

Visibly Pushdown Modular Games *

Ilaria De Crescenzo
Università degli Studi di Salerno
Dipartimento di Informatica

Salvatore La Torre
Università degli Studi di Salerno
Dipartimento di Informatica

Yaron Velner
Tel Aviv University
The Blavatnik School of Computer Science

Games on recursive game graphs can be used to reason about the control flow of sequential programs with recursion. In games over recursive game graphs, the most natural notion of strategy is the modular strategy, i.e., a strategy that is local to a module and is oblivious to previous module invocations, and thus does not depend on the context of invocation. In this work, we study for the first time modular strategies with respect to winning conditions that can be expressed by a pushdown automaton. We show that such games are undecidable in general, and become decidable for visibly pushdown automata specifications. Our solution relies on a reduction to modular games with finite-state automata winning conditions, which are known in the literature. We carefully characterize the computational complexity of the considered decision problem. In particular, we show that modular games with a universal Büchi or co-Büchi visibly pushdown winning condition are EXPTIME-complete, and when the winning condition is given by a CARET or NWTL temporal logic formula the problem is 2EXPTIME-complete, and it remains 2EXPTIME-hard even for simple fragments of these logics. As a further contribution, we present a different solution for modular games with finite-state automata winning condition that runs faster than known solutions for large specifications and many exits.

1 Introduction

Recursive state machines (RSMs) carefully model the control flow of systems with potentially recursive procedure calls [2]. A recursive state machine is composed of a set of modules, whose vertices can be standard vertices or can correspond to invocations of other modules. A large number of hardware and software systems fits into this class, such as procedural and object-oriented programs, distributed systems, communication protocols and web services.

In the open systems setting, i.e., systems where an execution depends on the interaction of the system with the environment, the natural counterpart of recursive state machines is two-player recursive game graphs. A recursive game graph (RGG) is essentially a recursive state machine where vertices are split into two sets each controlled by one of the players, and thus corresponds to pushdown games with an emphasis on the modules composing the system.

In this paper we focus on solving pushdown games on RGG in which the first player is restricted to *modular strategies* [7]. A strategy is a mapping that specifies, for each play ending into a controlled state, the next move. Modular strategies are formed of a set of strategies, one for each RGG module, that are *local* to a module and *oblivious* of the history of previous module activations, i.e., the next move in such strategies is determined by looking only at the local memory of the current module activation (by making the local memory *persistent* across module activations, deciding these games becomes undecidable already with reachability specifications [7]).

*This work was partially funded by the MIUR grants FARB 2011-2012-2013, Università degli Studi di Salerno (Italy).

The main motivation for considering modular strategies is related to the synthesis of controllers [18, 20]: given a description of the system where some of the choices depend upon the input and some represent uncontrollable internal non-determinism, the goal is to design a controller that supplies inputs to the system such that it satisfies the correctness specification. Synthesizing a controller thus corresponds to computing winning strategies in two-player games, and a modular strategy would correspond to a modular controller.

The notion of modular strategy is also of independent interest and has recently found application in other contexts, such as, the automatic transformation of programs for ensuring security policies in privilege-aware operating systems [11], and a general formulation of the synthesis problem from libraries of open components [10].

The problem of deciding the existence of a modular strategy in a recursive game graph has been already studied with respect to ω -regular specifications. The problem is known to be NP-complete for reachability specifications [7], EXPTIME-complete for specifications given as deterministic and universal Büchi or Co-Büchi automata, and 2EXPTIME-complete for LTL specifications [5].

In this paper, we study this problem with respect to several classes of specifications that can be expressed as a pushdown automaton. We show that in the general case the problem is undecidable. We thus focus on visibly pushdown automata (VPA) [8] specifications with Büchi or co-Büchi acceptance. In the following, we refer to this problem as the MVPG problem and omit the acceptance condition of the VPA by meaning either one of them.

Our main contributions are:

- We show a polynomial time reduction from the MVPG problem with deterministic or universal VPA specifications to recursive modular games over ω -regular specifications. By [5], we get that this problem is EXPTIME-complete. We then use this result to show the membership to 2EXPTIME for the MVPG problem with nondeterministic VPA specifications.
- We show that when the winning condition is expressed as a formula of the temporal logics CARET [3] and NWTTL [1] the MVPG problem is 2EXPTIME-complete, and hardness can be shown also for very simple fragments of the logics. In particular, we show a 2EXPTIME lower bound for the fragment containing only conjunctions of disjunctions of bounded-size path formulas (i.e., formulas expressing either the requirement that a given finite sequence is a subsequence of a word or its negation), that is in contrast with the situation in finite game graphs where PSPACE-completeness holds for larger significant fragments (see [4, 6]). On the positive side, we are able to show an exponential-time algorithm to decide the MVPG problem for specifications given as conjunctions of temporal logic formulas that can be translated into a polynomial-size VPA (such formulas include the path formulas).
- We also give a different solution for recursive games with finite-state automata specifications. Our approach yields an upper bound of $|G|2^{O(d^2(k+\log d)+\beta)}$ for the MVPG problem, where d is the number of P (the VPA) states, k is the number of G (the RGG) exits, and β is the number of *call edges* of G , i.e., the number of module pairs (m, m') such that there is a call from m to m' . The known solution [5] yields an $|G|2^{O(kd^2\log(kd))}$ upper bound. Thus, our solution is faster when k and d are large, and matches the known EXPTIME lower bound [5]. In addition we use one-way nondeterministic/universal tree automata instead of two-way alternating tree automata, thus we explicitly handle aspects that are hidden in the construction from [5].

Related work. Besides the already mentioned work that has dealt with modular games, but only for ω -regular specifications [5] or reachability [7], other research on pushdown games have focused on the

standard notion of winning strategy. We recall that determining the existence of a standard winning strategy (i.e., non-modular) in pushdown games with reachability specifications is known to be EXPTIME-complete [21]. Deciding such games is 2EXPTIME-complete for nondeterministic visibly pushdown specifications and 3EXPTIME-complete for LTL and CARET specifications [13].

The synthesis from recursive-component libraries defines a pushdown game which is *orthogonal* to the MVPG problem: there the modules are already synthesized and the game is on the function calls. Deciding such games for NWTL is 2EXPTIME-complete [15]. The synthesis from open recursive-component libraries combines both this synthesis problem and the MVPG synthesis. Deciding the related game problem with reachability specifications is EXPTIME-complete [10]. Other synthesis problems dealing with compositions of component libraries are [14, 9], and for modules expressed as terms of the λY -calculus, [19].

2 Preliminaries

Given two positive integers i and j , $i \leq j$, we denote with $[i, j]$ the set of integers k with $i \leq k \leq j$, and with $[j]$ the set $[1, j]$.

We fix a set of atomic propositions AP and a finite alphabet Σ . A ω -word over Σ is a mapping that assigns to each *position* $i \in \mathbb{N}$ a symbol $\sigma_i \in \Sigma$, and is denoted as $\{\sigma_i\}_{i \in \mathbb{N}}$ or equivalently $\sigma_1 \sigma_2 \dots$.

Recursive game graph. A recursive game graph (RGG) is composed of *game modules* that are essentially two-player graphs (i.e., graphs whose vertices are partitioned into two sets depending on the player who controls the outgoing moves) with *entry* and *exit* nodes and two different kind of vertices: the nodes and the boxes. A *node* is a standard graph vertex and a *box* corresponds to invocations of other game modules in a potentially recursive manner (in particular, entering into a box corresponds to a module *call* and exiting from a box corresponds to a *return* from a module).

As an example consider the RGG in Fig. 1, where the vertices of player 0 (pl_0) are denoted with rounds, those of player 1 (pl_1) with squares and the rectangles denote the vertices where there are no moves that can be taken by any of the players and correspond to calls and exits. Atomic propositions p_a, p_b, p_c and p_d label the vertices. Each RGG has a distinct game module which is called the *main* module (module M_{in} in the figure). In analogy to many programming languages, we require that the main module cannot be invoked by any other module. A typical play starts in vertex e_{in} . From this node, there is only one possible move to take and thus the play continues at the call to M_1 on box b , which then takes the play to the entry e_1 in M_1 .

