–1.6	5.3
$s_2a_4d_1w_6$	True	9.5	8.3	170	162	12.6	4.7
$s_2a_4d_1w_7$	True	11.9	11.7	190	182	1.7	4.2
$s_2a_4d_1w_8$	True	16.4	15.7	210	202	4.3	3.8

Access protocol. We model a simplified version of the fire alarm system from [4] as a Petri net game and use our tool to guarantee a reliable message passing even under interference. Table 1 shows that as the number of sensors and channels increases, we achieve an exponential speed up using ALT algorithm compared to LS. The reason is that the alternating simulation significantly prunes out the state space that is necessary to explore.

Table 3. Cat and mice use case

Model	Winning	Time (sec.)		States		Reduction in %	
		LS	ALT	LS	ALT	Time	States
m211	False	159.2	109.7	525	405	31.1	22.9
m222	False	562.1	298.7	798	579	46.9	27.4
m322	False	5354.2	1914.4	1674	1096	64.2	34.5
m332	False	9491.7	4057.6	2228	1419	57.3	36.3
m333	False	T.O.	8174.5	T.O.	1869	–	–
m2000	False	127.7	50.2	474	341	60.7	28.1
m2100	False	404.6	157.5	762	490	61.1	35.7
m2200	False	776.4	148.3	921	490	80.9	46.8

5.2 Student-Teacher Scheduling Use Case

In [5] the student semester scheduling problem is modelled as a workflow net. We extend this Petri net into a game by introducing a teacher that can observe the student behaviour (work on assignments) and there is a race between the two players. In Table 2 we see the experimental results. Model instances are of the form $s_i a_j d_k w_l$ where i is the number of students, j is the number of assignments, k is the number of deliverables in each assignment and l is the number of weeks in a semester. We can observe a general trend that with a higher number of assignments ($j = 6$ and $j = 7$) and two deliverables per assignment ($d = 2$), the alternating simulation reduces a significant amount of the state space leading to considerable speed up. As the number of assignment and deliverables gets lower, there is less and less to save when using the ALT algorithm, ending in a situation where only a few percents of the state space can be reduced at the bottom of the table. This leaves us only with computational overhead for the additional checks related to alternating simulation in the ATL algorithm, however, the overhead is quite acceptable (typically less than 20% of the overall running time).

5.3 Cat and Mouse Use Case

This Petri net game consists of an arena of grid shape with number of mice starting at the first row. The mice try to reach the goal place at the other end of the arena without being caught while doing so by the cats (that can move only on one row of the arena). In Table 3 we consider arena of size 2×3 and 2×4 and the model instances show the initial mice distribution in the top row of the grid. The table contains only instances where the mice do not have a winning strategy and again shows an improvement both in the running time as well as the number of explored states. On the positive instances there was no significant difference between ALT and LS algorithms as the winning strategy for mice were relatively fast discovered by both approaches.

6 Conclusion

We presented an on-the-fly algorithm for solving strictly alternating games and introduced the notion of alternating simulation that allows us to significantly speed up the strategy synthesis process. We formally proved the soundness of the method and instantiated it to the case of Petri net games where the problems are in general undecidable and hence the algorithm is not guaranteed to terminate. However, when using alternating simulation there are examples where it terminates for any search strategy, even though classical approaches like Liu and Smolka dependency graphs do not terminate at all. Finally, we demonstrated on examples of Petri net games with soft bounds on places (where the synthesis problem is decidable) the practical applicability of our approach. In future work, we can consider an extension of our framework to other types of formalism like concurrent automata or time Petri nets.

Acknowledgments. The work was supported by the ERC Advanced Grant LASSO and DFF project QASNET.

References

1. Abdulla, P.A., Bouajjani, A., d’Orso, J.: Deciding monotonic games. In: Baaz, M., Makowsky, J.A. (eds.) CSL 2003. LNCS, vol. 2803, pp. 1–14. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45220-1_1
2. Cassez, F., David, A., Fleury, E., Larsen, K.G., Lime, D.: Efficient on-the-fly algorithms for the analysis of timed games. In: Abadi, M., de Alfaro, L. (eds.) CONCUR 2005. LNCS, vol. 3653, pp. 66–80. Springer, Heidelberg (2005). https://doi.org/10.1007/11539452_9
3. Dalsgaard, A.E., et al.: Extended dependency graphs and efficient distributed fixed-point computation. In: van der Aalst, W., Best, E. (eds.) PETRI NETS 2017. LNCS, vol. 10258, pp. 139–158. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-57861-3_10
4. Feo-Arenis, S., Westphal, B., Dietsch, D., Muñoz, M., Andisha, S., Podelski, A.: Ready for testing: ensuring conformance to industrial standards through formal verification. *Formal Aspects Comput.* **28**(3), 499–527 (2016). <https://doi.org/10.1007/s00165-016-0365-3>
5. Juhásova, A., Kazlov, I., Juhás, G., Molnár, L.: How to model curricula and learn-flows by petri nets - a survey. In: 2016 International Conference on Emerging eLearning Technologies and Applications (ICETA), pp. 147–152, November 2016
6. Minsky, M.L.: Computation: finite and infinite machines. *Am. Math. Mon.* **75** (1968)
7. Liu, X., Smolka, S.A.: Simple linear-time algorithms for minimal fixed points. In: Larsen, K.G., Skyum, S., Winskel, G. (eds.) ICALP 1998. LNCS, vol. 1443, pp. 53–66. Springer, Heidelberg (1998). <https://doi.org/10.1007/BFb0055040>
8. Muñoz, M.: Model checking for time division multiple access systems. Ph.D. thesis, Freiburg University (2015). <https://doi.org/10.6094/UNIFR/10161>
9. Petri, C.A.: Kommunikation mit Automaten. Ph.D. thesis, Universität Hamburg (1962)
10. Raskin, J., Samuelides, M., Begin, L.V.: Games for counting abstractions. *Electr. Notes Theor. Comput. Sci.* **128**(6), 69–85 (2005). <https://doi.org/10.1016/j.entcs.2005.04.005>