



HAL
open science

On the controller synthesis problem for distributed systems with causal memory

Hugo Gimbert

► **To cite this version:**

Hugo Gimbert. On the controller synthesis problem for distributed systems with causal memory. 2016. hal-01259151v7

HAL Id: hal-01259151

<https://hal.science/hal-01259151v7>

Preprint submitted on 12 Dec 2016 (v7), 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.

On the controller synthesis problem for distributed systems with causal memory

Hugo Gimbert

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

Abstract. The decidability of the distributed version of the Ramadge and Wonham controller synthesis problem [12], where both the plant and the controllers are modelled as Zielonka automata [13,1] and the controllers have causal memory is a challenging open problem [8,9]. There exists three classes of plants for which the existence of a correct controller with causal memory 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. We provide a unified proof of these three results, and design a class of plants, called broadcast games, with a decidable controller synthesis problem. This leads to new examples of decidable architectures.

1 Introduction

The decidability of the distributed version of the Ramadge and Wonham control problem [12], where both the plant and the controllers are modelled as Zielonka automata [13,1] and the controller have causal memory is a challenging open problem. Very good introductions to this problem are given in [8,9].

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 it cannot block the *uncontrollable* actions from the environment. The choice of the local controller is based on two sources of information:

- First the controller monitors the sequence of states and actions of its process. This information is called the *local view* of the controller.
- Second when a shared action is played by several processes then all the controllers of these processes can exchange as much information as they want. In particular together they can compute their mutual view of the global execution: their *causal past*.

A correct controller restricts controllable actions so that every possible execution of the plant satisfies some specification. The controller synthesis problem aims at automatically computing a correct controller from its specification.

The difficulty of controller synthesis depends on several factors, e.g.:

- the size and architecture (pipeline, ring, ...) of the system,

- the information available to the controllers,
- the specification.

Assuming that processes can exchange information upon synchronization and use their causal past to take decisions is a one of the key aspects to get decidable synthesis problems [4]. In early work on distributed controller synthesis, for example in the setting of [11], the only source of information available to the controllers is their local view. In this setting, distributed synthesis is not decidable in general, except for very particular architectures like the pipeline architecture. The paper [3] proposes information forks as an uniform notion explaining the (un)decidability results in distributed synthesis. The idea of using causal past as a second source of information appeared in [4].

We adopt a modern terminology and call the plant a *distributed game* and the controllers are *distributed strategies* in this game. A distributed strategy is a function that maps the causal past of processes to a subset of controllable actions. In the present paper we focus on the *local reachability condition*, which is satisfied when each process is guaranteed to terminate its computation in finite time, in a final state. A distributed strategy is winning if it guarantees the local reachability condition, whatever non-deterministic choices are performed by the environment.

There exists three classes of plants for which the existence of a winning distributed strategy has been shown decidable:

1. when the dependency graph of actions is series-parallel [4],
2. when the processes are connectedly communicating [7],
3. and when the dependency graph of processes is a tree [5,10].

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 [4], for a different setup than ours: in the present paper we focus on process-based control while [4] was focusing on action-based control. Actually action-based control is more general than process-based control, see [8] for more details. 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 [4] was the first discovery of a class of asynchronous distributed system with causal memory for which controller synthesis is decidable

Connectedly communicating games have been introduced [7] 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 [7] which implies that controller synthesis is decidable for these games.

An acyclic game as defined in [5] is a game where processes are arranged as a tree and actions are either local or synchronize a father and its son. Even in this simple setting the synthesis problem is provably non-elementary hard.

Our contribution We develop a new proof technique to address the distributed controller synthesis problem, and provide a unified proof of decidability for series-parallel, connectedly communicating and acyclic games. We design a class of games, called *broadcast games*, which has a decidable controller synthesis problem. This leads to new examples of decidable architectures for controller synthesis, for example triangulated games and DAG games.

The new proof technique consists in simplifying a winning strategy by looking for useless parts to be removed in order to get a smaller winning 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 a cut-and-paste operation which makes the strategy smaller. By taking shortcuts again and again, we make the strategy smaller and smaller, until it does not have any useless thread anymore. Strategies with no useless thread 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 centralized game. In a centralized game with only one process, strategies are trees and one can cut a subtree from a point A and paste it to any other node B in A, and the operation makes sense as long as the unique process is in the same state in A and B. But in the case of a general distributed strategy, 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. The decidability of series-parallel games established [4] 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. Here we work without the series-parallel assumption.

This is the reason for introducing the notion of *broadcast*. A broadcast is a part of a strategy where an information is guaranteed to spread in a pool of processes before any of these processes synchronize with a process outside the pool. When two broadcasts are similar, they can be used to perform cut-and-paste: 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 *process ordering* of the set of actions. This notion is useful to define the new class of broadcast game and treat examples uniformly. However, to our opinion, the main contribution of the paper is not the notion of process ordering and broadcast games but rather the notion of useless threads and shortcuts and their use to simplify strategies.

The complexity of our algorithm is really bad, so probably this work has no immediate practical applications. This is not surprising since the problem is non-elementary [5]. Nevertheless we think this paper sheds new light on the difficult open problem of distributed synthesis.

Missing proofs and further examples can be found in the appendix and in [6].

2 Definitions and basic properties

2.1 Mazurkiewicz traces

The theory of Mazurkiewicz traces is very rich and extensively developed in [1]. 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 $\text{Alph}(u)$ the set of letters of u and we write:

$$u \mathbb{I} v$$

whenever $\text{Alph}(u) \times \text{Alph}(v) \subseteq \mathbb{I}$. 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:

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

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 $\text{Alph}(u)$ 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 $\text{Alph}(u) \subseteq B$. We denote B^* the set traces on an alphabet $B \subseteq A$.

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 $v = wu$.