This is a vertex of the adversary, who gets to pick the transition and thus can decide to visit either u_3 (generating p_a) or u_4 (generating p_b). In any of the cases, the play will evolve reaching the exit and then the control will return to module M_1 at the return vertex on box b . Here pl_0 gets to choose if generating p_c or p_d and so on back to the call to M_1 . Essentially, along any play alternatively pl_1 chooses one between p_a and p_b , and pl_0 chooses one between p_c and p_d . Formally, we have the following definitions.

Definition 1. (RECURSIVE GAME GRAPH) A recursive game graph G over AP is a triple $(M, m_{in}, \{S_m\}_{m \in M})$ where M is a finite set of module names, $m_{in} \in M$ denotes the main module and for each $m \in M$, S_m is a game module. A game module S_m is $(N_m, B_m, Y_m, En_m, Ex_m, \delta_m, \eta_m, P_m^0, P_m^1)$ where:

- N_m is a finite set of nodes and B_m is a finite set of boxes;
- $Y_m : B_m \rightarrow (M \setminus \{m_{in}\})$ maps every box to a module;




Figure 1: A sample RGG.

- $En_m \subseteq N_m$ is a non-empty set of entry nodes;
- $Ex_m \subseteq N_m$ is a (possibly empty) set of exit nodes;
- $\delta_m : N_m \cup Retns_m \rightarrow 2^{N_m \cup Calls_m}$ is a transition function where $Calls_m = \{(b, e) | b \in B_m, e \in En_{Y_m(b)}\}$ is the set of calls and $Retns_m = \{(b, e) | b \in B_m, e \in Ex_{Y_m(b)}\}$ is the set of returns;
- $\eta_m : V_m \rightarrow 2^{AP}$ labels in 2^{AP} each vertex from $V_m = N_m \cup Calls_m \cup Retns_m$;
- P_m^0 and P_m^1 form a partition of $(N_m \cup Retns_m) \setminus Ex_m$; P_m^0 is the set of the positions of pl_0 and P_m^1 is the set of the positions of pl_1 .

In the rest of the paper, we denote with: G an RGG as in the above definition; $V = \bigcup_m V_m$ (set of vertices); $B = \bigcup_m B_m$ (set of boxes); $Calls = \bigcup_m Calls_m$ (set of calls); $Retns = \bigcup_m Retns_m$ (set of returns); $Ex = \bigcup_m Ex_m$ (set of exits); $P^\ell = \bigcup_m P_m^\ell$ for $\ell \in [0, 1]$ (set of all positions of pl_ℓ); and $\eta : V \rightarrow 2^{AP}$ such that $\eta(v) = \eta_m(v)$ where $v \in V_m$.

To ease the presentation we make the following assumptions (with $m \in M$):

- there is only one entry point to every module S_m and we refer to it as e_m ;
- there are no transitions to an entry, i.e., $e_m \notin \delta_m(u)$ for every u ;
- there are no transitions from an exit, i.e., $\delta_m(x)$ is empty for every $x \in Ex_m$;
- a module is not called immediately after a return from another module, i.e., $\delta_m(v) \subseteq N_m$ for every $v \in Retns_m$.

A (global) *state* of an RGG is composed of a call stack and a vertex, that is, each state of G is of the form $(\alpha, u) \in B^* \times V$ where $\alpha = b_1 \dots b_h$, $b_1 \in B_m$, $b_{i+1} \in B_{Y(b_i)}$ for $i \in [h-1]$ and $u \in V_{Y(b_h)}$.

A *play* of G is a (possibly finite) sequence of states $s_0 s_1 s_2 \dots$ such that $s_0 = (\varepsilon, e_{in})$ and for $i \in \mathbb{N}$, denoting $s_i = (\alpha_i, u_i)$, one of the following holds:

- **Internal move:** $u_i \in (N_m \cup Retns_m) \setminus Ex_m$, $u_{i+1} \in \delta_m(u_i)$ and $\alpha_i = \alpha_{i+1}$;
- **Call to a module:** $u_i \in Calls_m$, $u_i = (b, e_{m'})$, $u_{i+1} = e_{m'}$ and $\alpha_{i+1} = \alpha_i.b$;
- **Return from a call:** $u_i \in Ex_m$, $\alpha_i = \alpha_{i+1}.b$, and $u_{i+1} = (b, u_i)$.

Fix an infinite play $\pi = s_0 s_1 \dots$ of G where $s_i = (\alpha_i, u_i)$ for each $i \in \mathbb{N}$.

With π_k we denote $s_0 \dots s_k$, i.e., the prefix of π up to s_k . For a finite play $\pi'.s$, with $ctr(\pi'.s)$ we denote the module m where the control is at s , i.e., such that $u \in V_m$ where $s = (\alpha, u)$. We define μ_π such that $\mu_\pi(i, j)$ holds iff for some $m \in M$, $u_i \in Calls_m$ and j is the smallest index s.t. $i < j$, $u_j \in Retns_m$ and $\alpha_i = \alpha_j$ (μ_π captures the matching pairs of calls and returns in π).

Modular strategies. Fix $\ell \in [0, 1]$. A *strategy* of pl_ℓ is a function f that associates a legal move to every play ending in a node controlled by pl_ℓ .

A modular strategy constrains the notion of strategy by allowing only to define the legal moves depending on the “local memory” of a module activation (every time a module is re-entered the local memory is reset).

Formally, a *modular strategy* f of pl_ℓ is a set of functions $\{f_m\}_{m \in M}$, one for each module $m \in M$, where $f_m : V_m^*.P_m^\ell \rightarrow V_m$ is such that $f_m(\pi.u) \in \delta_m(u)$ for every $\pi \in V_m^*$, $u \in P_m^\ell$.

The local successor of a position in π is: the successor according to the matching relation μ_π at matched calls, undefined at an exit or an unmatched call, and the next position otherwise. Formally, the *local successor* of j , denoted $succ_\pi(j)$, is: h if $\mu_\pi(j, h)$ holds; otherwise, is undefined if either $u_j \in Ex$ or $u_j \in Calls$ and $\mu_\pi(j, h)$ does not hold for every $h > j$; and $j+1$ in all the remaining cases.

For each $i \leq |\pi|$, the *local memory* of π_i , denoted $\lambda(\pi_i)$, is the maximal sequence $u_{j_1} \dots u_{j_k}$ such that $u_{j_k} = u_i$ and $j_{h+1} = \text{succ}_\pi(j_h)$ for each $h \in [k-1]$. (Note that since the sequence is maximal, $u_{j_1} = e_m$ where $m = \text{ctr}(\pi_i)$.)

A play π conforms to a modular strategy $f = \{f_m\}_{m \in M}$ of pl_ℓ if for every $i < |\pi|$, denoting $\text{ctr}(\pi_i) = m$, $u_i \in P_m^\ell$ implies that $u_{i+1} = f_m(\lambda(\pi_i))$.

Consider again the example from Fig. 1. A strategy of pl_0 that chooses alternatively to generate p_c and p_d is modular, in fact it requires as memory just to store the last move from the return of b , and thus is local to the current (sole) activation of module M_{in} . Instead, a strategy that attempts to match each p_a with p_c and each p_b with p_d is clearly non modular.

We remark that modular strategies are oblivious to the previous activations of a module. In the RGG of Fig. 1, a modular strategy for pl_1 would only allow either one of the behaviors: “ pl_1 always picks p_a ” or “ pl_1 always picks p_b ”.

Winning conditions and modular games. A modular game on RGG is a pair $\langle G, L \rangle$ where G is an RGG and L is a winning condition. A winning condition is a set L of ω words over a finite alphabet $\Sigma = 2^{AP}$, where AP is a set of propositions. Given an RGG G , for a play $\pi = s_0 s_1 \dots$ of G , with $s_i = (\alpha_i, u_i)$, we define the word $w_\pi = \eta(u_0)\eta(u_1)\dots$, which is the mapping that assigns to each position the correspondent symbol from Σ . A (modular) strategy f is winning if $w_\pi \in L$ for every play π of G that conforms to f . The *modular game problem* asks to determine the existence of a winning (modular) strategy of pl_0 in a given modular game. In the following sections, we consider L given by pushdown, visibly pushdown automata and by LTL, CARET or NCTL formulas.

