



HAL
open science

A Class of Zielonka Automata with a Decidable Controller Synthesis Problem

Hugo Gimbert

► **To cite this version:**

Hugo Gimbert. A Class of Zielonka Automata with a Decidable Controller Synthesis Problem. 2016.
hal-01259151v4

HAL Id: hal-01259151

<https://hal.science/hal-01259151v4>

Preprint submitted on 31 Jan 2016 (v4), last revised 3 Aug 2017 (v12)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Class of Zielonka Automata with a Decidable Controller Synthesis Problem

Hugo Gimbert

LaBRI, CNRS, Université de Bordeaux, France
hugo.gimbert@cnrs.fr

Abstract

The decidability of the distributed version of the Ramadge and Wonham control problem (Ramadge and Wonham 1989), where both the plant and the controllers are modelled as Zielonka automata (Zielonka 1987; Diekert and Rozenberg 1995) is a challenging open problem (Muscholl et al. 2008).

There exists three classes of plants for which the existence of a correct controller has been shown decidable in the distributed setting: when the dependency graph of actions is series-parallel, when the processes are connectedly communicating and when the dependency graph of processes is a tree.

We generalize these three results by showing that a larger class of plants, called broadcast plants, has a decidable controller synthesis problem. We give new examples of plants for which controller synthesis is decidable.

1. Introduction

The decidability of the distributed version of the Ramadge and Wonham control problem (Ramadge and Wonham 1989), where both the plant and the controllers are modelled as Zielonka automata (Zielonka 1987; Diekert and Rozenberg 1995) is a challenging open problem. A very good introduction to the distributed controller synthesis problem is given in (Muscholl et al. 2008).

We assume that the plant is distributed on several finite-state processes which interact asynchronously using shared actions. On every process, the local controller can choose to block some of the actions, called controllable actions, but he cannot block the uncontrollable actions from the environment. The choice of the local controller is based on several sources of information: first he can observe the local sequence of states and actions of the process and second when a shared action is played all the local controllers of the corresponding processes can exchange as much information as they want, in particular together they can compute their mutual view of the global execution.

Assuming that processes can exchange information upon synchronization is a game changer from the point of view of decidability. Actually, in the setting of (Pnueli and Rosner 1990), distributed synthesis is not decidable except for very simple architectures like the pipeline architecture. The comparison between the two assumptions is discussed in (Gastin et al. 2004). The pa-

per (Finkbeiner and Schewe 2005) proposes information forks as an uniform notion explaining the (un)decidability results in distributed synthesis.

A correct controller restricts controllable actions so that every possible execution of the plant satisfies some specification. In the present paper we focus on local reachability conditions: the set of plays authorized by the controller should be finite, and in each maximal play all processes should be in a final state.

We adopt a modern terminology and call the plant a game and the controllers are strategies in this game. Of course strategies should be distributed and the choice of actions to restrict should depend only of its local view of the global execution. Our goal is to decide, given a distributed game, whether there exists a winning strategy for the controllers in the game, which guarantees the local reachability condition to hold.

There exists three classes of plants for which the existence of a winning distributed strategy has been shown decidable: when the dependency graph of actions is series-parallel, when the processes are connectedly communicating and when the dependency graph of processes is a tree.

Connectedly communicating games have been introduced (Madhusudan et al. 2005) under the name of *connectedly communicating processes*. Intuitively, a game is connectedly communicating if there is a bound k such that if a process p executes k steps without hearing from process q , directly or indirectly, then p will never hear from q again. The event structure of a connectedly communicating games has a decidable MSO theory (Madhusudan et al. 2005) which implies that controller synthesis is decidable for these games.

A series-parallel game is a game such that the dependence graph (A, D) of the alphabet A is a co-graph. Series-parallel games were proved decidable in (Gastin et al. 2004), for a different setup than ours: in the present paper we focus on process-based control while (Gastin et al. 2004) was focusing on action-based synthesis. Actually action-based control is more general than process-based control, see (Gastin et al. 2004) for a proof. The results of the present paper could probably be extended to action-based control however we prefer to stick to process-based control in order to keep the model intuitive. To our knowledge, the result of (Gastin et al. 2004) was the first discovery of a class of asynchronous distributed system for which controller synthesis is decidable

An acyclic game as defined in (Genest et al. 2013) is a game where processes are arranged as a tree and actions are either local or synchronize a father and his son. Formally, the processes are arranged as a tree $T_{\mathbb{P}} = (\mathbb{P}, E_{\mathbb{P}})$, and each action is either a local action whose domain is a singleton or a synchronizing action such that $\text{dom}(a) = \{p, q\}$ and $(p, q) \in E_{\mathbb{P}}$ i.e. q is the father of p in the process tree.

We generalize these three results by showing that a larger class of games, called broadcast games, has a decidable controller synthesis problem, this is our main result. We give new examples

of distributed games where the existence of a winning strategy is decidable: acyclic games with arbitrary actions and triangulated games.

The proof of decidability of broadcast games is fairly simple and intuitive, at least when one is familiar with notations and concepts from Zielonka automata and distributed synthesis. In a nutshell, we start with a winning strategy and look for useless parts that we can remove in order to get a simpler strategy. These parts are called *useless threads*. Whenever a useless thread exists, we remove it using an operation called a *shortcut* in order to get a simpler strategy. Intuitively, a shortcut is like a cut and paste operation in the strategy from a point A to a point B, such that B is an ancestor of A, which makes the strategy smaller. By taking shortcuts again and again, we make the strategy smaller and smaller, until it does not have useless thread anymore. Strategies with no useless threads have bounded size and they can be enumerated which leads to decidability.

Performing cut-and-paste in a distributed strategy is not as easy as doing it in a synchronous game. In a synchronous game with only one process, strategies are trees and one can cut a subtree from a point A and paste it to an ancestor B of A. As long as the unique process is in the same state at A and B this will define another, shorter, strategy. The decidability of series-parallel games established (Gastin et al. 2004) relies also on some simplification of the winning strategies, in order to get *uniform* strategies. The series-parallel assumption is used to guarantee that the result of the replacement of a part of a strategy by a uniform strategy is still a strategy, as long as the states of all processes coincide. But in the case of a general distributed strategy, without series-parallel assumption, it is not sufficient that states of the processes coincide at the source and the destination, one has also to take into account the parallelism and the various information of the different processes, so that the result of the operation is still a distributed strategy.

This is the reason for introducing the notion of *broadcasts*. A broadcast is a part of a strategy where a pool of processes can synchronize and one of the processes of the pool can broadcast an information to the others such that this information is received by each process of the pool before it synchronizes with other processes outside the pool. When two broadcasts are similar in some sense made precise in the proof of the theorem, they can be used to perform cutting and pasting and create shortcuts in the strategy: upon arrival on A, a process of the pool broadcasts to other processes of the pool that they should jump to B, and play as if the path from A to B had been already taken.

The transformation of an arbitrary winning strategy to a simpler one is done by induction on the set of actions, which relies on a notion of *inductive decomposition* of the set of actions. This notion is useful to derive rather easily from our techniques the three known decidability results we mentioned above. However to our opinion the technical core of the paper is not this notion of inductive decomposition but rather the notion of useless threads and shortcuts, and the proof that taking a shortcut of a useless thread turns a distributed strategy into another distributed strategy.

The complexity of our algorithm is really bad, so it is not clear whether this work will have practical applications. This is not surprising since the problem is non-elementary (Genest et al. 2013). Nevertheless we think our contribution has some theoretical interest, since it sheds new light on the difficult open problem of distributed synthesis.

2. Definitions and basic properties

2.1 Mazurkiewicz traces

The theory of Mazurkiewicz traces is very rich and extensively developed in (Diekert and Rozenberg 1995). Here we only fix notations and recall the notions of traces, prime traces and views, and list a few elementary properties of traces that we will use throughout the paper.

We fix an alphabet A and a symmetric and reflexive dependency relation $D \subseteq A \times A$ and the corresponding independency relation $\mathbb{I} \subseteq A \times A$ defined by:

$$\forall a, b \in A, (a \mathbb{I} b) \iff (a, b) \notin D .$$

For $u, v \in A^*$, we denote $A(u)$ the set of letters of u and we write:

$$u \mathbb{I} v$$

whenever $A(u) \times A(v) \subseteq \mathbb{I}$ and for $B \subseteq A$ we write

$$u \mathbb{I} B$$

whenever $\forall b \in B, b \mathbb{I} u$.

A Mazurkiewicz trace on (A, \mathbb{I}) is an equivalence class of words for the smallest equivalence relation \approx on A^* such that:

$$\forall u, v \in A^*, \forall a, b \in A, ((a \mathbb{I} b) \implies (uabx \approx ubax)) .$$

In most of the paper, a Mazurkiewicz trace is simply called a *trace*. A word in a trace is called a *linearization* of the trace.

The empty trace denoted ϵ is the singleton which contains only the empty word.

All words of a trace have the same alphabet, thus the notation $A(u)$ extends to traces. The length of a trace u , denoted $|u|$, is the number of letters of any linearization of u .

For a subset $B \subseteq A$, a trace on B is a trace u such that $A(u) \subseteq B$. We abuse the notation and from now on we denote B^* the set traces on an alphabet $B \subseteq A$. We will use the notation $\text{words}(B)$ to denote the set of finite words on B .