2.2 Prime traces and views

A trace $u \in A^*$ is *prime* if all its linearizations 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 $v \perp B$. If B is a singleton $\{b\}$ then the B -view is also called the b -view and denoted $\partial_b(u)$.

2.3 Processes and automata

Zielonka automata are to traces what finite automata are to finite words.

Definition 1. A Zielonka automaton \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 $\text{dom}(a)$,
- $\Delta \subseteq \bigcup_{a \in A} \{a\} \times \prod_{p \in \text{dom}(a)} Q_p \times Q_p$ 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 $p \in \text{dom}(a)$.

The automaton \mathcal{A} defines a dependency relation D and its dual commutation relation \perp on A : two letters can commute if and only if they have no process in common in their domains:

$$\begin{aligned} ((a, b) \in D) &\iff (\text{dom}(a) \cap \text{dom}(b) \neq \emptyset) , \\ a \perp b &\iff \text{dom}(a) \cap \text{dom}(b) = \emptyset . \end{aligned}$$

This naturally defines a notion of Mazurkiewicz trace on 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 on each other,

$$\forall p \in \mathbb{P}, \forall u \in A^*, \partial_p(u) \text{ is prime.} \tag{1}$$

Moreover for $p \in \mathbb{P}$ and $u \in A^*$, $p \perp u$ is a notation for $A_p \perp u$. We extend the notion of domain to traces:

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

2.4 Plays

A play is an asynchronous computation of the automaton defined as follows.

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 $p \in \mathbb{P}$,

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

Intuitively $Q(u)$ denotes the last global state of the play u . 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 $Q(u)$. This holds because $\forall u \in \text{plays}(\mathcal{A}), Q_p(u) = Q_p(\partial_p(u))$, which can be easily proved inductively.

2.5 Strategies and games

Given an automaton \mathcal{A} , we would like the processes to choose actions so that every play eventually terminates in a final state of \mathcal{A} .

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

The choice of actions by processes is dynamic: at every step, a process p chooses a new set of controllable actions, depending on its (local) information about the way the play is (globally) going on.

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. In this case they exchange as much information about the play as they want. In particular it allows them to compute at this instant their common view of the play, i.e. their causal past. Formally, for every a -prime play $ua \in \text{plays}(\mathcal{A})$, this view is:

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

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

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

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

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 $p \in \text{dom}(a_i)$, $a_i \in \sigma_p(a_1 \cdots a_{i-1})$. A σ -play is maximal if it is not the strict prefix of another σ -play.

Note that a strategy is forced to allow every environment action to be executed at every moment. This may seem to be a huge strategic advantage for the environment. However depending on the current state, not every action can be effectively used in a transition because the transition function is not assumed to be total. So in general not every environment actions can occur. In particular it may happen that a process enters a final state with no outgoing transition, where no further action is possible.

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 of plays consistent with σ is finite and for every maximal σ -play u , every process is in a final state i.e.

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

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

3 Taking shortcuts

In this section we present an elementary operation called a *shortcut*, which can be used to simplify and reduce the duration of a winning strategy.

To create a shortcut, one selects a play $xy \in [\sigma]$ consistent with a strategy σ and modify the strategy so that as soon as a process sees the play x in its causal past in their causal past, they assume that not only x but also xy has actually occurred.

From a more formal point of view, a shortcut is a kind of *cut-and-paste* in the strategy tree: we glue on node x the subtree rooted at node xy . This guarantees the new tree to be smaller than the previous one, thus if the new tree is a winning strategy, its duration is guaranteed to be smaller than the duration of the original strategy.

However the choice of x and y should be carefully performed so that the new tree is well-defined and is still a winning strategy. We provide a sufficient condition for that: (x, y) should be a *useless thread*.

3.1 Duration and shortcuts

Winning strategies with shorter durations are better.

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*

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

Taking a shortcut is a convenient way to make a strategy shorter.

Definition 6 (Shortcut). *Let $x, y \in A^*$ be such that xy is a σ -play. Let $\phi_{x,y} : A^* \rightarrow A^*$ be the mapping:*

$$\phi_{x,y}(u) = \begin{cases} u & \text{if } x \not\sqsubseteq u \\ xyv & \text{if } u = xv \end{cases} , \quad (2)$$

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: there exists at most one 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 broadcasts and useless thread.

3.2 Broadcasts and useless threads

Intuitively, a *broadcast* is a prime play u in a strategy such that the maximal action of the play and the associated information about the play is broadcasted in priority to a pool of processes $\mathbb{Q} \subseteq \mathbb{P}$, before any of these processes synchronize with a process outside the pool. During a broadcast, as long as a process of the pool \mathbb{Q} plays in parallel of u then it synchronizes only with processes in \mathbb{Q} .

Definition 7 (broadcasts). *Let $\mathbb{Q} \subseteq \mathbb{P}$ be a subset of processes. We say that a prime play u is a \mathbb{Q} -broadcast if for every play w and every action a maximal in w ,*

$$(\text{dom}(a) \cap \mathbb{Q} \neq \emptyset) \text{ and } (\text{dom}(a) \cap \mathbb{P} \setminus \mathbb{Q} \neq \emptyset) \implies (u \sqsubseteq \partial_a(w)) . \quad (3)$$

In other words, a broadcast prevents a process in \mathbb{Q} to synchronize with a process out of \mathbb{Q} in parallel of u .

Two basic properties of broadcasts are listed in the following lemma.

Lemma 1. *Let u be a prime play with a maximal action b . Then u is a $\text{dom}(b)$ -broadcast and a \mathbb{P} -broadcast. Moreover, for every $\mathbb{Q} \subseteq \mathbb{P}$, u is a \mathbb{Q} -broadcast if and only if it is a $(\mathbb{P} \setminus \mathbb{Q})$ -broadcast. And u is a \mathbb{Q} -broadcast whenever \mathbb{Q} is a singleton.*

We give some more examples of broadcasts.