3 Pushdown specification

Pushdown modular games. A *pushdown modular game* is a pair $\langle G, \mathcal{P} \rangle$ where G is an RGG and \mathcal{P} is a pushdown automaton, whose accepted language defines the winning condition in G . A *pushdown automaton* \mathcal{P} is a tuple $(Q, q_0, \Sigma, \Gamma, \delta, \gamma^\perp, F)$ where Q is a finite set of states, $q_0 \in Q$ is the initial state, Σ is a finite alphabet, Γ is a finite stack alphabet, γ^\perp is the bottom-of-stack symbol, $F \subseteq Q$ defines an acceptance condition, and $\delta : Q \times \{\Sigma \cup \varepsilon\} \times \Gamma \rightarrow 2^{Q \times \Gamma^*}$ is the transition function. A pushdown automaton is *deterministic* if it satisfies the following two conditions: - $\delta(q, \alpha, \gamma)$ has at most one element for any $q \in Q$, $\gamma \in \Gamma$ and $\alpha \in \Sigma$ (or $\alpha = \varepsilon$); - if $\delta(q, \varepsilon, \gamma) \neq \emptyset$ for any $q \in Q$ and $\gamma \in \Gamma$ then $\delta(q, \alpha, \gamma) = \emptyset$ for any $\alpha \in \Sigma$.

Undecidability of pushdown specification. The modular game problem becomes undecidable if we consider winning conditions expressed as standard (deterministic) pushdown automata. This is mainly due to the fact that the stack in the specification pushdown automaton is not synchronized with the call-return structure of the recursive game graph.

We prove the undecidability of our problem with pushdown specification by presenting a reduction from the problem of checking the emptiness of the intersection of two deterministic context-free languages.

Consider two context-free languages L_1 and L_2 on an alphabet $\Sigma = \{\sigma_1, \sigma_2, \dots, \sigma_n\}$, which are accepted by two pushdown automata, \mathcal{P}_1 and \mathcal{P}_2 , respectively. We want to construct an instance $\langle G, \mathcal{P} \rangle$ of a deterministic pushdown modular game problem such that exists a winning modular strategy for pl_0 in $\langle G, \mathcal{P} \rangle$ if and only if the intersection of L_1 and L_2 is not empty.

The basic idea of the reduction is to construct a game where pl_1 challenges pl_0 to generate a word from either L_1 or L_2 , and pl_0 must match the choice of pl_1 without knowing it in order to win. We construct an RGG G with two modules, m_{in} and m (see the Fig. 2).

The module m_{in} is the main module and is composed of an entry e_{in} , two internal nodes u_1 and u_2 , and one box b labeled with m . The entry e_{in} belongs to pl_1 and has two transitions, one to each internal node. From u_1 and u_2 there is only one possible move, which leads to b . The labeling function associates the symbol a_1 to u_1 and the symbol a_2 to u_2 with $a_1, a_2 \notin \Sigma$. The node e_{in} and the call (b, e_m) are both labeled with $\sharp \notin \Sigma \cup \{a_1, a_2\}$. Observe that since the only choice of pl_1 is at e_{in} , for any strategy f of pl_0 there are only two plays conforming to it: one going through u_1 and the other through u_2 .

The module m is essentially a deterministic generator of any word in $\{\sharp\}.\Sigma^*.\{\sharp\}$. The module m has one entry e_m , $|\Sigma|$ internal nodes v_1, v_2, \dots, v_n and a sink node s (i.e., a node with only ingoing edges). All the vertices of m belong to pl_0 . There are only outgoing edges from e_m , which take to each of the other vertices of m . Moreover, there is a transition from v_i to v_j for any $i, j \in [n]$, and from any node there is a move to s . Each node v_i is labeled with σ_i . for $i \in [n]$. The symbol \sharp labels e_m and s .

As winning condition, we construct a deterministic pushdown automaton \mathcal{P} . In the initial state \mathcal{P} reads \sharp and moves into state q_0 . Fix $i \in [2]$. From q_0 and on input a_i , \mathcal{P} enters state q_i . From q_i , \mathcal{P} reads two occurrences of \sharp and enters the initial state of \mathcal{P}_i . From any \mathcal{P}_i state, \mathcal{P} behaves as \mathcal{P}_i and in addition, from each final state of \mathcal{P}_i , it has a move on input \sharp that takes to the only final state q_f .

Since the strategy must be modular, in the module m the player pl_0 has no information about the choice of pl_1 in m_{in} . Also, the local strategy in module m generates one specific word (there are no moves of pl_1 allowed in m) and thus this is the same independently of the moves of pl_1 in module m_{in} . Thus, the local strategy in m is winning if and only if it generates a word in the intersection of L_1 and L_2 , and therefore, the following theorem holds.

Theorem 1. *The (deterministic) pushdown modular game problem is undecidable.*




Figure 2: The module m_{in} and m

4 Solving modular games with VPA specifications

Visibly pushdown automata. Consider the finite alphabet Σ , and let $call$, ret , and int be new symbols. We denote with $\Sigma_{call} = \Sigma \times \{call\}$, $\Sigma_{ret} = \Sigma \times \{ret\}$ and $\Sigma_{int} = \Sigma \times \{int\}$ and with $\widehat{\Sigma} = \Sigma_{call} \cup \Sigma_{ret} \cup \Sigma_{int}$.

A *visibly pushdown automaton* (VPA) P is a tuple $(Q, Q_0, \widehat{\Sigma}, \Gamma \cup \{\gamma^\perp\}, \delta, F)$ where Q is a finite set of states, $Q_0 \subseteq Q$ is a set of initial states, $\widehat{\Sigma}$ is a finite alphabet, Γ is a finite stack alphabet, γ^\perp is the bottom-of-stack symbol, $F \subseteq Q$ defines an acceptance condition, and $\delta = \delta^{int} \cup \delta^{push} \cup \delta^{pop}$ where $\delta^{int} \subseteq Q \times \Sigma_{int} \times Q$, $\delta^{push} \subseteq Q \times \Sigma_{call} \times \Gamma \times Q$, and $\delta^{pop} \subseteq Q \times \Sigma_{ret} \times (\Gamma \cup \{\gamma^\perp\}) \times Q$.

A *configuration* (or global state) of P is a pair (α, q) where $\alpha \in \Gamma^*.\{\gamma^\perp\}$ and $q \in Q$. Moreover, (α, q) is *initial* if $q \in Q_0$ and $\alpha = \gamma^\perp$. We omit the semantics of the transitions of P being quite standard. It can be obtained similarly to that of RGG with the addition of the inputs. Here we just observe that we allow pop transitions on empty stack (a stack containing only the symbol γ^\perp). In particular, a pop transition do not change the stack when γ^\perp is at the top, and by the definition of δ^{push} , γ^\perp cannot be pushed onto the stack. A *run* ρ of P over the input $\sigma_0\sigma_1 \dots$ is an infinite sequence $C_0 \xrightarrow{\sigma_0} C_1 \xrightarrow{\sigma_1} \dots$ where C_0 is the initial configuration and such that, for each $i \in \mathbb{N}$, C_{i+1} is obtained from C_i by applying a transition on input σ_i .

Acceptance of an infinite run depends on the control states that are visited infinitely often. Fix a run $\rho = (\gamma^\perp, q_0) \xrightarrow{\sigma_0} (\alpha_1, q_1) \xrightarrow{\sigma_1} (\alpha_2, q_2) \dots$. With a *Büchi acceptance condition*, ρ is accepting if $q_i \in F$ for infinitely many $i \in \mathbb{N}$ (*Büchi VPA*). With a *co-Büchi acceptance condition*, ρ is accepting if there is a $j \in \mathbb{N}$ such that $q_i \notin F$ for all $i > j$ (*co-Büchi VPA*).

A VPA P is *deterministic* if: (1) $|Q_0| = 1$, (2) for each $q \in Q$ and $\sigma \in \Sigma_{call} \cup \Sigma_{int}$ there is at most one transition of δ from q on input σ , and (3) for each $q_1 \in Q$, $\sigma \in \Sigma_{ret}$, $\gamma \in \Gamma \cup \{\gamma^\perp\}$ there is at most one transition from q on input σ and stack symbol γ . Note that a deterministic VPA is such that for each word w there is at most a run over it.