The concatenation on words naturally extends to traces, given two traces $u, v \in A^*$, the trace uv is the equivalence class of any word $u'v'$ such that $u' \in u$ and $v' \in v$. Also the notion of prefix extends to traces. A trace $u \in A^*$ is a prefix of a trace $v \in A^*$, denoted

$$u \sqsubseteq v$$

if there exists $w \in A^*$ such that $uw = v$. And u is a suffix of v if there exists $w \in A^*$ such that $v = wu$. Not all properties of the concatenation operator and the prefix relation on words are preserved on traces, however the following are:

$$\forall u, v \in A^*, ((u \sqsubseteq v) \wedge (v \sqsubseteq u) \implies u = v) , \quad (1)$$

$$\forall u, v, w \in A^*, (uv = uw) \implies (v = w) , \quad (2)$$

$$\forall u, v, w \in A^*, (uv \sqsubseteq uw) \implies (v \sqsubseteq w) . \quad (3)$$

A trace $u \in A^*$ is *prime* if all its linearization have the same last letter. If this letter is $a \in A$, i.e. if $u \in A^*a$, u is said to be *a-prime*. Let $B \subseteq A$. If all linearization of u ends up with a letter in B then u is said to be *B-prime*.

Let $B \subseteq A$ and $u \in A^*$. Then there exists a shortest prefix $\partial_B(u)$ of u , called the *B-view* and denoted

$$\partial_B(u) .$$

such that u factorizes as $u = \partial_B(u) \cdot v$ with $v \mathbb{I} B$. If B is a singleton $\{b\}$ then the *B-view* is also called the *b-view* and denoted $\partial_b(u)$.

The following lemma lists some basic properties of traces that we use later in the paper. In case the reader is already familiar with trace theory, these properties and their proofs will probably seem obvious to him. However, if the reader is new to trace theory,

proving these properties is a nice exercise to get familiar with basic notions of a prefix of a trace, prime traces, views and their interplay.

Lemma 1. *For every trace $u, v, x \in A^*$ and $a \in A$ and $B \subseteq A$,*

$$(u \parallel B) \iff (\partial_B(u) = \epsilon) \quad (4)$$

$$(x \sqsubseteq uv) \implies \exists x_0 \sqsubseteq u, \exists x_1 \sqsubseteq v, x = x_0x_1 \quad (5)$$

$$(x \sqsubseteq uv) \implies \exists x_0, x_1, x_2, x_3 \in A^*, \quad (6)$$

$$(x = x_0x_1) \wedge (u = x_0x_2) \wedge (v = x_1x_3) \wedge (x_2 \parallel x_1)$$

$$uv \text{ is } B\text{-prime} \implies v \text{ is } B\text{-prime} \quad (7)$$

$$u \text{ and } v \text{ are } B\text{-prime} \implies uv \text{ is } B\text{-prime} \quad (8)$$

$$\text{If } ua \text{ is prime, } (av \text{ is } B\text{-prime} \iff uav \text{ is } B\text{-prime}) \quad (9)$$

$$(u \text{ is } B\text{-prime} \wedge \neg(a \parallel u)) \implies (au \text{ is } B\text{-prime}) \quad (10)$$

$$(u \sqsubseteq \partial_B(uv)) \iff (\partial_B(uv) = u \partial_B(v)) \quad (11)$$

$$(uw \sqsubseteq \partial_B(uv)) \implies (w \sqsubseteq \partial_B(v)) \quad (12)$$

$$\partial_B(\partial_B(u)) = \partial_B(u) \quad (13)$$

$$\partial_B(uv) = \partial_B(u \partial_B(v)) \quad (14)$$

If ua is prime,

$$uav \sqsubseteq \partial_B(uavw) \iff av \sqsubseteq \partial_B(avw) \quad (15)$$

Proof. The equivalence (4) is immediate from the definition of ∂_B .

Equation (5) is a corollary of (6) which is well-known, see (Diekert and Rozenberg 1995) for example. It can be proved by induction on $|x|$.

We prove (7). If the last letter of a word $v' \in v$ is not in B , then the same holds for every $u'v'$ where $u' \in u$ thus uv is not B -prime since $u'v' \in uv$.

We prove (8). Assume both u and v are B -prime. Every linearization of uv is a shuffle of a linearization of u and a linearization of v thus it terminates with a letter in B . Hence uv is B -prime.

We prove (9). Assume ua prime. The converse implication follows from (7). Assume av is B -prime. We prove that uav is B -prime by induction on $|u|$. If $|u| = 0$ then $u = \epsilon$ and $uav = av$ is B -prime by hypothesis. By induction let $n \in \mathbb{N}$ and assume $u'av$ is B -prime for all u' such that $|u'| \leq n$. Let u such that $|u| = n + 1$, we prove that uav is B -prime. Since $|u| = n + 1$, there exists $b \in A$ and $u' \in A^*$ such that $u = bu'$ and $|u'| = n$. Using (7) and the induction hypothesis so on one hand we know that $u'av$ is B -prime. By definition of a trace, for any trace w ,

$$bw = \{xbz \mid x, z \in \text{words}(A), xz \in w, b \parallel x\}. \quad (16)$$

Let y a linearization of $uav = bu'av$, we prove that the last letter of y is in B . According to (16), y factorizes as $y = xbz$ with $xz \in u'av$ and $x \parallel b$. Since $xz \in u'av$ and $u'av$ is B -prime, if z is not empty then it ends with a letter in B and so does y . Assume now that z is empty, then $y = xb$ with $x \in u'av$ and $x \parallel b$. Since $y \in bu'av$ then $A(u'a) \subseteq A(y)$ and $A(v) \subseteq A(y)$. Since $A(y) = A(x) \cup \{b\}$ and $x \parallel b$ every letter of $u'a$ and av commute with b thus $bu'a = u'ab$ and $bv = vb$. Since $bu'a = ua$ is prime, $bu'a = u'ab$ implies $a = b$. Since $bv = vb$ then $av = va$ and since av is B -prime, $a = b \in B$. Finally $b \in B$ and since $y = xb$ the last letter of y is in B , which terminates the proof of the inductive step, and the proof of (9).

We prove (10) by contradiction. Assume au is not B -prime then there exists a word v' and a letter $c \notin B$ such that $v'c \in au$. Let $u' \in u$ then $au' \in au$ and $v'c \approx au'$ thus $A(v') \cup \{c\} = A(u') \cup \{a\}$. If $a \notin A(v')$ then $a = c$ and $v'c \approx au'$ implies $a \parallel u$ which is false by hypothesis. Thus $a \in A(v')$. Let w' be the longest prefix of v' which does not contain a and x' the suffix of v' such that $v' = w'ax'c$. Then $au' \approx w'ax'c$ and $a \notin A(w')$ thus $w' \parallel a$. Then $w'ax'c \approx aw'x'c$ thus $aw'x'c \approx au'$ hence $w'x'c \approx u'$

and $w'x'c \in u$. Since $c \notin B$, this contradicts the hypothesis u is B -prime.

We prove (11). The converse implication in (11) is obvious so it is enough to prove the direct implication. Assume $u \sqsubseteq \partial_B(uv)$. According to (1) it is enough to prove both $\partial_B(uv) \sqsubseteq u \partial_B(v)$ and $u \partial_B(v) \sqsubseteq \partial_B(uv)$. We start with $u \partial_B(v) \sqsubseteq \partial_B(uv)$. Since $u \sqsubseteq \partial_B(uv)$, then $\partial_B(uv) = uw$ for some $w \in A^*$ and $uv = uww'$ for some $w' \parallel B$. Then $v = ww'$ according to (2) and since $w' \parallel B$, then $\partial_B(v) \sqsubseteq w$, thus $u \partial_B(v) \sqsubseteq uw = \partial_B(uv)$ and we got the first prefix relation. Now we prove the converse prefix relation. Since $\partial_B(v) \sqsubseteq w$ then by definition of ∂_B there exists $w'' \in A^*$ such that $w = \partial_B(v)w''$ and $w'' \parallel B$. Then $uv = u \partial_B(v)w''w'$ and $w''w' \parallel B$ thus by definition of ∂_B , $\partial_B(uv) \sqsubseteq u \partial_B(v)$. By definition of w this implies $uw \sqsubseteq u \partial_B(v)$ thus according to (3) $w \sqsubseteq \partial_B(v)$. Finally $w = \partial_B(v)$ and $u \partial_B(v) = uw = u \partial_B(v)$ which terminates the proof of (11).

Equation (12), is a direct corollary of (11). Let v', w' such that $\partial_B(uv) = uww'$ and $uv = \partial_B(uv)w'$. Then according to (11), $uww' = u \partial_B(vw'w')$ thus according to (2), $ww' = \partial_B(vw'w')$ and since $v = vw'w'$, we get $w \sqsubseteq \partial_B(v)$.

By definition $\partial_B(uv)$ is the shortest prefix of uv such that $uv = \partial_B(uv)v'$ with $v' \parallel B$, thus by hypothesis there exists w' such that $uv = uww'v'$. $v = ww'v'$ thus by definition of $\partial_B(v)$ again, $ww' \sqsubseteq \partial_B(v)$ thus $w \sqsubseteq \partial_B(v)$.

We prove (13). Since $\partial_B(u) \sqsubseteq u$, then $\partial_B(\partial_B(u)) \sqsubseteq \partial_B(u)$ according to (1) it is enough to prove $\partial_B(u) \sqsubseteq \partial_B(\partial_B(u))$. By definition of ∂_B , $u = \partial_B(u)u'$ with $u' \parallel B$ and $\partial_B(u) = \partial_B(\partial_B(u))u''$ with $u'' \parallel B$. Thus $u = \partial_B(\partial_B(u))u''u'$ with $u''u' \parallel B$ hence by definition of ∂_B , $\partial_B(u) \sqsubseteq \partial_B(\partial_B(u))$.