In an acyclic game, where processes are arranged in a tree organization (see the section with examples for a formal definition), every time a process p plays then it is a broadcast to its subtree: whenever a process in the subtree synchronizes with processes out of the subtree, it should synchronize with p as well.

In a series-parallel game, the dependency alphabet (A, D) is the product of two alphabets (A_0, D_0) and (A_1, D_1) , either parallel with $D = D_0 \cup D_1$ or synchronized with $D = D_0 \cup D_1 \cup A_0 \times A_1 \cup A_1 \times A_0$. Then every prime play whose maximal action is $a_0 \in A_0$ is a broadcast to $\text{dom}(a_0) \cup A_1$. The reason is that a_0 conflicts with every action in A_1 .

The definition of broadcasts can be equivalently reformulated as follows.

Proposition 1. *A prime play u is a \mathbb{Q} -broadcast if and only if for every uv such that v is prime, $(uv \text{ is prime})$ or $(v \parallel \mathbb{Q})$ or $(\text{dom}(v) \subseteq \mathbb{Q})$.*

We are interested in broadcasts occurring in *threads*.

Definition 8 (Thread). *Let $\mathbb{Q} \subseteq \mathbb{P}$ be a subset of processes (resp. $B \subseteq A$ a subset of actions). A \mathbb{Q} -thread (resp a B -thread) is a pair $(u, v) \in A^* \times A^*$ such that uv is a play and $\text{dom}(v) \subseteq \mathbb{Q}$ (resp. $v \in B^*$).*

Some threads are called *useless threads*, they can be deleted to reduce the delay of the strategy, for the same result.

Definition 9 (Useless thread). *Let σ be a strategy. A useless thread in σ is a \mathbb{Q} -thread (x, y) such that there exists an action b with the following properties:*

$$x \text{ and } xy \text{ are } b\text{-prime}, \quad (4)$$

$$x \text{ and } xy \text{ are } \mathbb{Q}\text{-broadcasts} \quad (5)$$

$$\text{every process } p \in \mathbb{P} \text{ has the same state in } x \text{ and } xy, \quad (6)$$

$$\text{for every play } xv \in A^*, (\text{dom}(v) \subseteq \mathbb{Q} \wedge v \parallel b) \implies (\sigma(xv) = \sigma(xyv)) \quad . \quad (7)$$

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

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) be 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*

$$\text{dur}(\sigma_{x,y}) < \text{dur}(\sigma) \quad . \quad (8)$$

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

$$xv \text{ is a } \sigma_{x,y}\text{-play} \iff xyv \text{ is a } \sigma\text{-play}. \quad (9)$$

According to this lemma, when a distributed game is winning, the winning strategy with the shortest duration has no useless thread. In the next section we introduce a class of games, called broadcast games, where broadcasts occur quite regularly. This is the key to obtain decidability of the synthesis problem: winning strategy with long durations can be simplified into simpler one by taking shortcuts. As a consequence we obtain a computable upper bound on the duration of the longest winning strategy.

3.3 Application: series-parallel games are decidable

We give an application of the results obtained so far : the synthesis problem for series-parallel games is decidable. Actually, for this result it is enough to use the notion of useless thread in the case where $\mathbb{Q} = \mathbb{P}$. This is summarized by the following corollary.

Corollary 1. *Let σ be a strategy with no useless \mathbb{P} -thread. Then for every primary play xy , if both x and y have the same maximal letter b and processes are in the same states after x and xy then*

$$\{v \in A^* \mid v \parallel b \wedge xv \in [\sigma]\} \neq \{v \in A^* \mid v \parallel b \wedge xyv \in [\sigma]\} .$$

Proof. Otherwise (x, y) would be a useless \mathbb{P} -thread. □

A series-parallel game is a game where 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 [4], for a slightly more general setup than ours called *action-based synthesis*, a setting more general than process-based control [8].

In a series-parallel game, either A is a singleton or there exists two cographs $G_0 = (A_0, D_0)$ and $G_1 = (A_1, D_1)$, A is the disjoint union of A_0 and A_1 and

$$\begin{aligned} &\text{either } D = D_0 \cup D_1 \text{ (parallel product)} \\ &\text{or } D = D_0 \cup D_1 \cup A_0 \times A_1 \text{ (synchronized product)} . \end{aligned}$$

We show by induction that if a strategy σ has no useless \mathbb{P} -thread then every prime play $u \in [\sigma]$ consistent has length at most K_A where K_A is inductively defined by:

$$K_A = \begin{cases} |Q| & \text{(singleton case)} \\ \max\{K_{A_0}, K_{A_1}\} & \text{(parallel case),} \\ K_{A_0} 2^{|A_0|^{K_{A_0}}} |A_0| |Q|^{|\mathbb{P}|} + K_{A_1} 2^{|A_1|^{K_{A_1}}} |A_1| |Q|^{|\mathbb{P}|} & \text{(synchronized case)} . \end{cases}$$

In the singleton case, this is a direct consequence of the corollary.

In the case of a parallel product, every prime play is either in A_0^* or in A_1^* , so the inductive step is easy.

Assume the product is synchronized. Any primary play u factorizes uniquely as $u = u_1 u_2 \cdots u_n$ where each prefix $u_1 u_2 \cdots u_i$ is primary and for every $i < n$, one of the two words u_i, u_{i+1} is a non-empty word in A_0^+ and the other a non-empty word in A_1^+ . Without loss of generality, assume $u_0 \in A_0^+$ (thus $u_1 \in A_1^+, u_2 \in A_0^+$ and so on).