For a word w , a *deterministic/nondeterministic* (resp. *universal*) VPA accepts w if there exists an accepting run over w (resp. all runs over w are accepting).

Visibly pushdown games. A *visibly pushdown game on RGG* (VPRG) is a pair $\langle G, P \rangle$ where G is an RGG and P is a visibly pushdown automaton (see [8]). Consider a VPRG $\langle G, P \rangle$ where G is an RGG and P is a VPA. For a play $\pi = s_0 s_1 \dots$ of G , with $s_i = (\alpha_i, u_i)$, we define the word w_π as $\sigma_0 \sigma_1 \dots$ such that for $i \in \mathbb{N}$, $\sigma_i = (\eta_m(s_i), t_i)$ where $ctr(\pi_i) = m$ and t_i is *call* if $u_i \in Calls$, *ret* if $u_i \in Retns$, and *int* otherwise. The *visibly pushdown (modular) game problem* asks to determine the existence of a (winning) modular strategy of pl_0 in a given VPRG such that w_π is accepted by P for every play π that conforms to f . We denote the visibly pushdown modular game problem as MVPG problem.

When the VPA is a finite state automaton $\mathcal{B} = (Q, q_0, \Sigma, \delta, F)$ we denote with ω -MGP the ω -*modular game problem* that asks to determine the existence of a winning modular strategy of pl_0 in a given $\langle G, \mathcal{B} \rangle$ game (ω -MG).

Solving games with VPA specifications We consider games with winning conditions that are given by a VPA with different acceptance conditions. We present a reduction from recursive games with VPA specifications to recursive games with specifications that are given as finite state automata. The reduction is almost independent of the acceptance condition, and it works for reachability and safety conditions as well as for Büchi and co-Büchi acceptance conditions.

The reduction transforms a recursive game graph with a visibly pushdown automaton specification (with some acceptance condition) to a slightly different recursive game graph with a finite-state automaton specification (with the same type of acceptance condition). The key idea is to embed the top stack symbol of a VPA P in the states of a finite-state automaton A . In addition, the states of A will simulate the corresponding states of P and thus we will get that the winning conditions are equivalent. Clearly, a finite-state automaton cannot simulate an unbounded stack. While it is easy to keep track of the top symbol after a *push* operation, extracting the top symbol after a *pop* operation requires infinite memory. For this purpose we exploit the fact that the stacks of the VPA P and the game graph G are synchronized and we introduce a new *dummy* module d_m for every module in G . Recall that the invocation of a module m in G correspond to a *push* operation in P . We replace every invocation of m by a call to d_m . In d_m (see Figure 3) pl_1 first has to *declare* the value of the top symbol in P before the *push* operation by going to the corresponding $v_{\gamma_1}, \dots, v_{\gamma_g}$ state in d_m , and A can verify that pl_1 is honest since it keeps track of the current top-symbol (if the player is not honest then A goes to a sink accepting state and pl_1 loses). After the declaration, the

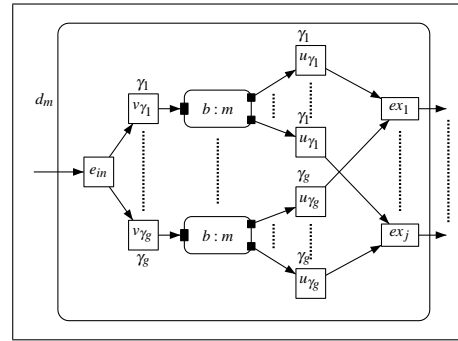


Figure 3: The module d_m . All the vertices of d_m are controlled by pl_1

module invokes the actual module m and when m terminates, then pl_1 must declare again the top-symbol γ_i of P , visiting the vertex u_{γ_i} , and A changes his simulated top symbol accordingly.

Denoting with k the number of exits of the starting RGG, with g the number of stack symbols, and d the number of states of the specification, we get that the resulting RGG G has $2k$ exits and the resulting automaton A has $O(dg)$ states. Thus combining this with the solution from [5], we get an upper bound linear in $|G|$ and exponential in $2k(dg)^2 \log(2kdg)$. We have:

Theorem 2. *The MVPG problem with winning conditions expressed as either a deterministic Büchi VPA or a deterministic co-Büchi VPA is EXPTIME-complete.*

The proposed reduction can be extended for universal VPA specification. W.l.o.g we assume that in the non-deterministic VPA for every state, stack letter and labeling there are exactly two possible transitions. In this case we add a dummy module e , that is composed only by pl_1 nodes and has one exit. Each transition from a node v to a node u is splitted in two transitions, $v \rightarrow e$ and $e \rightarrow u$. In the module e , pl_1 resolves the nondeterminism, selecting one of the two possible transitions for the VPA specification. The choices of pl_1 in e are oblivious to pl_0 . Hence, the universal VPA accepts if and only if pl_0 has a strategy that wins against all pl_1 choices in e , and we get the next theorem.

Theorem 3. *The MVPG problem with winning conditions expressed as a universal Büchi (resp. co-Büchi) VPA is EXPTIME-complete.*

We can handle nondeterministic VPAs in the following way: Let P be a nondeterministic Büchi VPA P . By [8], we can construct a nondeterministic Büchi VPA P' that accepts a word w iff P does not accept it, and such that the size of P' is exponential in the size of P . Complete P' with transitions that take to a rejecting state such that for each word there is at least a run of P' over it. Let P'' be the dual automaton of P' , i.e., P has the same components of P' except that acceptance is now universal and the set of accepting states is now interpreted as a co-Büchi condition. Clearly, P'' accepts exactly the same words as P and has size exponential in $|P|$. Similarly, we can repeat the above reasoning starting from a co-Büchi VPA P . Therefore, we have:

Theorem 4. *The MVPG problem with winning conditions expressed as a nondeterministic Büchi (resp. co-Büchi) VPA is in 2EXPTIME.*

5 Temporal logic winning conditions

By [3], we know that given a CARET formula φ it is possible to construct a nondeterministic Büchi VPA of size exponential in $|\varphi|$ that accepts exactly all the words that satisfy φ . From [1], we know that the same holds for the temporal logic NWTL. Thus, given a formula φ in any of the two logics, we construct a Büchi VPA P for its negation $\neg\varphi$. By dualizing as in the case of nondeterministic VPA specifications, we get a co-Büchi VPA that accepts all the models of φ and which size is exponential in $|\varphi|$. Since both CARET and NWTL subsume LTL [17], and LTL games are known to be 2EXPTIME-hard [18] already on standard finite game graphs, we get:

Theorem 5. *The MVPG problem with winning conditions expressed as CARET and NWTL formulas is 2EXPTIME-complete.*

The complexity of the temporal logic MVPG problem remains 2EXPTIME-hard even if we consider simple fragments.

A *path formula* is a formula expressing either the requirement that a given sequence appears as a subsequence in an ω -word or its logical negation. Path formulas are captured by LTL formulas of the

form $\diamond(p_1 \wedge \diamond(p_2 \wedge \dots \diamond(p_{n-1} \wedge \diamond p_n) \dots))$ and by their logical negation, where each p_i is state predicate, $\diamond\psi$ (*eventually ψ*) denotes that ψ holds at some future position, and \wedge is the Boolean conjunction. We denote such a fragment of LTL as PATH-LTL.

We present a reduction from exponential-space alternating Turing machines. We only give here the general idea. We use a standard encoding of computations, where cell contents are preceded by the cell number written in binary ($2N$ atomic propositions suffice to encode 2^N cell numbers) and configurations are sequences of cells encodings ended with a marker (the tape head and the current control state are encoded as cell contents).

Denote with Q and Σ respectively the control states and the input alphabet, and let 2^N be the number of cells used in each configuration. A configuration encoding is a sequence of the form $\langle 0 \rangle \sigma_0 \dots \langle 2^N \rangle \sigma_{2^N}$ where there is a i s.t. $\sigma_i \in Q \times \Sigma_i$ (this denotes the current state, the symbol of cell i and that the tape head is on cell i), $\sigma_j \in \Sigma$ for all $j \neq i$ (symbol in cell j), and $\langle h \rangle$ is the binary encoding of h (cell number) over new symbols d_r^\top and d_r^\perp for $r \in [N]$ (d_r^\top is equivalent to 1 and d_r^\perp to 0 in the binary encoding). A path of a computation (computations of alternating TM can be seen as trees of configurations) is encoded as a sequence $C_0 d_0 \$ \dots C_i d_i \$ \dots$ where each C_i is a configuration (C_0 is initial) and d_i is the transition taken from C_i to C_{i+1} .