We prove (14). Since $u \partial_B(v) \sqsubseteq uv$, then $\partial_B(u \partial_B(v)) \sqsubseteq \partial_B(uv)$ and according to (1) it is enough to prove $\partial_B(uv) \sqsubseteq \partial_B(u \partial_B(v))$. By definition of $\partial_B(v)$ there exists $v' \parallel B$ such that $v = \partial_B(v)v'$ then $uv = u \partial_B(v)v'$ thus $\partial_B(uv) \sqsubseteq u \partial_B(v)$. According to (13), $\partial_B(\partial_B(uv)) = \partial_B(uv)$ thus $\partial_B(uv) \sqsubseteq \partial_B(u \partial_B(v))$ which terminates the proof of (14).

We prove (15). Assume ua is prime. The direct implication is immediate using (12). For the converse implication, assume $av \sqsubseteq \partial_B(avw)$. Since $\partial_B(uavw) = \partial_B(uav \partial_B(w))$, without loss of generality we can assume $w = \partial_B(w)$ thus $\partial_B(avw) = avw$, thus we can replace v with vw and assume $w = \epsilon$. Then $\partial_B(av) = av$ thus according to (11) $v = \partial_B(v)$ and $a = \partial_B(a)$. Then uav factorizes as $uav = w_0w_1$ with $w_0 = \partial_B(uav)$ and $w_1 \parallel B$. Then according to (6), $ua = u_0u_1$ such that $w_0 = u_0z_0$ and $w_1 = u_1z_1$. Since ua is a -prime then either $u_1 = \epsilon$ or u_1 is a -prime. In case $u_1 = \epsilon$ then $w_0 = uaz_0$ thus $ua \sqsubseteq \partial_B(uav)$ and according to (11), $\partial_B(uav) = u \partial_B(av) = uav$ so the proof is done. Otherwise, u_1 is a -prime but $w_1 \in B$ thus $a \parallel B$, a contradiction with $a = \partial_B(a)$. This terminates the proof of (15). \square

2.2 Processes and automata

Definition 1. *A Zielonka automata \mathcal{A} on the alphabet A and the set of processes \mathbb{P} is a tuple*

$$\mathcal{A} = (A, (Q_p)_{p \in \mathbb{P}}, (i_p)_{p \in \mathbb{P}}, (F_p)_{p \in \mathbb{P}}, (A_p)_{p \in \mathbb{P}}, \Delta),$$

where

- \mathbb{P} is a finite set called the set of processes,
- Q_p is the set of states of process p ,
- $i_p \in Q_p$ is the initial state of p ,
- $F_p \subseteq Q_p$ is the set of final states of p ,
- A_p is the set of actions of process p ,
- $A = \bigcup_{p \in \mathbb{P}} A_p$ and for $a \in A$, the set $\{p \in \mathbb{P} \mid a \in A_p\}$ is called the domain of a and denoted $\text{dom}(a)$,
- $\mathcal{T} \subseteq \bigcup_{a \in A} \{a\} \times \prod_{p \in \text{dom}(a)} Q_p \times Q_p$ is the set of transitions,

We assume that transitions are deterministic i.e. for every $a \in A$, if $(a, (q_p, q'_p)_{p \in \text{dom}(a)}) \in \Delta$ and $(a, (q_p, q''_p)_{p \in \text{dom}(a)}) \in \Delta$ then $q'_p = q''_p$ for every $p \in \text{dom}(a)$.

For the rest of the paper we fix a Zielonka automaton \mathcal{A} .

The automaton \mathcal{A} defines a reflexive and symmetric dependency relation D on a defined by:

$$((a, b) \in D) \iff (\text{dom}(a) \cap \text{dom}(b) \neq \emptyset) ,$$

and the associate independency relation \mathbb{I}

$$a \mathbb{I} b \iff \text{dom}(a) \cap \text{dom}(b) = \emptyset .$$

This naturally defines a notion of Mazurkiewicz trace on alphabet A .

We extend the notion of views and independence to processes. Let $p \in \mathbb{P}$ then the p -view of a trace $u \in A^*$ is

$$\partial_p(u) = \partial_{A_p}(u) ,$$

and since all letters of A_p are dependent from each other,

$$\forall p \in \mathbb{P}, \forall u \in A^*, \partial_p(u) \text{ is prime.} \quad (17)$$

Moreover for $p \in \mathbb{P}$ and $u \in A^*$,

$$p \mathbb{I} u$$

is a notation for $A_p \mathbb{I} u$. We extend the notion of domain to traces:

$$\text{dom}(a_1 \cdots a_n) = \bigcup_{i=1}^n \text{dom}(a_i) .$$

Definition 2 (Plays and maximal plays). *The set of plays of the automaton \mathcal{A} denoted $\text{plays}(\mathcal{A})$ is defined inductively, together with a mapping $Q : \text{plays}(\mathcal{A}) \rightarrow \prod_{p \in \mathbb{P}} Q_p$. The set $\text{plays}(\mathcal{A}) \subseteq A^*$ is the smallest set of traces on A such that:*

- ϵ is a play and $Q(\epsilon) = (i_p)_{p \in \mathbb{P}}$,
- if $u \in \text{plays}(\mathcal{A})$, $a \in A$ and there exists $(a, (q_p, q'_p)_{p \in \text{dom}(a)}) \in \Delta$ such that $\forall p \in \text{dom}(a)$, $q_p = Q_p(u)$ then $ua \in \text{plays}(\mathcal{A})$ and for every $p \in \mathbb{P}$,

$$Q_p(ua) = \begin{cases} Q_p(u) & \text{if } p \notin \text{dom}(a), \\ q'_p & \text{otherwise.} \end{cases}$$

The definition makes sense because for every $u \in \text{plays}(\mathcal{A})$, whatever linearization of u is chosen to compute $Q(u)$ does not change the value of $Q(u)$, since

$$\forall u \in \text{plays}(\mathcal{A}), Q_p(u) = Q_p(\partial_p(u)) ,$$

which can be easily proved inductively.

2.3 Strategies and games

Given an automaton \mathcal{A} , we would like the processes to choose actions so that the length of plays in \mathcal{A} is bounded and in every maximal play of \mathcal{A} , all states are final.

Not all actions are controllable by processes, and we assume that A is partitioned in $A = Ac \sqcup Ae$ where Ac is the set of controllable actions and Ae the set of environment actions. Intuitively, processes cannot prevent their environment to play actions in Ae , while they can forbid some of the actions that are in Ac .

The choice of actions by processes is not made once for all at the beginning of the play, it is dynamic and at every step, a process p can choose a new set of actions, depending on the information p has on the play.

This information of a process p on a play u is assumed to be the p -view $\partial_p(u)$: intuitively two processes cannot communicate together unless they synchronize on a common action and in this case they exchange as much information about the play as they

want, which allows them to compute a common p -view of the play, common indeed because for every a -prime play $ua \in \text{plays}(\mathcal{A})$,

$$\forall p, q \in \text{dom}(a), \partial_p(ua) = \partial_q(ua) = ua .$$

We adopt a modern terminology and call the automaton \mathcal{A} together with the partition $A = Ac \sqcup Ae$ a *distributed game*, or simply a game in this paper, in which the processes play distributed strategies, defined as follows.

Definition 3 (Distributed strategy). *A distributed strategy σ_p for process $p \in \mathbb{P}$ in the game \mathcal{A} is a mapping $\sigma_p : A^* \rightarrow 2^A$ such that for every $u \in A^*$,*

$$\begin{aligned} Ae &\subseteq \sigma_p(u) , \\ \sigma_p(u) &= \sigma_p(\partial_p(u)) . \end{aligned}$$

A distributed strategy in \mathcal{A} is a tuple $\sigma = (\sigma_p)_{p \in \mathbb{P}}$ where each σ_p is a strategy of process p .

A play $u = a_1 \cdots a_n$ is a σ -play if $u \in \text{plays}(\mathcal{A})$ and for every $i \in 1..n$ and every $p \in \text{dom}(a_i)$, $a_i \in \sigma_p(a_1 \cdots a_i)$. A σ -play is maximal if it is not the strict prefix of another σ -play.

Our goal is to synthesize winning strategies, which ensure that the game terminates and all processes are in final state.

Definition 4 (Winning strategy). *A strategy σ is winning if the set $\{u \mid u \text{ is a } \sigma\text{-play}\}$ is bounded and for every maximal σ -play u ,*

$$Q(u) \in \prod_{p \in \mathbb{P}} F_p .$$

Actually, not all winning strategies are equivalent, we prefer those which have small duration, in the following sense.

Definition 5 (Duration of a strategy). *The duration $\text{dur}(\sigma)$ of a strategy σ is an integer in $\mathbb{N} \cup \{\infty\}$ defined as follows. If the set of σ -plays is infinite then $\text{dur}(\sigma) = \infty$. Otherwise*

$$\text{dur}(\sigma) = \sum_{u \text{ maximal } \sigma\text{-play}} |u| .$$

The distributed synthesis problem asks, given a game $G = (\mathcal{A}, Ac, Ae)$, whether the game is winning, in the sense where there is a winning strategy in G . If yes such a strategy should be computed.

We do not know whether the distributed synthesis problem is decidable in the general case, but we know it is decidable when the game is a broadcast game.

2.4 Broadcast games

The notion of broadcast game relies on the notion of B -broadcast of a prime trace, with $B \subseteq A$.

Intuitively, a B -broadcast is a prime play in a strategy such that the maximal action of the play and the associated information about the play can be broadcasted in priority to a pool of processes using a pool of actions B . Processes of the pool play no action outside B until they observe the maximal action.

Formally, a B -broadcast is a prime play whose last action is in B and such that every parallel play is either a B -thread or is independent of B .

Definition 6 (σ -broadcast). *Let $B \subseteq A$ a subset of actions. We say that a prime play u is a B -broadcast if $u \in A^* B$ and for every play uv such that v is prime,*

$$(uv \text{ is prime}) \vee (v \mathbb{I} B) \vee (v \in B^*) . \quad (18)$$

We say that a prime σ -play u is a B -broadcast in σ if (18) holds for every σ -play uv such that v is prime.