Since σ has no useless \mathbb{P} -thread, then for every i , the strategy

$$\sigma[u_0 u_1 \cdots u_i] : w \rightarrow \sigma(u_0 u_1 \cdots u_i w)$$

has no useless \mathbb{P} -thread either. Thus, according to Corollary 1, for every pair u_i, u_j with $i < j$, if u_i and u_j have the same maximal action b and every process is in the same state after u_i and u_j then there exists a prime play $v \parallel b$ such that $\sigma(u_0 u_1 \cdots u_i v) \neq \sigma(u_0 u_1 \cdots u_j v)$. Without loss of generality, assume $b \in A_0$. Since $v \parallel b$ then $v \in (A_0 \setminus \text{dom}(b))^*$ because the product is synchronized thus $\{b\} \times A_1 \subseteq D$. By induction hypothesis, $|v| \leq K_{A_0}$.

As a consequence, there are at most $|A_0| |Q|^{|P|} 2^{|A_0|^{K_{A_0}}}$ different possible values for the words u_i whose maximal action is in A_0 . Thus $n \leq |A_0| |Q|^{|P|} 2^{|A_0|^{K_{A_0}}} + |A_1| |Q|^{|P|} 2^{|A_1|^{K_{A_1}}}$. Since every prefix $u_1 u_2 \dots u_i$ is primary then every u_i is primary. Moreover, by definition of K_{A_0} and K_{A_1} , every word $u_i \in A_0^+$ (resp $u_i \in A_1^+$) has length at most K_{A_0} (resp. K_{A_1}). This terminates the proof of the inductive step.

When the game is winning, the winning strategy with the shortest duration has no useless \mathbb{P} -thread. Thus winning strategies can be looked for in a finite set. This gives decidability of the synthesis problem for series-parallel games.

4 Broadcast games

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

The notion of broadcast games is defined so that decidability results of [7,5] can be retrieved quite easily. The definition of a broadcast game is rather ad-hoc and technical, we hope further research will lead to a more general and cleaner definition.

4.1 Definition

The proof of decidability of broadcast games is performed inductively, and relies on the notion of an process ordering of the dependency alphabet (A, D) .

Definition 10 (Process ordering). *A process ordering of \mathbb{P} is a partial order ¹ \preceq on \mathbb{P} such that for every prime trace u , $\text{dom}(u)$ has a \preceq -maximum.*

¹ i.e. a transitive, reflexive and anti-symmetric binary relation.

In a distributed game, there may exist several process orderings, for example any total order on \mathbb{P} is a process ordering of \mathbb{P} .

A broadcast game is a game where, periodically, processes create broadcasts. The pool of processes where the broadcast occurs is computed by downward-closure of the set of processes, with respect to the process ordering \preceq .

Definition 11 (Process closure). *Let \preceq be a process ordering of \mathbb{P} and $\mathbb{Q} \subseteq \mathbb{P}$. The process-closure of \mathbb{Q} is:*

$$\mathbb{Q}_{\preceq} = \{p \in \mathbb{P} \mid p \preceq q \text{ for some } q \in \mathbb{Q}\} .$$

We are especially interested in broadcasts consistent with process ordering.

Definition 12 (Well-ordered broadcast). *A \mathbb{Q} -broadcast u is well-ordered if $\mathbb{Q} = \mathbb{Q}_{\preceq}$ and the maximal process of $\text{dom}(u)$ plays in the last action of u .*

For every trace $u \in A^*$ and process $p \in \mathbb{P}$, $|u|_p$ denotes the number of actions of process p in u , i.e. the number of letters in u whose domain contains p .

Definition 13 (Broadcast games). *Let $N \in \mathbb{N}$. A game G is a (N, \preceq) -broadcast game if, for every prime play $uv \in A^*$, whenever*

$$\forall q \in \text{dom}(v), |v|_q \geq N$$

then there exists a prefix $v' \sqsubseteq v$ such that uv' is a well-ordered $\text{dom}(v)_{\preceq}$ -broadcast.

In case G is an (N, \preceq) -broadcast game for some N and \preceq we also say that G is an N -broadcast game or simply a broadcast game.

The property of being a broadcast game is decidable.

Proposition 2. *It is decidable whether a game G is a broadcast game. In case it is then G is an N -broadcast-game for some $N \leq \prod_{p \in \mathbb{P}} |Q_p|$.*

4.2 Decidability

Theorem 1. *Whether a broadcast game is winning or not is decidable.*

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 (10) below. The bound is quite large, which is not surprising since the problem is non-elementary [5]. We provide some examples and applications in the next section.

The proof is easy to sketch but harder to implement because distributed systems are not so easy to handle. For every subset of processes $\mathbb{Q} \subseteq \mathbb{P}$, we compute inductively a bound $K_{\mathbb{Q}}$ such that any winning strategy which has a \mathbb{Q} -thread of duration more than $K_{\mathbb{Q}}$ can be simplified in a shorter winning strategy, by taking a shortcut associated to a useless thread.

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

$$K_{\mathbb{Q}} = |\mathbb{Q}| \cdot R \left(2^{|\mathbb{Q}|}, N2^{|A|} |A| |\mathbb{Q}| 2^{|A|^{\max_{\mathbb{Q}' \subseteq \mathbb{Q}} K_{\mathbb{Q}'}}} \right) . \quad (10)$$

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

Lemma 4. *Let σ be a distributed strategy of a (N, \preceq) -broadcast game. Assume that for some $\mathbb{Q} \subseteq \mathbb{P}$, σ has a \mathbb{Q} -thread of length more than $K_{\mathbb{Q}}$. Then there is a useless thread in σ .*

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. By Lemma 4, every play of σ has length less than $K_{\mathbb{P}}$. 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 Examples of broadcast games

In this section we provide several examples of broadcast games. The first example is the class of connectedly communicating games, whose decidability was already known. The second example are acyclic games whose decidability was already known in the case where all actions are local or binary. The last example is the class of triangulated games. We terminate with a discussion about dynamic broadcast games.