We construct an RGG G with two modules M_{in} and M_1 . In M_{in} , initially, pl_0 generates an encoding of an initial configuration, then, a transition is selected by pl_0 , if the initial state is existential, or by pl_1 , otherwise. In both cases, an end-of-configuration marker $\$$ is generated and then pl_0 is in charge to generate again a configuration encoding, and so on. A call to M_1 is placed before generating the first cell encoding of each configuration and after generating each cell encoding. In M_1 , pl_1 selects one among a series of actions that can either state that everything is fine (*ok*) or that some check is required (by raising one over nine objections). M_1 has only one exit.

The goal of pl_0 is to build an encoding of an accepting run of a TM \mathcal{A} on a given input word, while the goal of pl_1 is to point out errors in such encoding by raising objections to delimit the cell encodings where the check has to take place. There are two possible mistakes that can occur: the i^{th} cells of two consecutive configurations do not conform the transition relation of \mathcal{A} and the number of a cell is not the successor of the number of the preceding cell in the configuration. We use separate groups of objections to point out each of these mistakes. These specifications can be captured with a formula φ defined as $\{\neg \diamond obj \vee [(\psi_{wr1} \vee \psi_\Delta) \wedge (\psi_{wr2} \vee \psi_\#)]\} \wedge (\diamond obj \vee \diamond F)$ where: (1) *obj* denotes that an objection has been raised; (2) F denotes a state predicate that is true on symbols of the encoding that correspond to final states of \mathcal{A} ; (3) ψ_{wr1} and ψ_{wr2} capture all the illegal uses of respectively the first and the second type of objections; (4) ψ_Δ checks the transition relation between two consecutive configurations on the cells selected by the raised objections of type 1; and (5) $\psi_\#$ checks the correct encoding of the cell numbers of two consecutive cells selected by the raised objections of type 2. All the above formulas can be written with disjunctions of path formulas except for $\psi_\#$ that is a conjunction of disjunctions of path formulas. Using De Morgan laws, the total formula can be transformed into an equivalent formula of size polynomial in $|\varphi|$, which is a disjunction of conjunctions of path formulas. Also note that all the used path formulas are of bounded size (the most complex one uses eleven occurrences of \diamond). Moreover, in a modular strategy pl_0 cannot use the fact that pl_1 has raised an objection to decide the next move since the objections are raised in a different module (which has just one exit). Therefore, in order to win, pl_0 must correctly generate the computations of the TM. We get the following:

Lemma 6. *The MVPG problem with winning conditions expressed as a conjunction of disjunctions of bounded-size PATH-LTL formulas is 2EXPTIME-hard.*

It is known that each formula φ from PATH-LTL admits a deterministic Büchi word automaton accepting all the models of φ and which is linear in its size [4]. The same can be shown with Büchi

VPA, by extending PATH-LTL allowing the versions of the \diamond operator of CARET and NWTL. By the closure properties of universal visibly pushdown automata we can easily extend Theorem 3 to winning conditions given as intersection of deterministic VPAs and thus:

Theorem 7. *The MVPG problem with winning conditions expressed as a conjunction of CARET and NWTL formulas that admit a deterministic Büchi or co-Büchi VPA generator of polynomial size is EXPTIME-complete.*

Now consider the larger fragment of formulas $\bigvee_{i=1}^h \bigwedge_{j=1}^k \varphi_{i,j}$ where for each $\varphi_{i,j}$ we can construct either a deterministic Büchi or a deterministic co-Büchi VPA $P_{i,j}$ of polynomial size that generates all the models of $\varphi_{i,j}$. We are only able to show an EXPTIME lower bound using a construction similar to that used in the reduction of Lemma 6. However, we observe that a matching upper bound cannot be shown using automata constructions, since we would need to manage the union of N specifications without an exponential blow-up, and since intersection is easy, this would contradict Lemma 6.

6 Improving the tree automata construction to solve ω -MGP with Büchi condition

We assume the standard definitions of trees and nondeterministic/universal tree automata with Büchi and co-Büchi acceptance (universality refers to the fact that all runs must be accepting in order to accept. See [5] for definitions).

General structure of the construction. Fix a ω -MG $\langle G, \mathcal{B} \rangle$ where $\mathcal{B} = (Q, q_0, \Sigma, \delta, F)$ is a deterministic Büchi automaton and G is as in Section 2. We construct a Büchi tree automaton $\mathcal{A}_{G, \mathcal{B}}$ that accepts a tree if and only if pl_0 has a winning modular strategy in the game $\langle G, \mathcal{B} \rangle$.

The trees accepted by $\mathcal{A}_{G, \mathcal{B}}$ must encode G and a modular strategy on it (*strategy trees*). Each such tree essentially has a subtree rooted at a child of the root for each module of G and each such subtree is the unwinding of the corresponding module along with a labeling encoding the strategy.

The general idea is to check on each subtree of the root some properties of the corresponding local function of the encoded modular strategy, by assuming some other properties on the local functions of the other modules (as in an assume-guarantee reasoning). These assumptions concern: a call structure CG (*Büchi call graph*), to handle acceptance on plays involving infinitely many unreturned calls; a set \mathcal{E} for each module, each giving a superset of the exits that can be visited during the plays; a set of extended pre-post conditions \mathcal{C} , that for each module m , each exit x and each possible state q of \mathcal{B} , carries the requirement that if m is entered with q and the play exits at x , then this must happen with state q' such that $(m, q, x, q') \in \mathcal{C}$ and after visiting an accepting state if this is required by \mathcal{C} . To ensure correctness all modules must share the same assumptions, thus $\mathcal{A}_{G, \mathcal{B}}$ guesses the assumptions at the root and then passes them onto the children of the root.

The tasks of $\mathcal{A}_{G, \mathcal{B}}$ thus are: recognizing the strategy trees; ensuring the correctness of the extended pre-post conditions; ensuring that all the plays according to the modular strategy conform the Büchi call graph and do not exit a module from an exit different from those listed in \mathcal{E} ; checking the acceptance condition of \mathcal{B} on all the plays encoded in the strategy tree, using the pre-post condition, the acceptance condition of \mathcal{B} and the Büchi call graph.

The above tasks are split among the Büchi tree automata \mathcal{A}_G , $\mathcal{A}_{CG}^{\mathcal{E}}$ and $\mathcal{A}_{\mathcal{B}, \mathcal{C}, CG}$. The automata \mathcal{A}_G and $\mathcal{A}_{CG}^{\mathcal{E}}$ are nondeterministic and check the fulfillment of the properties related to the game graph. \mathcal{A}_G is in charge of verifying that the input tree is a valid strategy tree. $\mathcal{A}_{CG}^{\mathcal{E}}$ is parameterized over the set of exits \mathcal{E} and a Büchi call graph CG . The automaton $\mathcal{A}_{\mathcal{B}, \mathcal{C}, CG}$, which is a universal Büchi tree

automaton, checks the extended pre-post condition \mathcal{C} , simulates \mathcal{B} and checks the fulfillment of its winning conditions.

$\mathcal{A}_{G,\mathcal{B}}$ captures the intersection of \mathcal{A}_G and the automaton that, at the root, nondeterministically guesses \mathcal{E} , CG , \mathcal{C} , and then at the children of the root, captures the intersection of $\mathcal{A}_{\mathcal{B},\mathcal{C},CG}$ and $\mathcal{A}_{CG}^{\mathcal{E}}$.

The automaton \mathcal{A}_G . Let k be the maximum over the number of exits of G modules and the out-degree of G vertices. Denote with Ω_G the set $\{\text{dummy}, \text{root}\} \cup (V \setminus P^0) \cup (P^0 \times [k])$ (recall V denotes the set of vertices of G). \mathcal{A}_G accepts strategy trees, i.e., Ω_G -labeled k -trees that encode modular strategies of pl_0 .