The first clause in the disjunction (18) can be reformulated in several ways.

Proposition 1. *Let $B \subseteq A$ and $a, b \in A$ and $u, v \in A^*$ such that u is a -prime and v is b -prime. Then the following conditions are equivalent:*

$$\begin{aligned} (uv \text{ is } b\text{-prime}) &\iff \neg(a \parallel v) \\ &\iff a \sqsubseteq \partial_b(av) \\ &\iff u \sqsubseteq \partial_b(uv) . \end{aligned}$$

Proof. We prove one by one all implications from the bottom to the top and finally the implication from the very top to the very bottom. Assume $(u \sqsubseteq \partial_b(uv))$ then according to (15) $(a \sqsubseteq \partial_b(av))$. Assume $a \parallel v$ then in particular $a \parallel b$ and $\partial_b(av) = \partial_b(v) = v$ because v is b -prime thus $(a \not\sqsubseteq \partial_b(av))$. Assume $\neg(a \parallel v)$. Then av is b -prime according to (10) thus uv is b -prime according to (9). Assume uv is b -prime then $uv = \partial_b(uv)$ thus $u \sqsubseteq \partial_b(uv)$. \square

The proof of our main theorem is by induction on the set of actions, and relies on inductive decompositions of the alphabet A .

Definition 7 (Inductive decomposition of A). *Let A be an alphabet equipped with a reflexive and symmetric dependency relation $D \subseteq A \times A$. For every $B \subseteq A$ we denote $G_B = (B, D \cap B \times B)$ the induced graph of $G_A = (A, D)$. An inductive decomposition of A is a collection $\mathcal{C} \subseteq 2^A \times 2^A$ such that for every connected component A' of G_A , $\exists C', (A', C') \in \mathcal{C}$ and $\forall (B, C) \in \mathcal{C}$,*

$$C \subseteq B \quad (19)$$

$$G_C \text{ is a clique} \quad (20)$$

$$\forall \text{ connected component } B' \text{ of } G_{B \setminus C}, \exists C', (B', C') \in \mathcal{C} \quad (21)$$

We can now introduce our decidable class of games, broadcast games. The notion is defined so that decidability results of (Gastin et al. 2004; Madhusudan et al. 2005; Genest et al. 2013) can be retrieved quite easily, as will be done in the next section. However to our opinion the main contribution of the paper are the notions of useless threads and shortcuts introduced in section 4.

Definition 8 (Broadcast games). *Let $N \in \mathbb{N}$ and \mathcal{C} an inductive decomposition of A . A game G is a (N, \mathcal{C}) -broadcast game if for every $(B, C) \in \mathcal{C}$ and every $u, v_1, \dots, v_N \in A^*$ such that*

$$uv_1v_2 \dots v_N \text{ is a play} \quad (22)$$

$$\forall i \in 1..N, v_i \in B^*C, \quad (23)$$

$$A(v_1) = A(v_2) = \dots = A(v_N) \quad (24)$$

$$\forall i \in 1..N, uv_1v_2 \dots v_i \text{ is a prime play} , \quad (25)$$

there exists $(B', C') \in \mathcal{C}$ and $j \in 1..N$ such that $B' \subseteq B$ and $uv_1v_2 \dots v_j$ is a B' -broadcast.

A game G is a N -broadcast game if there exists an inductive decomposition \mathcal{C} of A such that G is a (N, \mathcal{C}) -broadcast game.

A game G is a broadcast game if there exists $N \in \mathbb{N}$ such that G is a N -broadcast game.

Being a broadcast game is a decidable property.

Proposition 2. *It is decidable whether a game G is a broadcast game. In case G is a broadcast-game then there exists $N \leq \prod_{p \in \mathbb{P}} |Q_p|$ such that G is a N -broadcast game.*

Proof. Let $M = \prod_{p \in \mathbb{P}} |Q_p|$. Let \mathcal{C} be an inductive decomposition of A and $(B, C) \in \mathcal{C}$. A standard argument of automata theory shows that the conditions in Definition 8 are satisfied for some N and only if they are satisfied when $N = M$ and the trace u as well as each trace v_i has length less than M .

Let $I_b = \{a \in A \mid a \parallel b\}$ and $I_B = \{a \in A \mid a \parallel C\}$. Then we prove that a b -primary play u is a B -broadcast if and only if there

does not exist a prime play $v \in A^*$ such that uv is a play and

$$v \notin I_b^* \quad (26)$$

$$v \notin B^* \quad (27)$$

$$v \notin I_B^* . \quad (28)$$

The conjunction of these three conditions is indeed equivalent to the opposite of (18), since according to Proposition 1,

$$(uv \text{ is prime}) \iff \neg(v \parallel b) \iff (v \notin I_b^*).$$

Again, a standard pumping argument shows that if there exists $v \in A^*$ which satisfies (60), (61) and (62) and such that uv is a play then v can be chosen of length at most $3M$.

Thus, the proposition holds since whether G is a broadcast game or not can be decided by enumerating all $N \leq M$ and every inductive decomposition \mathcal{C} of A and for each of those, enumerating all u, v_0, \dots, v_N of length less than M , and for those which satisfies the conditions in Definition 8, enumerating all $i \in 1..N$ and words w of length less than $3M$ such that $uv_0 \dots v_i w$ is a play and check whether (60), (61) and (62) are satisfied with u replaced by $uv_0 \dots v_i$ and v replaced by w . If a witness is found then G is not a broadcast game, otherwise G is a broadcast game. \square

2.5 Main result

We can now state our main result.

Theorem 1. *It is decidable whether a distributed game is a winning broadcast game.*

The algorithm consists in enumerating all possible strategies whose plays have length less than some computable bound, and check whether any of these strategies is winning. This bound is defined by equation (??) in Section 4, where the proof of the theorem can be found as well. The bound is quite large, which is not surprising since the problem is non-elementary (Genest et al. 2013). Before giving the proof of the main theorem, we provide some examples and applications in the next section.

3. Examples of N -broadcast games

In this section we provide several examples of N -broadcast games, and according to Theorem 1, each of them is an example of a decidable class of systems for the distributed synthesis problem.

The first two classes are connectedly communicating games and series-parallel games whose decidability was already known. The present paper provides an alternate proof to these results. The third example are acyclic games whose decidability was already known in the case where all actions are local or binary. The fourth example is the class of triangulated games and the special case of three-player games.

3.1 Connectedly communicating games are broadcast games

Connectedly communicating games have been introduced (Madhusudan et al. 2005) under the name of *connectedly communicating processes*, and the authors did establish the decidability of the MSO theory of the corresponding event structure, which implies that controller synthesis is decidable.

A game is connectedly communicating if it is k -communicating for some $k \in \mathbb{N}$, which holds if in every play, if a process q never plays while process p plays k times, then p and q will stay forever in separate threads. Formally, for a play u and a process p we denote $|u|_p$ the number of letters of u whose domain contains p . Then a game is k -communicating for some $k \in \mathbb{N}$ if for every processes $p, q \in \mathbb{P}$ and play uvw in G ,

$$((|v|_p \geq k) \wedge (|v|_q = 0) \wedge w \text{ prime}) \implies (|w|_p = 0 \vee |w|_q = 0) .$$

Every k -communicating game is a k -broadcast game. For every set of processes $Q \in \mathbb{P}$ we denote

$$A_Q = \{a \in A \mid \text{dom}(a) \subseteq Q\} .$$

Let

$$\mathcal{C} = \{(A_Q, A_q \cap A_Q) \mid Q \subseteq \mathbb{P}, q \in Q\} .$$

Then \mathcal{C} is clearly an inductive decomposition of A .

We show that G is a (k, \mathcal{C}) -broadcast game. Let $u \in A^*$ and $v_1, \dots, v_k \in A_Q^*$ such that $uv_1 \dots v_k$ is a play, $A(v_1) = A(v_2) = \dots = A(v_k) \forall i \in 1..k, v_i \in B^*$ and $uv_1 \dots v_k$ is prime for every i . Let $Q' = \cup_{a \in A(v_1)} \text{dom}(a)$. Then we show that $uv_1 \dots v_k$ is a $A_{Q'}$ -broadcast. Since $A(v_i) = A(v_1)$ for each i , then $\forall p \in Q', |v_1 \dots v_k|_p \geq k$ and $\forall p \in \mathbb{P} \setminus Q', |v_1 \dots v_k|_p = 0$. Thus by definition of k -communicating game, for every prime play w such that $uv_1 \dots v_k w$ is a play,

$$(\forall p \in Q', |w|_p = 0) \vee (\forall p \in \mathbb{P} \setminus Q', |w|_p = 0) .$$

In the case where $\forall p \in Q', |w|_p = 0$ then $\forall p \in Q', p \mathbb{I} w$ thus $w \mathbb{I} A_{Q'}$. In the case where $\forall p \in \mathbb{P} \setminus Q', |w|_p = 0$, then $w \in A_{Q'}^*$. This shows that $uv_1 \dots v_k$ is an $A_{Q'}$ broadcast.

Finally, every connectedly communicating game is a broadcast game.

3.2 Series-parallel games are 1-broadcast games

A series-parallel game is a game such that the dependence graph (A, D) of the alphabet A is a co-graph i.e. it belongs to the smallest class of graphs containing singletons and closed under parallel product and complementation. Series-parallel games were proved decidable in (Gastin et al. 2004), for a different setup than ours: in the present paper we focus on process-based control while (Gastin et al. 2004) was focusing on action-based synthesis. Actually action-based control is more general than process-based control, see (Muscholl et al. 2008) for a proof.