5.1 Connectedly communicating games are broadcast games

Connectedly communicating games have been introduced [7] 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 there exists some bound $k \in \mathbb{N}$ such that whenever a process q never plays while process p plays k times, then p and q will stay forever in separate threads. Formally, 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 ,

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

It is quite clear that every k -communicating game is a k -broadcast game. The process ordering is an arbitrary total order \preceq on \mathbb{P} . Then, for every play uv , the

hypothesis $\forall q \in \text{dom}(v), |v|_q \geq k$ in the definition of broadcast games implies that no process $q \in \text{dom}(v)$ will ever synchronize with a process $p \notin \text{dom}(v)$ after uv thus uv itself is a (well-ordered) $\text{dom}(v)_{\preceq}$ -broadcast.

Thus every connectedly communicating is a k -broadcast game.

5.2 DAG games are 1-broadcast games

An acyclic game as defined in [5] is a game where processes are arranged as a tree and actions are either local or synchronize a father and its 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 [5] to the case of non-binary actions and we relax the assumption about the tree structure into the existence of a process ordering $\leq_{\mathbb{P}}$ on \mathbb{P} such that for every action $a \in A$ and processes $p_0, p_1, p_2 \in \mathbb{P}$,

$$(p_0 \in \text{dom}(a) \wedge p_1 \in \text{dom}(a) \wedge p_0 \leq_{\mathbb{P}} p_2) \implies (p_1 \leq_{\mathbb{P}} p_2 \vee p_2 \in \text{dom}(a)) \quad . \quad (11)$$

This condition holds in particular in the case where there is a tree structure $T_{\mathbb{P}} = (\mathbb{P}, E_{\mathbb{P}})$ on \mathbb{P} , $\leq_{\mathbb{P}}$ is the ancestor relation, and the domain of every action is a connected subset of $T_{\mathbb{P}}$. Indeed, in this case if $p_0 \leq_{\mathbb{P}} p_2$ then either $p_2 \in \text{dom}(a)$ or p_2 is an ancestor of all processes in $\text{dom}(a)$ thus in particular $p_1 \leq_{\mathbb{P}} p_2$. Also this shows that four player games, ($\mathbb{P} = \{1, 2, 3, 4\}$) where all actions are allowed except those with supports $\{1, 4\}$, $\{1, 3\}$ and $\{2, 4\}$ are decidable.

Every DAG game is a $(1, \preceq)$ -broadcast game, whenever property (11) is satisfied. Let uv be a primary play. We shall find a prefix $v' \sqsubseteq v$ such that v' is a $\text{dom}(v)_{\preceq}$ -broadcast. Let p be the maximal process in $\text{dom}(v)$ and v' the shortest prime prefix of v whose last action is an action of p . Then we show that v' is a $\text{dom}(v)_{\preceq}$ -broadcast. Let $uv'w$ be a play and c be a maximal action of $uv'w$ such that $\text{dom}(c) \cap \text{dom}(v)_{\preceq} \neq \emptyset$ and $\text{dom}(c) \not\subseteq \text{dom}(v)_{\preceq}$. Let $p_0 \in \text{dom}(c) \cap \text{dom}(v)_{\preceq}$ and $p_1 \in \text{dom}(c) \setminus \text{dom}(v)_{\preceq}$. Then p is the $\leq_{\mathbb{P}}$ -maximum of $\text{dom}(v)_{\preceq}$ thus $p_0 \leq_{\mathbb{P}} p$ and $\neg(p_1 \leq_{\mathbb{P}} p)$. Thus according to (11) we get $p \in \text{dom}(c)$. As a consequence $uv' = \partial_c(uv') \sqsubseteq uv'w$.

Thus every DAG game with property (11) is decidable.

5.3 Triangulated games and three player games

Any 3-player game G with processes $\{1, 2, 3\}$ is a 1-broadcast game, with the order $1 \preceq 2 \preceq 3$.

We show that any game G with alphabet A is a $(1, \preceq)$ -broadcast game. Let uv be a prime trace. We shall find a prefix $v' \sqsubseteq v$ such that uv' is a well-ordered $\text{dom}(v)_{\preceq}$ broadcast. The case where $\text{dom}(v)$ is a singleton is obvious. In the remaining case we select any primary prefix v' of v whose last action a is binary or ternary. This is a $\text{dom}(a)$ -broadcast because all players in $\text{dom}(a)$ can immediately observe uv' . Consequently, all 3-player games are 1-broadcast games and are decidable.

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

$$\forall a \in A, \text{dom}(a) \text{ is connected in } G_{\mathbb{P}}. \quad (12)$$

This definition is inspired by [2]. If we shrink all triangles of a triangulated game then we get a tree, thus we can fix a root of this tree and attribute a depth to these triangles. Then we order processes with respect to the average depth of the triangles it belongs to, and there is a natural notion of descendants, that we do not detail here.

Like in an acyclic game, every action of a process p is broadcasted to its descendants. The reason is every synchronisation of a descendant with a non-descendant process should include p , according to hypothesis (12) and by absence of simple cycles of length 4. Consequently, all triangulated games with property (12) are 1-broadcast games and are decidable.

5.4 Dynamic broadcast games and strategies

The condition for being a broadcast game is *dynamic*, by opposition to *static* restrictions on the architecture of processes or the dependency alphabet: a broadcast game may well behave for some time as a ring or another architecture whose decidability is unknown, as long as it performs a broadcast from time to time. Actually probably the requirement of the existence of broadcasts could be put on strategies rather than on the games themselves.

Conclusion

We have presented a theorem and proof techniques 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.

Probably the distributed synthesis problem is still decidable if the hypothesis about the regular occurrences of broadcasts is not put on the game itself but rather used to restrict the class of strategies used by the players.

The decidability of distributed synthesis in the general case is still open to our knowledge, even in the simple case of ring games with five processes. Another intriguing open problem is the case of *weakly k -connectedly communicating* plants. In such a plant, whenever two processes play both k times in a row without hearing from each other, they will never hear from each other anymore. Although this assumption seems only slightly lighter than the *k -connectedly communicating* one, we do not know how to use techniques of this paper to solve the controller synthesis for such plants.