Intuitively, in a strategy tree, the label root is associated with the root of the tree. The children of the root are labeled with the entries of each module in G . A subtree rooted in one of these vertices corresponds to the unrolling of a module. If a vertex is labeled with a node that belongs to pl_0 , the move according to the encoded strategy is annotated with the index of the selected successor. If a node is associated to a call, then its children are labeled with the matching returns. The *dummy* nodes are used to complete the k -tree. A similar formal definition is given [5].

In Fig. 4 we depict the top fragment of a strategy tree for pl_0 of the RGG from Fig. 1




Figure 4: A fragment of a strategy tree.

Given a tree T , the automaton \mathcal{A}_G accepts T iff it is a strategy tree for G . A construction for \mathcal{A}_G can be easily obtained from G , and thus we omit it (see [5] for a similar construction).

Proposition 8. *There exists an effectively constructible Büchi tree automaton of size $O(|G|)$ that accepts a Ω_G -labeled k -tree if and only if it is a strategy tree.*

Directly from the definitions, the following holds:

Proposition 9. *For a ω -MG and fixed a player pl , there exists a one-to-one mapping between the modular strategies of pl and the strategy trees.*

The automaton $\mathcal{A}_{CG}^{\mathcal{E}}$. Fix a modular strategy f of pl_0 in G . A *call graph* (of G) according to f is a directed graph (V, \rightarrow) , $V \subseteq M$, such that for each play π which conforms to f , if a module $m \in M$ is reachable on π then $m \in V$ and if a call from m' to m'' is done on π then $m' \rightarrow m''$ holds.

A *Büchi call graphs* CG of G according to f and \mathcal{B} is $(V, \rightarrow, \rightarrow_F)$ where $V \subseteq M \times Q$, $\rightarrow_F \subseteq \rightarrow$ and: (1) denoting with $\xi((m, q)) = m$, the graph defined by all the edges $\xi(v) \rightarrow' \xi(v')$ s. t. $v \rightarrow v'$ is a call graph of G according to f (denoted $\xi(CG)$ in the following), and (2) for each cycle $v_1 \rightarrow v_2 \rightarrow \dots \rightarrow v_h \rightarrow v_{h+1}$ (with $v_1 = v_{h+1}$), there exists at least a $j \in [1, h]$ such that $(v_j \rightarrow_F v_{j+1})$ (in order to fulfill a Büchi condition, in a cycle of calls there must exist at least a module m_{i_j} among m_{i_1}, \dots, m_{i_h} in which an F state of \mathcal{B} is visited infinitely often).

In a strategy tree, a node is *enabled* if it is the root or the corresponding vertex of G is reachable in the encoded strategy. For a given Büchi call graph CG and a selection of exits \mathcal{E} , by assuming that the input tree is a strategy tree, the automaton $\mathcal{A}_{CG}^{\mathcal{E}}$ checks that indeed the input tree conforms to CG and \mathcal{E} .

Lemma 10. *There exists an effectively constructible Büchi tree automaton $\mathcal{A}_{CG}^{\mathcal{E}}$ such that if T is a strategy tree of G and f is the corresponding modular strategy, then $\mathcal{A}_{CG}^{\mathcal{E}}$ accepts T unless either (1) an exit not in \mathcal{E} is enabled, or (2) $\xi(CG)$ is not a call graph of G according to f (i.e., there exists an enabled call from m to m' but no edge from m to m' in $\xi(CG)$). The size of $\mathcal{A}_{CG}^{\mathcal{E}}$ is linear in $|G|$.*

The automaton $\mathcal{A}_{\mathcal{B},\mathcal{C},CG}$. The automaton $\mathcal{A}_{\mathcal{B},\mathcal{C},CG}$ is parameterized over the Büchi automaton \mathcal{B} , an extended pre-post condition \mathcal{C} (which essentially summarizes the effects of \mathcal{B} executions in each module of G) and a Büchi call-graph CG . It is quite complex and its tasks are:

1. to simulate \mathcal{B} on a strategy tree (and in this, it uses \mathcal{C} s.t. when simulating \mathcal{B} in a module it is not needed to follow the calls to other modules and the simulation can continue at a matching return);
2. to check the correctness of the pre-post condition \mathcal{C} ;
3. to check that the accepting states of \mathcal{B} are entered consistently with CG on the cycles of calls;
4. to check the fulfillment of \mathcal{B} acceptance conditions.

We first construct an automaton $\mathcal{A}_{\mathcal{B}}^{\mathcal{C}}$ which ensures task 1, then on the top of $\mathcal{A}_{\mathcal{B}}^{\mathcal{C}}$ we construct three different automata $\mathcal{A}_{\mathcal{B},\mathcal{C}}, \mathcal{A}_{\mathcal{B},CG}$ and $\mathcal{A}_{\mathcal{B},win}$, one for each of the remaining three tasks. We then get $\mathcal{A}_{\mathcal{B},\mathcal{C},CG}$ by taking the usual cross product for the intersection of these automata (note that an efficient construction can be obtained by discarding all the states that do not agree on the $\mathcal{A}_{\mathcal{B}}^{\mathcal{C}}$ part, thus avoiding a cubic blow-up in the size of $\mathcal{A}_{\mathcal{B},\mathcal{C},CG}$). Under the assumption that the input tree is a strategy tree and $\xi(CG)$ is consistent with it, we get that $\mathcal{A}_{\mathcal{B},\mathcal{C},CG}$ accepts only winning strategy trees (i.e., strategy trees that correspond to winning modular strategies) that conform to CG and \mathcal{C} . The details on all the above automata are given in the rest of this section. Thus, we get:

Lemma 11. *Let $\langle G, \mathcal{B} \rangle$ be a ω -MG and \mathcal{B} be a Büchi automaton. Given a Büchi call graph CG and a pre-post condition \mathcal{C} , there exists an effectively constructible universal Büchi tree automaton $\mathcal{A}_{\mathcal{B},\mathcal{C},CG}$ s.t.: if T is a strategy tree of G , f is the corresponding modular strategy and $\xi(CG)$ is a call graph of G according to f , then $\mathcal{A}_{\mathcal{B},\mathcal{C},CG}$ accepts T iff T is winning in $\langle G, \mathcal{B} \rangle$ and consistent with CG and \mathcal{C} . The size of $\mathcal{A}_{\mathcal{B},\mathcal{C},CG}$ is quadratic in the number of \mathcal{B} states.*

Winning strategy trees. For a strategy tree T of G , a *play* of T is an ω -sequence of T -nodes $x_1 x_2 \dots$ such that x_1 is the child of the root corresponding to m_{in} , $\alpha_1 = \varepsilon$ (*call-stack*) and for $i \in \mathbb{N}$, x_i is an enabled node and: (1) if x_i is labeled with a call to m , then x_{i+1} is the child of the T root corresponding to module m (and thus is labeled with the entry e_m), and $\alpha_{i+1} = \alpha_i.x_i$; (2) if x_i is labeled with an exit ex , then $\alpha_i = \alpha_{i+1}.y$, y is a node labeled with a call (b, e_m) and x_{i+1} is the child of y labeled with the return (b, ex) ; (3) otherwise, x_{i+1} is an enabled child of x_i and $\alpha_{i+1} = \alpha_i$.

Note that any play of a strategy tree T corresponds to a play of G (conforming to the modular strategy defined by T). A *winning strategy tree* T w.r.t. \mathcal{B} is such that for all the plays v of T , w_v is accepted by \mathcal{B} . From Proposition 9, we get:

Lemma 12. *Given a ω -MG $\langle G, \mathcal{B} \rangle$, a modular strategy is winning iff the corresponding strategy tree is winning.*

Pre-post conditions. A *pre-post condition* on the graph G is a pair $\langle \mathcal{C}_{pre}, \mathcal{C}_{post} \rangle$ where $\mathcal{C}_{pre} \subseteq M \times Q$ (set of *pre-conditions*), $\mathcal{C}_{post} \subseteq M \times Q \times Ex \times Q$ (set of *pre-post conditions*), and such that for each $(m, q, ex_j, q') \in \mathcal{C}_{post}$, also $(m, q) \in \mathcal{C}_{pre}$ (i.e., tuples of \mathcal{C}_{post} add a post-condition to some of the pre-conditions of \mathcal{C}_{pre}).