Any series-parallel game is a 1-broadcast game. We define inductively for every co-graph (A, D) an inductive decomposition $\mathcal{C}(A, D)$ as follows. If A is a singleton then $\mathcal{C}(A, D) = \{(A, A)\}$. If A is not a singleton then A can be partitioned into A_0 and A_1 such that, denoting $D_0 = D \cap A_0 \times A_0$ and $D_1 = D \cap A_1 \times A_1$ both induced subgraphs (A_0, D_0) and (A_1, D_1) are co-graphs and

$$\text{either } D = D_0 \cup D_1 \quad (29)$$

$$\text{or } D = D_0 \cup D_1 \cup A_0 \times A_1 . \quad (30)$$

Let \mathcal{C}_0 and \mathcal{C}_1 the two inductive decompositions associated with (A_0, D_0) and (A_1, D_1) respectively. The definition of $\mathcal{C}(A, D)$ depends whether we are in case (29) or (30). In case (29) holds we define

$$\mathcal{C}(A, D) = \mathcal{C}_0 \cup \mathcal{C}_1 ,$$

and in case (30) holds we define

$$\begin{aligned} \mathcal{C}(A, D) = \mathcal{C}_0 \cup \mathcal{C}_1 \\ \cup \{(A_1 \cup B_1, \{b\}) \mid B_1 \subseteq A_1 \wedge b \in B_1\} . \end{aligned}$$

Then $\mathcal{C}(A, D)$ is an inductive decomposition. In case of a parallel product (29), this comes from the fact that every subset of A connected in G_A is either included in A_0 or in A_1 thus since \mathcal{C}_0 and \mathcal{C}_1 are inductive decomposition, \mathcal{C} also is. In case of a serial product (30), then the only connected component of G_A is A and again the properties are inherited inductively.

If the alphabet of a game G is a co-graph (A, D) then G is a $(1, \mathcal{C}(A, D))$ -broadcast game because the following property

holds, independently of which game G is played:

$$\begin{aligned} \forall (B, \{b\}) \in \mathcal{C}(A, D), \exists (B', \{b'\}) \in \mathcal{C}(A, D) \\ \text{such that } B' \subseteq B \text{ and } \forall uvw \in A^*, \\ ((v \in B'^* b' \wedge uv \text{ is prime} \wedge w \text{ is prime}) \\ \implies (v \mathbb{I} w \vee w \in (B')^* \vee uvw \text{ is prime})) . \quad (31) \end{aligned}$$

We prove (31) by induction on the co-graph A . Assume inductively that the property holds when $A = A_0$ or $A = A_1$.

In case of a parallel product (29), every prime trace is either in A_0^* or A_1^* . Without loss of generality assume that $uv \in A_0^*$. Then either $w \in A_1^*$ or $w \in A_0^*$. In the first case $w \mathbb{I} v$. In the second case $uvw \in A_0^*$ thus (29) holds according to the inductive hypothesis.

Assume now we are in the case of a serial product (30) and let $(B, \{b\}) \in \mathcal{C}$. If $(B, \{b\}) \in \mathcal{C}_0 \cup \mathcal{C}_1$, we choose $(B', \{b'\}) = (B, \{b\})$. Otherwise we choose $(B', \{b'\}) = (A'_0, \{b'\})$ where A'_0 is a connected component of G_{A_0} and $b' \in A'_0$.

We distinguish between three cases whether $(A(w) \cap A_0 \neq \emptyset \wedge A(w) \cap A_1 \neq \emptyset)$ or $w \in A_0^*$ or $w \in A_1^*$. In the first case since either $b' \in A_0$ or $b' \in A_1$ then according to (30) $\neg(w \mathbb{I} a)$ thus according to (10), uvw is prime. In the second case, $w \in A_0^*$. If $b \in A_1$ then $\neg(b' \mathbb{I} w)$ thus uvw is prime according to (10). If $b' \in A_0$ then by definition of \mathcal{C} $(B, \{b'\}) \in \mathcal{C}_0$ and $B \subseteq A_0$ thus $v \in A_0^*$. Since $vw \in A_0^*$ we can apply inductively (31) for $u = \epsilon$, then $(v \mathbb{I} w \vee w \in B'^* \vee vw \text{ is prime})$. In case vw is prime then uvw is as well because uv is prime thus (31) holds. In the third case, $w \in A_1^*$. If $b' \in A_0$ then $\neg(b' \mathbb{I} w)$ thus uvw is prime according to (10). If $b' \in A_1$ then necessarily by choice of $(B', \{b'\})$, $(B', \{b'\}) \in \mathcal{C}_1$, thus $v \in A_1^*$. Since $vw \in A_1$ we can apply inductively (31) for $u = \epsilon$ and $(v \mathbb{I} w \vee w \in B'^* \vee vw \text{ is prime})$. Again, in case vw is prime then uvw is as well because uv is prime thus (31) holds.

Finally, according to (31), every series-parallel game is a 1-broadcast game.

3.3 Acyclic games are 1-broadcast games

Intuitively, an acyclic game as defined in (Genest et al. 2013) is a game where processes are arranged as a tree and actions are either local or synchronize a father and his son. Formally, the processes are arranged as a tree $T_{\mathbb{P}} = (\mathbb{P}, E_{\mathbb{P}})$, and each action is either a local action whose domain is a singleton or a binary synchronizing action such that $\text{dom}(a) = \{p, q\}$ and $(p, q) \in E_{\mathbb{P}}$ i.e. q is the father of p in the process tree.

We extend the definition of (Genest et al. 2013) to the case of non-binary actions, and we assume that:

$$\forall a \in A, \text{dom}(a) \text{ is connected in } T_{\mathbb{P}}, \quad (32)$$

in other words if an action synchronizes two processes p_1, p_2 it synchronizes as well all processes on the shortest path from p_1 to p_2 in $T_{\mathbb{P}}$.

As a consequence, in an acyclic game for every process p , information from a descendant of p to an ancestor of p has to flow through p . Formally for every prime play $u = a_1 \dots a_n$ and process p , if $\text{dom}(u) = \bigcup_{i=1}^n \text{dom}(a_i)$ contains both an ancestor and a descendant of p , it contains p as well.

We associate with every acyclic game, an inductive decomposition \mathcal{C} such that G is a $(1, \mathcal{C})$ -broadcast game. for a process p , we denote T_p the subtree of $T_{\mathbb{P}}$ rooted at p and we denote $B_p = \cup_{q \in T_p} A_q$. Then

$$\mathcal{C} = \{(B_p, A_p \mid p \in \mathbb{P}\} .$$

Then \mathcal{C} is obviously an inductive decomposition.

We show that any game G with alphabet A is a $(1, \mathcal{C})$ -broadcast game. Let $(B_p, A_p) \in \mathcal{C}$ and $u, v \in A^*$ such that $uv \in A_p^* B_p$ and uv is prime. Let $b \in A_p$ be the last letter of v such that v is

b -prime. Then for every prime $w \in A^*$,

$$w \parallel v \vee w \in B_p^* \vee uvw \text{ is prime.} \quad (33)$$

We distinguish between three cases, depending on the set $\text{dom}(w)$ of processes that are involved in w . First assume that $p \in \text{dom}(w)$. Then since $b \in A_p$, then $\neg(b \parallel w)$ thus according to Proposition 1 uvw is prime and (33) holds. Now assume $p \notin \text{dom}(w)$. Since w is prime, $\text{dom}(w)$ is a connected set of nodes of $T_{\mathbb{P}}$. Thus in the case where $p \notin \text{dom}(w)$ then either all processes in $\text{dom}(w)$ or no process in $\text{dom}(w)$ belong to the subtree T_p . In the first case, $w \in B_p^*$ by definition of B_p . In the second case, $w \parallel B_p$. Finally, (33) holds in all cases.

Consequently, all acyclic games with arbitrary actions are 1-broadcast games.

3.4 Three player games are 1-broadcast games

Any 3-player game G with processes $\{1, 2, 3\}$ is a 1-broadcast game. We order the actions A with a total order \preceq such that for every $a, b \in A$, $a \prec b$ whenever $|\text{dom}(a)| < |\text{dom}(b)|$ or $|\text{dom}(a)| = |\text{dom}(b)|$ and $\max \text{dom}(a) < \max \text{dom}(b)$. Then we define

$$\mathcal{C} = \{(B, \{b\}) \mid (B \subseteq A) \wedge (b = \max_{\preceq} B)\},$$

which is clearly an inductive decomposition.

We show that any game G with alphabet A is a $(1, \mathcal{C})$ -broadcast game. Let $(B, \{b\}) \in \mathcal{C}$ and $uv \in A^*$ such that uv is prime and $uv \in B^*b$. Let $w \in A^*$ such that w is prime. We show

$$(w \parallel B \vee w \in B^* \vee uvw \text{ is prime}). \quad (34)$$

According to Proposition 1, if uvw is not prime then $b \parallel w$, which we assume now. If $|\text{dom}(b)| = 3$ then $b \parallel w$ implies $w = \epsilon$ thus $w \in B^*$. If $|\text{dom}(b)| = 2$ then $b \parallel w$ implies $|\text{dom}(w)| = 1$, thus all letters of w are \preceq -inferior to b hence $w \in B^*$ and (33) holds. If $|\text{dom}(b)| = 1$ then b is a local action of a process p and since $b = \max_{\preceq} B$ then by choice of \preceq , $\forall c \in B$, $|\text{dom}(c)| \leq 1$ i.e. there are only local actions in B . Since $A(v) = B$ and v is prime then $B = \{b\}$ because two local actions are either equal or independent. Then, since $w \parallel b$ then $w \parallel C$ thus (34) holds.

Consequently, all 3-player games are 1-broadcast games.