Acknowledgements

We thank Blaise Genest for interesting discussions on the topic.

References

1. Diekert, V., Rozenberg, G.: The Book of Traces. World Scientific (1995), <https://books.google.co.uk/books?id=vNFL0E2pjuAC>
2. Diekert, V., Muscholl, A.: A note on Métivier's construction of asynchronous automata for triangulated graphs. *Fundamenta Informaticae* 25(3), 241–246 (1996)
3. Finkbeiner, B., Schewe, S.: Uniform distributed synthesis. In: Logic in Computer Science, 2005. LICS 2005. Proceedings. 20th Annual IEEE Symposium on. pp. 321–330. IEEE (2005)
4. Gastin, P., Lerman, B., Zeitoun, M.: 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. pp. 275–286 (2004), http://dx.doi.org/10.1007/978-3-540-30538-5_23
5. Genest, B., Gimbert, H., Muscholl, A., Walukiewicz, I.: Asynchronous games over tree architectures. In: Automata, Languages, and Programming - 40th International Colloquium, ICALP 2013, Riga, Latvia, July 8–12, 2013, Proceedings, Part II. pp. 275–286 (2013), http://dx.doi.org/10.1007/978-3-642-39212-2_26
6. Gimbert, H.: On the controller synthesis problem for distributed systems with causal memory. CoRR abs/1601.05176 (2016), <http://arxiv.org/abs/1601.05176>
7. Madhusudan, P., Thiagarajan, P.S., Yang, S.: 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. pp. 201–212 (2005), http://dx.doi.org/10.1007/11590156_16
8. Muscholl, A., Walukiewicz, I., Zeitoun, M.: A look at the control of asynchronous automata. In: K. Lodaya, M.M., N. Kumar, e. (eds.) Perspectives in Concurrency Theory. Universities Press, CRC Press (2009)
9. Muscholl, A.: Automated synthesis of distributed controllers. In: Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6–10, 2015, Proceedings, Part II. pp. 11–27 (2015), http://dx.doi.org/10.1007/978-3-662-47666-6_2
10. Muscholl, A., Walukiewicz, I.: Distributed synthesis for acyclic architectures. In: 34th International Conference on Foundation of Software Technology and Theoretical Computer Science, FSTTCS 2014, December 15–17, 2014, New Delhi, India. pp. 639–651 (2014), <http://dx.doi.org/10.4230/LIPIcs.FSTTCS.2014.639>
11. Pnueli, A., Rosner, R.: Distributed reactive systems are hard to synthesize. In: Foundations of Computer Science, 1990. Proceedings., 31st Annual Symposium on. pp. 746–757. IEEE (1990)
12. Ramadge, P.J., Wonham, W.M.: The control of discrete event systems. *Proceedings of the IEEE* 77(1), 81–98 (1989)
13. Zielonka, W.: Notes on finite asynchronous automata. *ITA* 21(2), 99–135 (1987)

Appendix

6 Elementary properties of traces

Not all properties of the concatenation operator and the prefix relation on words are preserved on traces, however the following are:

$$\forall u, v \in A^*, ((u \sqsubseteq v) \wedge (v \sqsubseteq u) \implies u = v) , \quad (13)$$

$$\forall u, v, w \in A^*, (uv = uw) \implies (v = w) , \quad (14)$$

$$\forall u, v, w \in A^*, (uv \sqsubseteq uw) \implies (v \sqsubseteq w) . \quad (15)$$

The following lemma lists some basic properties of traces used in the proofs.

Lemma 5. *For every trace $u, v, x \in A^*$ and $a \in A$ and $B \subseteq A$,*

$$(u \mathbb{I} B) \iff (\partial_B(u) = \epsilon) \quad (16)$$

$$(x \sqsubseteq uv) \implies \exists x_0 \sqsubseteq u, \exists x_1 \sqsubseteq v, x = x_0x_1 \quad (17)$$

$$(x \sqsubseteq uv) \implies \exists x_0, x_1, x_2, x_3 \in A^*, \quad (18)$$

$$(x = x_0x_1) \wedge (u = x_0x_2) \wedge (v = x_1x_3) \wedge (x_2 \mathbb{I} x_1)$$

$$uv \text{ is } B\text{-prime} \implies v \text{ is } B\text{-prime} \quad (19)$$

$$u \text{ and } v \text{ are } B\text{-prime} \implies uv \text{ is } B\text{-prime} \quad (20)$$

$$\text{If } ua \text{ is prime, } (av \text{ is } B\text{-prime} \iff uav \text{ is } B\text{-prime}) \quad (21)$$

$$(u \text{ is } B\text{-prime} \wedge \neg(a \mathbb{I} u)) \implies (au \text{ is } B\text{-prime}) \quad (22)$$

$$(u \sqsubseteq \partial_B(uv)) \iff (\partial_B(uv) = u \partial_B(v)) \quad (23)$$

$$(uw \sqsubseteq \partial_B(uv)) \implies (w \sqsubseteq \partial_B(v)) \quad (24)$$

$$\partial_B(\partial_B(u)) = \partial_B(u) \quad (25)$$

$$\partial_B(uv) = \partial_B(u \partial_B(v)) \quad (26)$$

If ua is prime,

$$uav \sqsubseteq \partial_B(uavw) \iff av \sqsubseteq \partial_B(avw) \quad (27)$$

Proof. The equivalence (16) is immediate from the definition of ∂_B .

Equation (17) is a corollary of (18) which is well-known, see [1] for example. It can be proved by induction on $|x|$.