Intuitively, a pre-post condition is meant to summarize all the \mathcal{B} locations that can be reached on entering each module of G along any play of T , and for each reachable exit ex , all the pairs of \mathcal{B} locations (q, q') s.t. there exists a play of T along with \mathcal{B} enters a module at q and exits it from ex at q' .

Fix a pre-post condition $\langle \mathcal{C}_{pre}, \mathcal{C}_{post} \rangle$. $\langle \mathcal{C}_{pre}, \mathcal{C}_{post} \rangle$ is *consistent* with a strategy tree T if for each play $v = x_1 x_2 \dots$ of T and for each x_i which is labeled with a call to module $m \in M$, denoting with q the location at which the only run of \mathcal{B} over w_{v_i}

ends: (1) (m, q) belongs to \mathcal{C}_{pre} and (2) if v reaches the matching return at x_{j+1} , x_j is labeled with exit ex and the location at which the run of \mathcal{B} over $w_{v_{j_1}}$ is q' (i.e., the location when reading the symbol of ex at x_j), then (m, q, ex, q') belongs to \mathcal{C}_{post} .

$\langle \mathcal{C}_{pre}, \mathcal{C}_{post} \rangle$ is *consistent* with a set of exits $\langle \mathcal{E}^m \rangle_{m \in M}$ iff: (1) $\forall m \in M, \forall ex \in \mathcal{E}^m$, if there is a $(m, q) \in \mathcal{C}_{pre}$ then there is at least a tuple of the form $(m, q, ex, q') \in \mathcal{C}_{post}$; (2) $\forall (m, q, ex, q') \in \mathcal{C}_{post}$, then $ex \in \mathcal{E}^m$ holds.

We extend pre-post conditions with a function $Fin : \mathcal{C}_{post} \rightarrow \{true, false\}$. An *extended pre-post condition* $\mathcal{C} = \langle \mathcal{C}_{pre}, \mathcal{C}_{post}, Fin \rangle$ is *consistent* with a strategy tree T if $\langle \mathcal{C}_{pre}, \mathcal{C}_{post} \rangle$ is consistent with T and for each play $v = x_1 x_2 \dots$ of T s.t. x_i is labeled with a call to module $m \in M$, x_{j+1} is labeled with its matching return, and the portion of run of \mathcal{B} from i to j starts at location q and ends at location q' : whenever $Fin(m, q, ex, q') = true$ then a location in F must be visited on this portion of run (acceptance-condition).

Construction of $\mathcal{A}_{\mathcal{B}}^{\mathcal{C}}$. Fix an extended pre-post condition $\mathcal{C} = \langle \mathcal{C}_{pre}, \mathcal{C}_{post}, Fin \rangle$. Let m_i be the module mapped to the i^{th} child of the root of a strategy tree.

We construct $\mathcal{A}_{\mathcal{B}}^{\mathcal{C}}$ such that the automaton simulates \mathcal{B} on an input strategy tree T by using \mathcal{C} . In particular, starting from the i^{th} child, the automaton $\mathcal{A}_{\mathcal{B}}^{\mathcal{C}}$ runs in parallel a copy of \mathcal{B} from each control state q such that $(m_i, q) \in \mathcal{C}_{pre}$. When reading a node labeled with a call, $\mathcal{A}_{\mathcal{B}}^{\mathcal{C}}$ starts at each matching return a copy of \mathcal{B} according to the applicable tuples in \mathcal{C} and performs updates according to Fin . On all the other enabled nodes, the control state of \mathcal{B} is updated for each copy according to \mathcal{B} transitions.

The states of $\mathcal{A}_{\mathcal{B}}^{\mathcal{C}}$ are: an initial state q_0 , an accepting state q_a , a rejecting state q_r , and states of the form $(q, d, f, q_{m_i}, \mathcal{C})$ where $q, q_{m_i} \in Q$, q is the control state which is updated in the simulation of \mathcal{B} , q_{m_i} is the current pre-condition, and $d, f \in \{0, 1\}$ are related to the winning conditions. Namely, d is used to check the acceptance-conditions along all the plays that conform to the strategy, and f is used to expose that a final state of \mathcal{B} was seen between a call and its matching return. A task of $\mathcal{A}_{\mathcal{B}}^{\mathcal{C}}$ is to handle the correct update of these bits, but they will be used to determine the acceptance by $\mathcal{A}_{\mathcal{B}_{win}}$. The states q_a and q_r are sinks, i.e., once reached, the automaton cycles forever on them.

Construction of $\mathcal{A}_{\mathcal{B}, \mathcal{C}}$, $\mathcal{A}_{\mathcal{B}, CG}$ and $\mathcal{A}_{\mathcal{B}_{win}}$. The automaton $\mathcal{A}_{\mathcal{B}, \mathcal{C}}$ is in charge of checking that the input tree is consistent with the extended pre-post condition \mathcal{C} . We construct it from $\mathcal{A}_{\mathcal{B}}^{\mathcal{C}}$ by modifying the transitions from a state s of the form $(q, d, f, q_{m_i}, \mathcal{C})$ at a tree-node labeled with an exit. In particular, in this case, we let $\mathcal{A}_{\mathcal{B}, \mathcal{C}}$ enter q_a if there exists a tuple $(m_i, q_{m_i}, ex_h, q) \in \mathcal{C}_{post}$, and q_r otherwise.

The purpose of $\mathcal{A}_{\mathcal{B}, CG}$ is to check that the Büchi call graph CG is indeed consistent with the input strategy tree. We construct $\mathcal{A}_{\mathcal{B}, CG}$ from $\mathcal{A}_{\mathcal{B}}^{\mathcal{C}}$ by modifying the transitions from calls. Namely, when on a node u labeled with a call to a module m_j in the subtree of the i^{th} child of the root, at a state of the form $(q', d, f, q, \mathcal{C})$ and suppose there is a transition $(q', \eta_{m_i}(u), q'') \in \delta$: if $(m_i, q) \rightarrow_F (m_j, q'')$ holds in the Büchi call graph CG and $d = 0$, then $\mathcal{A}_{\mathcal{B}, CG}$ enters the rejecting state q_r . In all the other cases it behaves as $\mathcal{A}_{\mathcal{B}}^{\mathcal{C}}$.

The purpose of $\mathcal{A}_{\mathcal{B}_{win}}$ is to check that the winning conditions of \mathcal{B} are satisfied along all plays of the input strategy tree. Again, we can modify $\mathcal{A}_{\mathcal{B}}^{\mathcal{C}}$ to ensure this. In particular, when the automaton reaches a tree-node labeled with exit ex of module m in a state $(q', d, f, q, \mathcal{C})$, then it enters the accepting state q_a , whenever $(m, q, ex, q') \in \mathcal{C}_{post}$ and $Fin(m, q, ex, q') = true$ implies $b = 1$, and q_r otherwise. Moreover the accepting states of $\mathcal{A}_{\mathcal{B}_{win}}$ are q_a and all the states of the form $(q, d, f, q', \mathcal{C})$ such that either $q \in F$ or $f = 1$.

Reducing to emptiness of Büchi tree automata. $\mathcal{A}_{\mathcal{B}, \mathcal{C}, CG}$ can be translated into an equivalent non-deterministic Büchi tree automaton with $2^{O(|Q|^2 \log |Q|)}$ states [16]. Denoting with k the number of G exits

and β the number of *call edges* of G (i.e., the number of module pairs (m, m') such that there is a call from m to m'), the number of different choices for an extended pre-post condition is $2^{O(k|Q|^2)}$, for a Büchi call graph is 2^{2^β} , and for a set of exits \mathcal{E} is 2^k . Since \mathcal{A}_G and $\mathcal{A}_{CG}^{\mathcal{E}}$ are both of size $O(|G|)$, the automaton $\mathcal{A}_{G,\mathcal{B}}$ (obtained as described earlier in this section) is of size $|G|^2 2^{O(|Q|^2(k+\log|Q|)+\beta)}$. We can reduce the factor $|G|^2$ to $|G|$, by combining \mathcal{A}_G and $\mathcal{A}_{CG}^{\mathcal{E}}$ into the same automaton (they are essentially based on G transitions). Therefore, we can get an efficient construction of $\mathcal{A}_{\mathcal{B},CG}$ of size $|G| 2^{O(|Q|^2(k+\log|Q|)+\beta)}$. Thus, by Propositions 8 and 9, and Lemmas 10 and 11