3.5 Triangulated games are 1-broadcast games

A triangulated game is a game where processes are arranged as an undirected graph $G_{\mathbb{P}} = (\mathbb{P}, E_{\mathbb{P}})$ such that all simple cycles in the graph have length 3, and moreover we assume that

$$\forall a \in A, \text{dom}(a) \text{ is connected in } G_{\mathbb{P}}. \quad (35)$$

This definition is inspired by (Diekert and Muscholl 1996).

We build by induction an inductive decomposition $\mathcal{C}(\mathbb{P})$ of A such that $G_{\mathbb{P}}$ is a $(1, \mathcal{C}(\mathbb{P}))$ -broadcast game. Since all simple cycles have length 3 in $G_{\mathbb{P}}$, then either $G_{\mathbb{P}}$ is disconnected or $G_{\mathbb{P}}$ has three vertices or there exists a vertex $p \in \mathbb{P}$ whose removal disconnects $G_{\mathbb{P}}$. If $G_{\mathbb{P}}$ is disconnected then let $\mathbb{P}_1, \dots, \mathbb{P}_j$ be the connected components and let

$$\mathcal{C}(\mathbb{P}) = \bigcup_{i=1}^j \mathcal{C}(\mathbb{P}_i).$$

If $G_{\mathbb{P}}$ has three vertices then let $\mathcal{C}(\mathbb{P})$ be defined like in the previous subsection. Finally, if for some process $p \in \mathbb{P}$ the removal of p from $G_{\mathbb{P}}$ disconnects $G_{\mathbb{P}}$ into several components $\mathbb{P}_1, \dots, \mathbb{P}_j$. Then we define:

$$\mathcal{C}(\mathbb{P}) = \left\{ \left(A_p \cup \bigcup_{q \in \mathbb{P}_1} A_q, A_p \right) \right\} \cup \bigcup_{i=2}^j \mathcal{C}(\mathbb{P}_i).$$

If one of the \mathbb{P}_i is a three player game then we give number 3 to the vertex connecting \mathbb{P}_i to p and $\mathcal{C}(\mathbb{P}_i)$ is build using 3 has the leading vertex.

Then \mathcal{C} is obviously an inductive decomposition.

Moreover, it has the following extra property: for every connected subset $X \subseteq \mathbb{P}$ such that X is a connected subset of $G_{\mathbb{P}}$, and for every $(B, A_p) \in \mathcal{C}$,

$$(X \cap \text{dom}(B) \neq \emptyset) \implies (p \in X \vee X \subseteq \text{dom}(B)). \quad (36)$$

We show that any game G with alphabet A is a $(1, \mathcal{C})$ -broadcast game. In the case where $G_{\mathbb{P}}$ is disconnected, this is equivalent to showing that for each $i \in 1..j$, the game G_i restricted to \mathbb{P}_i is a $(1, \mathcal{C})$ -broadcast game, which is obvious.

In the case where $G_{\mathbb{P}}$ has three processes, this was proved in the previous subsection.

Now assume $G_{\mathbb{P}}$ can be disconnected by removing process p . Let $(B, C) \in \mathcal{C}$ and $u, v \in A^*$ such that $uv \in B^*A_q$ and uv is prime. Let $b \in A_q$ be the last letter of v such that v is b -prime. Then for every prime $w \in A^*$, we show that

$$w \parallel v \vee w \in B_p^* \vee uvw \text{ is prime.} \quad (37)$$

Assume first that $(B, C) \in \mathcal{C}_{\mathbb{P}_i}$ for some set \mathbb{P}_i of three processes. Then since w is prime, $\text{dom}(w)$ is connected, thus either $\text{dom}(w)$ does contains the leader of \mathbb{P}_i or $w \parallel \mathbb{P}_i$. Moreover the projection w' of w on $\cup_{q \in \mathbb{P}_i} A_q$ has property (34). So either B contains the leader of \mathbb{P}_i and then uvw is prime, or B does not and $w' \parallel B$ thus $w \parallel B$. This implies (37).

Otherwise (B, C) is of the form $(B, A_q) \in \mathcal{C}$. Let $u, v \in A^*$ such that $uv \in B^*A_q$ and uv is prime. Let $b \in A_q$ be the last letter of v such that v is b -prime.

First assume that $q = p$. If $p \in \text{dom}(w)$ then $\neg(b \parallel w)$ thus according to Proposition 1 uvw is prime and (37) holds. Otherwise assume $p \notin \text{dom}(w)$. Since w is prime, $\text{dom}(w)$ is a connected set of nodes of $G_{\mathbb{P}}$ thus $\text{dom}(w)$ is included in a connected component of $G_{\mathbb{P}} - p$ and there exists $i \in 1..j$ such that $\text{dom}(w) \subseteq \mathbb{P}_i$. If $i = 1$ then $w \in B^*$ thus (37) holds. If $i \neq 1$ then $w \parallel B$ thus (37) holds as well.

Now, assume $q \neq p$. Then by definition of \mathcal{C} , $\text{dom}(v) \subseteq \mathbb{P}_i$ for some $i \in 2..j$ and $B = A_q \cup \cup_{r \in \mathbb{P}_i} A_r$. Since w is prime, $X = \text{dom}(w)$ is a connected subset of $G_{\mathbb{P}}$ thus we can apply property (36) hence either $q \in \text{dom}(w)$ in which case w has a letter in A_q and uvw is prime or $\text{dom}(w) \subseteq \mathbb{P}_i$ in which case $w \in B^*$ or $\text{dom}(w) \cap \mathbb{P}_i = \emptyset$ in which case $w \parallel v$. In all three cases (37) holds.

Consequently, all triangulated games with arbitrary actions are 1-broadcast games.

3.6 Other examples

Other examples that are hybrid between triangulated games and connectedly communicating games can be designed.

4. Proof of the main theorem

The proof is easy to sketch, but harder to implement because distributed systems are not so easy to handle. For every subset of actions $C \subseteq A$, we compute inductively a bound K_C such that any winning strategy which has a C -thread of duration more than K_C can be simplified in a shorter winning strategy. The simplification consists in removing a *useless thread* from the strategy. This operation is called a *shortcut*, and has to be carefully done so that the new object obtained after removal of the shortcut is still a winning distributed strategy.

Taking a shortcut in a strategy consists in playing the strategy until a particular play x happens, and then jump to a continuation xy of this play, and keep playing σ as if y had actually happened after x , although it did not.

Definition 9 (Shortcut). Let $x, y \in A^*$ such that xy is a σ -play. Let $\phi : A^* \rightarrow A^*$ be the mapping:

$$\phi_{x,y}(u) = \begin{cases} u & \text{if } x \not\sqsubseteq u \\ xyv & \text{if } u = xv \end{cases}, \quad (38)$$

Then the (x, y, σ) -shortcut is the mapping $\sigma_{x,y} : A^* \rightarrow A^*$ defined by:

$$\sigma_{x,y} = \sigma \circ \phi_{x,y}.$$

The mapping $\phi_{x,y}$ is well-defined since according to (2), there is a unique v such that $u = xv$.

There is a priori no reason in general for $\sigma_{x,y}$ to be a distributed strategy. For that we need extra conditions on the pair (x, y) and we introduce the notion of useless thread.

Definition 10 (Threads). Let $B \subseteq A$ a non-empty set of letters, a B -thread is a pair $(u, v) \in A^* \times A^*$ such that uv is a play and $v \in B^*$. In case uv is a σ -play, (u, v) is called a B -thread of strategy σ .

Some threads are called useless threads, we will see later that they can be removed from the strategy.

Definition 11 (Useless thread). Let σ be a strategy. A useless thread in σ is a B -thread (x, y) such that there exists $b \in B$ with the following properties:

$$x \text{ and } xy \text{ are } b\text{-prime}, \quad (39)$$

$$x \text{ and } xy \text{ are } B\text{-broadcasts in } \sigma, \quad (40)$$

$$\text{every process } p \in \mathbb{P} \text{ has the same state in } x \text{ and } xy, \quad (41)$$

$$\forall v \in B^*, (v \parallel b) \implies (\sigma(xv) = \sigma(xyv)). \quad (42)$$

Taking a shortcut of a useless thread in a distributed strategy makes sense because the result is still a distributed strategy.

Lemma 2. Let (x, y) a useless thread in a distributed strategy σ . Then the (x, y, σ) -shortcut $\sigma_{x,y}$ is a distributed strategy.

Proof. We denote $\tau = \sigma_{x,y} = \sigma \circ \phi_{x,y}$ the (x, y, σ) -shortcut. To prove that τ is a distributed strategy, we take any process $p \in \mathbb{P}$ and $u \in A^*$ and prove that

$$\tau_p(u) = \tau_p(\partial_p(u)).$$

By definition of v and the shortcut τ , $\tau_p(u) = \tau_p(xv) = \sigma_p(xyv)$ and since σ is a distributed strategy $\sigma_p(xyv) = \sigma_p(\partial_p(xyv))$, thus it is enough to prove:

$$\sigma_p(\partial_p(xyv)) = \tau_p(\partial_p(xv)). \quad (43)$$

We distinguish between three cases.

First case: assume $(x \not\sqsubseteq u \wedge x \not\sqsubseteq \partial_p(u))$. Then $\tau_p(u) = \sigma_p(u) = \sigma_p(\partial_p(u)) = \tau_p(\partial_p(u))$, where the first and third equality hold by definition of a shortcut, and the second equality holds because σ is a distributed strategy. Thus (43) holds in the first case.