We prove (19). 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 (20). 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 (21). Assume ua prime. The converse implication follows from (19). 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 (19) 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 ,

$$bw = \{xbz \mid x, z \text{ words on } A, xz \in w, b \mathbb{I} x\} . \quad (28)$$

Let y a linearization of $uav = bu'av$, we prove that the last letter of y is in B . According to (28), y factorizes as $y = xbz$ with $xz \in u'av$ and $x \mathbb{I} b$. 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 $x \mathbb{I} b$. 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 $x \mathbb{I} b$ 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 (21).

We prove (22) 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 $a \mathbb{I} u$ 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 $w' \mathbb{I} a$. 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 (23). The converse implication in (23) is obvious so it is enough to prove the direct implication. Assume $u \sqsubseteq \partial_B(uv)$. According to (13) 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 $w' \mathbb{I} B$. Then $v = ww'$ according to (14) and since $w' \mathbb{I} B$, 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 $w'' \mathbb{I} B$. Then $uv = u \partial_B(v)w''w'$ and $w''w' \mathbb{I} B$ 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 (15) $w \sqsubseteq \partial_B(v)$. Finally $w = \partial_B(v)$ and $u \partial_B(v) = uw = u \partial_B(v)$ which terminates the proof of (23).

Equation (24), is a direct corollary of (23). Let v', w' such that $\partial_B(uv) = uww'$ and $uv = \partial_B(uv)w'$. Then according to (23), $uww' = u \partial_B(wv'w')$ thus according to (14), $wv' = \partial_B(wv'w')$ and since $v = wv'w'$, 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 $v' \mathbb{I} B$, 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 (25). Since $\partial_B(u) \sqsubseteq u$, then $\partial_B(\partial_B(u)) \sqsubseteq \partial_B(u)$ according to (13) it is enough to prove $\partial_B(u) \sqsubseteq \partial_B(\partial_B(u))$. By definition of ∂_B , $u = \partial_B(u)u'$ with $u' \mathbb{I} B$ and $\partial_B(u) = \partial_B(\partial_B(u))u''$ with $u'' \mathbb{I} B$. Thus $u = \partial_B(\partial_B(u))u''u'$ with $u''u' \mathbb{I} B$ hence by definition of ∂_B , $\partial_B(u) \sqsubseteq \partial_B(\partial_B(u))$.

We prove (26). Since $u \partial_B(v) \sqsubseteq uv$, then $\partial_B(u \partial_B(v)) \sqsubseteq \partial_B(uv)$ and according to (13) 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 (25), $\partial_B(\partial_B(uv)) = \partial_B(uv)$ thus $\partial_B(uv) \sqsubseteq \partial_B(u \partial_B(v))$ which terminates the proof of (26).

We prove (27). Assume ua is prime. The direct implication is immediate using (24). For the converse implication, assume $av \sqsubseteq \partial_B(aw)$. Since $\partial_B(uavw) = \partial_B(uav \partial_B(w))$, without loss of generality we can assume $w = \partial_B(w)$ thus $\partial_B(aw) = aw$, thus we can replace v with vw and assume $w = \epsilon$. Then $\partial_B(av) = av$ thus according to (23) $v = \partial_B(v)$ 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 (18), $ua = u_0u_1$ 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 (23), $\partial_B(uav) = u \partial_B(av) = uav$ 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 (27). \square

7 Taking shortcuts

Proposition 1 A prime play u is a \mathbb{Q} -broadcast if and only if for every uv such that v is prime,

$$(uv \text{ is prime}) \text{ or } (v \parallel \mathbb{Q}) \text{ or } (\text{dom}(v) \subseteq \mathbb{Q}) . \quad (29)$$

Proof. For the direct implication, assume that u is a \mathbb{Q} -broadcast, v is prime and uv is not prime. Let b be the last action of v . Then $u \not\sqsubseteq \partial_b(uv)$. Thus by definition of a broadcast, either $\text{dom}(b) \subseteq \mathbb{Q}$ or $\text{dom}(b) \cap \mathbb{Q} = \emptyset$. And by induction the same holds for every letter b of v . Since v is prime then either $\text{dom}(v) \subseteq \mathbb{Q}$ or $\text{dom}(v) \cap \mathbb{Q} = \emptyset$. Which terminates the proof of the direct implication.

Conversely, assume that for every uv such that v is prime, condition (29) holds. Let a be a maximal action of uv .

If uv is prime then $u \sqsubseteq uv = \partial_a(uv)$ thus the right handside of the implication (3) in the definition of broadcasts is satisfied.

If uv is not prime then according to (29) either $v \parallel \mathbb{Q}$ or $\text{dom}(v) \subseteq \mathbb{Q}$. In the first case in particular $\text{dom}(a) \cap \mathbb{Q} = \emptyset$. In the second case in particular $\text{dom}(a) \subseteq \mathbb{Q}$. In both cases the implication (3) in the definition of broadcasts is satisfied because the assumption is false.

In both cases the condition defining a \mathbb{Q} -broadcast is satisfied. \square

The first clause in the disjunction (29) can be reformulated in several ways.

Proposition 3. 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 (27) $(a \sqsubseteq \partial_b(av))$. Assume $a \mathbb{I} v$ then in particular $a \mathbb{I} b$ and $\partial_b(av) = \partial_b(v) = v$ because v is b -prime thus $(a \not\sqsubseteq \partial_b(av))$. Assume $\neg(a \mathbb{I} v)$. Then av is b -prime according to (22) thus uv is b -prime according to (21). Assume uv is b -prime then $uv = \partial_b(uv)$ thus $u \sqsubseteq \partial_b(uv)$. \square

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:

$$\sigma_p(\partial_p(xyv)) = \tau_p(\partial_p(xv)) . \quad (30)$$

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 (30) 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

$$\partial_p(xyv) = xy \partial_p(v) . \quad (31)$$

Since $x \sqsubseteq \partial_p(xw)$, (23) implies

$$\partial_p(xw) = x \partial_p(w) . \quad (32)$$

Since (x, y) is a useless thread, according to (4), both x and xy are b -prime. Since moreover $\partial_p(xw) = x \partial_p(w)$ we can apply (27) twice and get first $\partial_p(bw) = b \partial_p(w)$ and then (31). Now that (31) is proved we can conclude the second case:

$$\sigma_p(\partial_p(xyw)) = \sigma_p(xy \partial_p(w)) \quad (33)$$

$$= \tau_p(x \partial_p(w)) \quad (34)$$

$$= \tau_p(\partial_p(xw)) , \quad (35)$$

where (33) comes from (31), (34) hold by definition of shortcuts and τ , and (35) comes from (32). Thus (30) holds in the second case.