Theorem 13. *For an RGG G and a deterministic Büchi automaton \mathcal{B} pl_0 has a winning modular strategy in $\langle G, \mathcal{B} \rangle$ iff the nondeterministic Büchi tree automaton $\mathcal{A}_{G,\mathcal{B}}$ accepts a non-empty language. Moreover, $\mathcal{A}_{G,\mathcal{B}}$ is of size $|G| 2^{O(|Q|^2(k+\log|Q|)+\beta)}$, where k is the number of G exits.*

7 Discussion

In this paper, we have considered modular games with winning condition expressed by pushdown, visibly pushdown and temporal logic specifications. We have proved that the modular game problem with respect to standard pushdown specifications is undecidable. Then we have presented a number of results that give a quite accurate picture of the computational complexity of the MVPG problem with visibly pushdown winning conditions. With some surprise, we have found that MVPG with temporal logic winning conditions becomes immediately hard. In fact, while the complexity for LTL specifications is 2EXPTIME-complete both for MVPG and games on finite graphs, for the fragment consisting of all the Boolean combinations of PATH-LTL formulas, solving the corresponding games on finite graphs is PSPACE-complete while the MVPG problem is already 2EXPTIME-complete. As a consequence, the computational complexity of many interesting fragments of LTL, that have a better complexity than full LTL on finite game graphs, collapses at the top of the complexities (see [4, 6]). This also differs with the scenario of the complexities of model-checking RSMs in LTL fragments (see [12]). As a final remark, we observe that the tree automaton construction proposed in Section 6 can be easily adapted to handle visibly pushdown winning conditions to get a direct solution of the MVPG problem. We only need to modify the transition rules to synchronize the calls and returns of the RGG with the pushes and pops of the specification automaton, and this would be possible since they share the same visibly alphabet. The change does not affect the overall complexity, however it will slightly improve on the approach presented in Section 4 that causes doubling the number of exits and gives a complexity with an exponential dependency in the number of stack symbols.

References

- [1] Rajeev Alur, Marcelo Arenas, Pablo Barceló, Kousha Etessami, Neil Immerman & Leonid Libkin (2008): *First-Order and Temporal Logics for Nested Words*. *Logical Methods in Computer Science* 4(4), doi:10.2168/LMCS-4(4:11)2008.
- [2] Rajeev Alur, Michael Benedikt, Kousha Etessami, Patrice Godefroid, Thomas W. Reps & Mihalis Yannakakis (2005): *Analysis of recursive state machines*. *ACM Trans. Program. Lang. Syst.* 27(4), pp. 786–818, doi:10.1145/1075382.1075387.
- [3] Rajeev Alur, Kousha Etessami & P. Madhusudan (2004): *A Temporal Logic of Nested Calls and Returns*. In Kurt Jensen & Andreas Podelski, editors: *TACAS, Lecture Notes in Computer Science* 2988, Springer, pp. 467–481, doi:10.1007/978-3-540-24730-2_35.

- [4] Rajeev Alur & Salvatore La Torre (2004): *Deterministic generators and games for Ltl fragments*. *ACM Trans. Comput. Log.* 5(1), pp. 1–25, doi:10.1145/963927.963928.
- [5] Rajeev Alur, Salvatore La Torre & P. Madhusudan (2003): *Modular Strategies for Infinite Games on Recursive Graphs*. In Warren A. Hunt Jr. & Fabio Somenzi, editors: *CAV, Lecture Notes in Computer Science 2725*, Springer, pp. 67–79, doi:10.1007/978-3-540-45069-6_6.
- [6] Rajeev Alur, Salvatore La Torre & P. Madhusudan (2003): *Playing Games with Boxes and Diamonds*. In Roberto M. Amadio & Denis Lugiez, editors: *CONCUR, Lecture Notes in Computer Science 2761*, Springer, pp. 127–141, doi:10.1007/978-3-540-45187-7_8.
- [7] Rajeev Alur, Salvatore La Torre & P. Madhusudan (2006): *Modular strategies for recursive game graphs*. *Theor. Comput. Sci.* 354(2), pp. 230–249, doi:10.1016/j.tcs.2005.11.017.
- [8] Rajeev Alur & P. Madhusudan (2009): *Adding nesting structure to words*. *J. ACM* 56(3), doi:10.1145/1516512.1516518.
- [9] Benjamin Aminof, Fabio Mogavero & Aniello Murano (2014): *Synthesis of hierarchical systems*. *Sci. Comput. Program.* 83, pp. 56–79, doi:10.1016/j.scico.2013.07.001.
- [10] Ilaria De Crescenzo & Salvatore La Torre (2013): *Modular Synthesis with Open Components*. In Parosh Aziz Abdulla & Igor Potapov, editors: *RP, Lecture Notes in Computer Science 8169*, Springer, pp. 96–108, doi:10.1007/978-3-642-41036-9_10.
- [11] William R. Harris, Somesh Jha & Thomas W. Reps (2012): *Secure Programming via Visibly Pushdown Safety Games*. In P. Madhusudan & Sanjit A. Seshia, editors: *CAV, Lecture Notes in Computer Science 7358*, Springer, pp. 581–598, doi:10.1007/978-3-642-31424-7_41.
- [12] Salvatore La Torre & Gennaro Parlato (2007): *On the Complexity of LtlModel-Checking of Recursive State Machines*. In Lars Arge, Christian Cachin, Tomasz Jurdzinski & Andrzej Tarlecki, editors: *ICALP, Lecture Notes in Computer Science 4596*, Springer, pp. 937–948, doi:10.1007/978-3-540-73420-8_80.
- [13] Christof Löding, P. Madhusudan & Olivier Serre (2004): *Visibly Pushdown Games*. In Kamal Lodaya & Meena Mahajan, editors: *FSTTCS, Lecture Notes in Computer Science 3328*, Springer, pp. 408–420, doi:10.1007/978-3-540-30538-5_34.
- [14] Yoad Lustig & Moshe Y. Vardi (2009): *Synthesis from Component Libraries*. In Luca de Alfaro, editor: *FOSSACS, Lecture Notes in Computer Science 5504*, Springer, pp. 395–409, doi:10.1007/978-3-642-00596-1_28.
- [15] Yoad Lustig & Moshe Y. Vardi (2011): *Synthesis from Recursive-Components Libraries*. In Giovanna D’Agostino & Salvatore La Torre, editors: *GandALF, EPTCS 54*, pp. 1–16, doi:10.4204/EPTCS.54.1.
- [16] David E. Muller & Paul E. Schupp (1995): *Simulating Alternating Tree Automata by Nondeterministic Automata: New Results and New Proofs of the Theorems of Rabin, McNaughton and Safra*. *Theor. Comput. Sci.* 141(1&2), pp. 69–107, doi:10.1016/0304-3975(94)00214-4.
- [17] Amir Pnueli (1977): *The Temporal Logic of Programs*. In: *FOCS*, IEEE Computer Society, pp. 46–57, doi:10.1109/SFCS.1977.32.
- [18] Amir Pnueli & Roni Rosner (1989): *On the Synthesis of a Reactive Module*. In: *POPL*, ACM Press, pp. 179–190, doi:10.1145/75277.75293. Available at <http://dl.acm.org/citation.cfm?id=75277>.
- [19] Sylvain Salvati & Igor Walukiewicz (2013): *Evaluation is MSOL-compatible*. In Anil Seth & Nisheeth K. Vishnoi, editors: *FSTTCS, LIPIcs 24*, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 103–114, doi:10.4230/LIPIcs.FSTTCS.2013.103. Available at <http://drops.dagstuhl.de/portals/extern/index.php?semnr=13018>.
- [20] Wolfgang Thomas (2002): *Infinite Games and Verification (Extended Abstract of a Tutorial)*. In Ed Brinksma & Kim Guldstrand Larsen, editors: *CAV, Lecture Notes in Computer Science 2404*, Springer, pp. 58–64, doi:10.1007/3-540-45657-0_5.
- [21] Igor Walukiewicz (2001): *Pushdown Processes: Games and Model-Checking*. *Inf. Comput.* 164(2), pp. 234–263, doi:10.1006/inco.2000.2894.