Second case: assume $x \sqsubseteq \partial_p(u)$. Since $\partial_p(u) \sqsubseteq u$ this implies $x \sqsubseteq u$, hence there exists $w \in A^*$ such that $u = xw$. We start with proving

$$\partial_p(xyv) = xy \partial_p(v). \quad (44)$$

Since $x \sqsubseteq \partial_p(xw)$, (11) implies

$$\partial_p(xw) = x \partial_p(w). \quad (45)$$

Since (x, y) is a useless thread, according to (39), both x and xy are b -prime. Since moreover $\partial_p(xw) = x \partial_p(w)$ we can apply (15) twice and get first $\partial_p(bw) = b \partial_p(w)$ and then (44). Now that (44)

is proved we can conclude the second case:

$$\sigma_p(\partial_p(xyv)) = \sigma_p(xy \partial_p(v)) \quad (46)$$

$$= \tau_p(x \partial_p(v)) \quad (47)$$

$$= \tau_p(\partial_p(xw)), \quad (48)$$

where (46) comes from (44), (47) hold by definition of shortcuts and τ , and (48) comes from (45). Thus (43) holds in the second case.

We are now left with the third and last case:

$$x \not\sqsubseteq \partial_p(u) \wedge x \sqsubseteq u, \quad (49)$$

which we assume until the end of the proof. Then $u = xv$ for some $v \in A^*$.

We first take care of the special case where $\partial_p(v) = \epsilon$. Then $v \parallel p$ according to (4) thus $\partial_p(xyv) = \partial_p(xy)$. Hence,

$$\sigma_p(\partial_p(xyv)) = \sigma_p(\partial_p(xy)) \quad (50)$$

$$= \sigma_p(xy) \quad (51)$$

$$= \tau_p(x) \quad (52)$$

$$= \sigma_p(x) \quad (53)$$

$$= \sigma_p(\partial_p(x)) \quad (54)$$

$$= \sigma_p(\partial_p(xv)) \quad (55)$$

$$= \tau_p(\partial_p(xv)), \quad (56)$$

where (50) and (55) hold because $v \parallel p$, (51) and (54) hold because σ_p is a distributed strategy, (52) holds by definition of τ , (53) holds because (x, y) is a useless thread and according to (42), (56) holds by definition of τ and because by hypothesis $x \not\sqsubseteq \partial_p(u)$. This shows that (43) holds when $\partial_p(v) = \epsilon$.

Now assume that $\partial_p(v) \neq \epsilon$ (and we keep assuming (49) as well). Since (x, y) is a useless thread in σ , then according to (40), x is a B -broadcast in σ . We can apply the definition of a B -broadcast to $x \partial_p(v)$ because $\partial_p(v)$ is prime according to (17), and $x \partial_p(v)$ is a σ -play because it is a prefix of the σ -play xv . Thus by definition of broadcasts and Proposition 1, one of the three following properties holds:

$$x \sqsubseteq \partial_p(x \partial_p(v)) \quad (57)$$

$$\text{or } \partial_p(v) \in B^* \quad (58)$$

$$\text{or } \partial_p(v) \parallel B. \quad (59)$$

Since $x \not\sqsubseteq \partial_p(x \partial_p(v))$ by hypothesis, (57) is not possible and we are left with the two other cases (58) and (59).

We assume first that (59) holds. Since $\partial_p(v) \neq \epsilon$, it implies that $p \parallel B$. Since (x, y) is a B -thread then $y \in B^*$ thus $p \parallel y$. We can conclude the proof of (43) in the case where (59) holds:

$$\sigma_p(\partial_p(xyv)) = \sigma_p(\partial_p(xy \partial_p(v))) \quad (60)$$

$$= \sigma_p(\partial_p(x \partial_p(v) y)) \quad (61)$$

$$= \sigma_p(\partial_p(x \partial_p(v))) \quad (62)$$

$$= \sigma_p(\partial_p(xv)) \quad (63)$$

$$= \tau_p(\partial_p(xv)), \quad (64)$$

where (60) and (63) hold according to (14), (61) holds because $\partial_p(v) \parallel B$ and $y \in B^*$, (62) holds because $p \parallel y$, and (64) holds by definition of τ_p , since by hypothesis $x \not\sqsubseteq \partial_p(u)$ and $u = xv$. This proves (43) in the case where (59) holds.

Now we are left with the case where (58) holds, i.e. $\partial_p(v) \in B^*$ (and we keep assuming $\partial_p(v) \neq \epsilon$ and (49) as well). We first establish

$$\partial_p(v) \in (B \setminus \{b\})^*. \quad (65)$$

Since by hypothesis $x \not\sqsubseteq \partial_p(xv)$ and x is b -prime then according to (15), $b \not\sqsubseteq \partial_p(bv)$ thus according to (10), $b \parallel \partial_p(v)$ which

implies (65). Since (x, y) is a useless thread in σ , we can apply (42) to $\partial_p(v)$ hence

$$\sigma_p(x \partial_p(v)) = \sigma_p(xy \partial_p(v)) . \quad (66)$$

Finally,

$$\sigma_p(\partial_p(xyv)) = \sigma_p(\partial_p(xy \partial_p(v))) \quad (67)$$

$$= \sigma_p(xy \partial_p(v)) \quad (68)$$

$$= \sigma_p(x \partial_p(v)) \quad (69)$$

$$= \sigma_p(\partial_p(x \partial_p(v))) \quad (70)$$

$$= \sigma_p(\partial_p(xv)) \quad (71)$$

$$= \tau_p(\partial_p(xv)), \quad (72)$$

where equalities (67) and (71) hold according to (14), equalities (68) and (70) hold because σ is a distributed strategy, (69) comes from (66), and finally (72) is by definition of τ and because by hypothesis $x \not\sqsubseteq \partial_p(xv)$. This terminates the proof of (43) in the last case.

As a consequence, τ is a distributed strategy. \square

Taking shortcuts of useless threads is really useful for making winning strategies smaller: it transforms a *winning* distributed strategy into another, *shorter*, *winning* distributed strategy.

Lemma 3. *Let (x, y) a useless thread in a winning distributed strategy σ . Then the (x, y, σ) -shortcut $\sigma_{x,y}$ is a winning distributed strategy as well, and*

$$\text{dur}(\sigma_{x,y}) < \text{dur}(\sigma) . \quad (73)$$

Moreover for every $v \in A^*$,

$$xv \text{ is a } \sigma_{x,y}\text{-play} \iff xyv \text{ is a } \sigma\text{-play}. \quad (74)$$

Proof. We denote $\tau = \sigma_{x,y} = \sigma \circ \phi_{x,y}$ the (x, y, σ) -shortcut.

We first prove property (74). Let $xv \in A^*$ be a τ -play, we prove that xyv is a σ -play by induction on v . When $v = \epsilon$ then xy is a σ -play because by hypothesis (x, y) is a thread. For the inductive step, assume xyv is a σ -play, let $c \in A$ such that xvc is a τ -play, and let us prove that $xyvc$ is a σ -play. Since xvc is a τ -play, $c \in \tau_p(xv)$ for every $p \in \text{dom}(c)$. Thus by definition of τ , $c \in \sigma_p(xyv)$ for every $p \in \text{dom}(c)$ hence $xyvc$ is a σ -play by definition of σ -plays.

Now we prove that τ is winning. Since σ is winning, the set of σ -plays is finite. According to property (74) and the definition of τ , every τ -play is either a σ -play or is a subword of a σ -play thus K is also an upper bound on the length of τ -plays, hence every maximal τ -play is finite. Let u be a maximal τ -play. If $x \not\sqsubseteq u$ then u is a maximal σ -play and since σ is winning u is a winning play. Assume now that $x \sqsubseteq u$ and $u = xw$. According to (74), since xw is a maximal τ -play, xyw is a maximal σ -play, and since σ is winning, all processes are in a final state xyw . Since (x, y) is a useless shell, (41) states that all processes are in the same state in x and xy , and since transitions are deterministic, all processes are in the same state in xw and xyw . So finally all processes are in a final state in xw . Thus τ is winning.

Now we prove property (73). According to (74), the mapping $\phi_{x,y}$ used to define $\tau = \sigma_{x,y}$ in (38) maps maximal τ -plays to maximal σ -plays. Moreover, according to (2), Φ is an injection, and by definition it preserves the length on $\{u \mid x \not\sqsubseteq u\}$ and increases the length of $|y|$ on $\{u \mid x \sqsubseteq u\}$. This shows that $\text{len}(\sigma) \geq \text{len}(\tau) + q \cdot |y|$ where q is the number of maximal τ -plays prefixed by x . Since x is a τ -play, $q \geq 1$. According to (39), $y \neq \epsilon$ thus we get property (73).

This terminates the proof of Lemma 3. \square

If we apply this lemma in the case where (x, y) is a B -thread with $B = A$, we get the following corollary.

Corollary 1. *Let σ be a strategy and xy be a σ -play such that both x and y are b -prime for some letter $b \in A$. Assume that:*

- every process $p \in \mathbb{P}$ has the same state in x and xy ,
- $\forall v \in A^*, (v \parallel b) \implies (\sigma(xv) = \sigma(xyv))$.

Then the mapping τ defined by:

$$\tau(u) = \begin{cases} \sigma(u) & \text{if } x \not\sqsubseteq u \\ \sigma(xyv) & \text{if } u = xv . \end{cases}$$

is a distributed strategy. Moreover $\text{dur}(\tau) < \text{dur}(\sigma)$ and if σ is winning then τ is winning as well.

There is a limit on the length of threads of a strategy, such that above this limits useless threads start appearing in the strategy. An upper bound on this limit is K_A , which is computed as follows.

With every $B \subseteq A$ we associate a constant $K_B \in \mathbb{N}$ as follows. According to Ramsey theorem, for every $m, n \in \mathbb{N}$, there exists a constant $R(m, n)$ such that every undirected complete graph with at least $R(m, n)$ vertices whose edges are labelled with m different colors contains a monochromatic clique of size n . Then we define inductively $K_\emptyset = 0$ and

$$K_B = |Q| \cdot (1 + K') \cdot R \left(|Q|^{\mathbb{P}} \times 2^{|\mathcal{A}|(|B|-1)^{K'}} \times 2^{|B|}, N \cdot (1 + 2^{|B|}) \right),$$

where $K' = \sum_{(B', C') \in \mathcal{C}, B' \subseteq B} K_{B'}$.

Next lemma states that in a (N, \mathcal{C}) -broadcast game, very long strategies have useless threads.

Lemma 4. *Let σ be a distributed strategy of a (N, \mathcal{C}) -broadcast game. Assume that for some $(B, C) \in \mathcal{C}$, σ has a B -thread of length more than K_B . Then there is a useless thread in σ .*

Proof. Without losing generality, we can choose B -minimal for the inclusion so that for every B' -thread in σ with $(B', C') \in \mathcal{C}$ for some C' has length less than $K_{B'}$.

By hypothesis there exists $(B, C) \in \mathcal{C}$ and a prime σ -play uv such that $A(v) = B$ and $|v| \geq \frac{K_B}{|Q|}$.

We factorize v along the occurrences of C :

$$v = v_0 c_0 v_1 c_1 \cdots v_n c_n v_{n+1} \text{ with } v_i \in (B \setminus C)^* ,$$

and for $0 \leq i \leq n+1$, denote

$$z_i = v_0 c_0 v_1 c_1 \cdots v_i c_i .$$

We can choose the factorization such that each z_i is prime. The existence of such a factorization is shown by induction on n , for the base case, i.e. $n = 1$, if $v_0 = \partial_{c_0}(v_0)v'$ then $v = \partial_{c_0}(v_0)c_0v'v_1$ is a suitable factorization. The induction step uses the same trick and the fact that G_C is a clique so the c_i do not commute with each other and (10) applies recursively.

For every $0 \leq i \leq n$, we denote

$$q_i = (q_{p,i})_{p \in \mathbb{P}} \text{ the states of processes after play } uz_i,$$

$$P_i = \{w \in (B \setminus C)^* \mid uz_i c_i w \text{ is a } \sigma\text{-play}\} ,$$

And for every $0 \leq i, j \leq n$,

$$T_{i,j} = (c_i, q_i, (\sigma(uz_i c_i w))_{w \in P_i}, A(v_i c_i v_{i+1} c_{i+1} \cdots v_j c_j))$$

We establish an upper bound on the lengths of traces in P_i . Let $B \setminus C = B_1 \cup \dots \cup B_j$ the decomposition of $B \setminus C$ into connected components of $G_{B \setminus C}$. Then by definition of a recursive decomposition of A , for every $i \in 1..j$ there exists $(B_i, C_i) \in \mathcal{C}$.

Let

$$K' = \sum_{k=1}^j K_{B_k} .$$

Then

$$\forall k \in 1..j, \forall w \in P_k, |w| \leq K'$$

because each $w \in P_i$ it is the union of parallel threads whose alphabet is included in one of the B_k .

Thus $T_{i,j}$ can take at most $m = |A| \times |Q|^{|P|} \times 2^{|A|(B-1)^{K'}} \times 2^{|B|}$ different values.

Let $M = 2^{|B|}$. Each v_i has length at most K' thus $|v| \leq n \cdot (1 + K')$. Moreover by hypothesis, $|v| \geq \frac{K_B}{|Q|} = (1 + K') \cdot R(m, N \cdot (1 + M))$ thus $n \geq R(m, N \cdot (1 + M))$. By definition of Ramsey numbers, there are indices $0 \leq i_1 < i_2 < \dots < i_{(1+M) \cdot N} \leq n$ such that:

$$\forall 0 \leq k < l \leq 2N, T_{i_k, i_l} = T_{i_1, i_2}. \quad (75)$$

For every $k \in 1..N \cdot (1 + M)$, let

$$w_k = v_{i_k} c_{i_k} v_{i_k+1} c_{i_k+1} \dots v_{i_{k+1}-1} c_{i_{k+1}-1} \in B^*,$$

We are going to apply the definition of (N, \mathcal{C}) -broadcast game, $M + 1$ times, to each sequence of threads

$$S_l = (u, w_0 w_1 \dots w_i)_{i \in 1+l \cdot N..(l+1) \cdot N},$$

for every $l \in 0..M$. All conditions are met for this. First, each of the σ -plays $(u w_0 w_1 \dots w_i)_{i \in 1..2N}$ is prime because each $u w_0 w_1 \dots w_i$ can be rewritten $w_0 v_1 \dots v_j$. Second, (75) guarantees $\forall i, A(w_i) = A(w_0)$ and all w_i are $b = c_{i_1}$ -prime. Let $B' = A(w_0)$ then $b \in B'$ and $B' \subseteq B$.

By definition of (N, \mathcal{C}) -broadcasts game, applied consecutively to each sequence $(S_l)_{l \in 0..M}$, for every $l \in 0..M$, there exists $B_l'' \subseteq B'$ and $k_l \in 1 + l \cdot N..(l + 1) \cdot N$ such that $u w_0 \dots w_{k_l}$ is a B_l'' -broadcast in σ . Since $M = 2^{|B|}$, there exists two indices $0 \leq l_1 < l_2 \leq M$ such that $B_{l_1}'' = B_{l_2}''$. We denote B'' this common value.

Let $x = u w_0 \dots w_{l_1}$ and $y = w_{l_1+1} \dots w_{l_2}$. Then we show that (x, y) is a useless thread in σ , with the pair (B'', b) as a witness, for that we have to prove that properties (39), (41) and (42) in the definition of useless threads are satisfied. Properties (39) and (40) are satisfied by choice of x and y . Property (41) holds because of (75). To show that property (42) holds, take any $w \in (B'')^*$ such that $w \Vdash b$, we prove that $\sigma(xw) = \sigma(xyw)$. Since by hypothesis G_C is a clique and $b \in \mathcal{C}$ then $\forall c \in \mathcal{C}, \neg(b \Vdash c)$ thus

$$w \in (B \setminus \mathcal{C})^*.$$

According to (75), $T_{i_{l_1}, m} = T_{i_{l_2}, n}$ thus by definition of $T_{i,j}$, for every $w' \in (B \setminus \mathcal{C})^*$, $\sigma(xw') = \sigma(xyw')$, in particular when $w = w'$. This proves (42) and terminates the proof of the lemma. \square

Proof of Theorem 1. Let σ be a winning strategy of minimal duration. By minimality, according to Lemma 3, strategy σ does not contain any useless thread. Let (A_1, \dots, A_j) the connected components of (A, D) . Then since \mathcal{C} is an inductive decomposition, for every $i \in 1..j$ there exists $(A_i, C_i) \in \mathcal{C}$. Thus by Lemma 4, every play of σ has length less than $\sum_{i=1}^j K_{A_j}$. There is a finite number of distributed strategies with this property, and for each such strategy σ , there is a simple algorithm that checks whether σ is winning or not: look non-deterministically for a losing play consistent with σ . Thus the existence of a winning strategy is decidable. \square

5. Conclusion

We have presented a theorem that unifies several known decidability results for distributed games, and presented new examples of distributed games for which the existence of a winning strategy is decidable.

The decidability of distributed synthesis in the general case is still open to our knowledge, even in the simple case of ring games

where $G_{\mathbb{P}}$ is a simple cycle of length 5, or in the case where the automaton is synchronizing.

References

- V. Diekert and A. Muscholl. A note on Métivier's construction of asynchronous automata for triangulated graphs. *Fundamenta Informaticae*, 25(3):241–246, 1996.
- V. Diekert and G. Rozenberg. *The Book of Traces*. World Scientific, 1995. ISBN 9789810220587. URL <https://books.google.co.uk/books?id=vNFL0E2pjuAC>.
- B. Finkbeiner and S. Schewe. Uniform distributed synthesis. In *Logic in Computer Science, 2005. LICS 2005. Proceedings. 20th Annual IEEE Symposium on*, pages 321–330. IEEE, 2005.
- P. Gastin, B. Lerman, and M. Zeitoun. Distributed games with causal memory are decidable for series-parallel systems. In *FSTTCS 2004: Foundations of Software Technology and Theoretical Computer Science, 24th International Conference, Chennai, India, December 16-18, 2004, Proceedings*, pages 275–286, 2004. doi: 10.1007/978-3-540-30538-5_23. URL http://dx.doi.org/10.1007/978-3-540-30538-5_23.
- B. Genest, H. Gimbert, A. Muscholl, and I. Walukiewicz. Asynchronous games over tree architectures. In *Automata, Languages, and Programming - 40th International Colloquium, ICALP 2013, Riga, Latvia, July 8-12, 2013, Proceedings, Part II*, pages 275–286, 2013. doi: 10.1007/978-3-642-39212-2_26. URL http://dx.doi.org/10.1007/978-3-642-39212-2_26.
- P. Madhusudan, P. S. Thiagarajan, and S. Yang. The MSO theory of connectedly communicating processes. In *FSTTCS 2005: Foundations of Software Technology and Theoretical Computer Science, 25th International Conference, Hyderabad, India, December 15-18, 2005, Proceedings*, pages 201–212, 2005. doi: 10.1007/11590156_16. URL http://dx.doi.org/10.1007/11590156_16.
- A. Muscholl, I. Walukiewicz, and M. Zeitoun. A look at the control of asynchronous automata. 2008. URL <http://www.labri.fr/perso/anca/Publications/mwz08thiagu.pdf>.
- A. Pnueli and R. Rosner. Distributed reactive systems are hard to synthesize. In *Foundations of Computer Science, 1990. Proceedings., 31st Annual Symposium on*, pages 746–757. IEEE, 1990.
- P. J. Ramadge and W. M. Wonham. The control of discrete event systems. *Proceedings of the IEEE*, 77(1):81–98, 1989.
- W. Zielonka. Notes on finite asynchronous automata. *ITA*, 21(2):99–135, 1987.