We are now left with the third and last case:

$$x \not\sqsubseteq \partial_p(u) \wedge x \sqsubseteq u, \quad (36)$$

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 (16) thus $\partial_p(xyv) = \partial_p(xy)$. Hence,

$$\sigma_p(\partial_p(xyv)) = \sigma_p(\partial_p(xy)) \quad (37)$$

$$= \sigma_p(xy) \quad (38)$$

$$= \tau_p(x) \quad (39)$$

$$= \sigma_p(x) \quad (40)$$

$$= \sigma_p(\partial_p(x)) \quad (41)$$

$$= \sigma_p(\partial_p(xv)) \quad (42)$$

$$= \tau_p(\partial_p(xv)) \quad (43)$$

where (37) and (42) hold because $v \parallel p$, (38) and (41) hold because σ_p is a distributed strategy, (39) holds by definition of τ , (40) holds because (x, y) is a useless thread and according to (7), (43) holds by definition of τ and because by hypothesis $x \not\sqsubseteq \partial_p(u)$. This shows that (30) holds when $\partial_p(v) = \epsilon$.

Now assume that $\partial_p(v) \neq \epsilon$ (and we keep assuming (36) as well). Since (x, y) is a useless thread in σ , then according to (5), x is a \mathbb{Q} -broadcast in σ . We apply the characterisation of broadcasts given by Proposition 1, to $x \partial_p(v)$, which is allowed because $\partial_p(v)$ is prime and $x \partial_p(v)$ is a σ -play because it is a prefix of the σ -play xv . Thus according to Proposition 3, one of the three following properties holds:

$$x \sqsubseteq \partial_p(x \partial_p(v)) \quad (44)$$

$$\text{or } \text{dom}(\partial_p(v)) \subseteq \mathbb{Q} \quad (45)$$

$$\text{or } \text{dom}(\partial_p(v)) \parallel \mathbb{Q} \quad (46)$$

Since $x \not\sqsubseteq \partial_p(x \partial_p(v))$ by hypothesis, (44) is not possible and we are left with the two other cases (45) and (46).

We assume first that (46) holds. Since $\partial_p(v) \neq \epsilon$, it implies that $p \notin \mathbb{Q}$. Since (x, y) is a \mathbb{Q} -thread then $\text{dom}(y) \subseteq \mathbb{Q}$ thus $p \parallel y$. We can conclude the proof of (30) in the case where (46) holds:

$$\sigma_p(\partial_p(xyv)) = \sigma_p(\partial_p(xy \partial_p(v))) \quad (47)$$

$$= \sigma_p(\partial_p(x \partial_p(v) y)) \quad (48)$$

$$= \sigma_p(\partial_p(x \partial_p(v))) \quad (49)$$

$$= \sigma_p(\partial_p(xv)) \quad (50)$$

$$= \tau_p(\partial_p(xv)) \quad (51)$$

where (47) and (50) hold according to (26), (48) holds because $\partial_p(v) \parallel \mathbb{Q}$ and $\text{dom}(y) \subseteq \mathbb{Q}$, (49) holds because $p \parallel y$, and (51) holds by definition of τ_p , since by hypothesis $x \not\sqsubseteq \partial_p(u)$ and $u = xv$. This proves (30) in the case where (46) holds.

Now we are left with the case where (45) holds, i.e. $\text{dom}(\partial_p(v)) \subseteq \mathbb{Q}^*$ (and we keep assuming $\partial_p(v) \neq \epsilon$ and (36) as well). We first establish

$$\partial_p(v) \in (A \setminus \{b\})^* \quad (52)$$

Since by hypothesis $x \not\sqsubseteq \partial_p(xv)$ and x is b -prime then according to (27), $b \not\sqsubseteq \partial_p(bv)$ thus according to (22), $b \parallel \partial_p(v)$ which implies (52). Since (x, y) is a useless thread in σ , we can apply (7) to $\partial_p(v)$ hence

$$\sigma_p(x \partial_p(v)) = \sigma_p(xy \partial_p(v)) . \quad (53)$$

Finally,

$$\sigma_p(\partial_p(xyv)) = \sigma_p(\partial_p(xy \partial_p(v))) \quad (54)$$

$$= \sigma_p(xy \partial_p(v)) \quad (55)$$

$$= \sigma_p(x \partial_p(v)) \quad (56)$$

$$= \sigma_p(\partial_p(x \partial_p(v))) \quad (57)$$

$$= \sigma_p(\partial_p(xv)) \quad (58)$$

$$= \tau_p(\partial_p(xv)), \quad (59)$$

where equalities (54) and (58) hold according to (26), equalities (55) and (57) hold because σ is a distributed strategy, (56) comes from (53), and finally (59) is by definition of τ and because by hypothesis $x \not\sqsubseteq \partial_p(xv)$. This terminates the proof of (30) in the last case.

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

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*

$$\text{dur}(\sigma_{x,y}) < \text{dur}(\sigma) . \quad (60)$$

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

$$xv \text{ is a } \sigma_{x,y}\text{-play} \iff xyv \text{ is a } \sigma\text{-play}. \quad (61)$$

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

We first prove property (61). 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 (61) 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 (61), 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, (6) 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 (60). According to (61), the mapping $\phi_{x,y}$ used to define $\tau = \sigma_{x,y}$ in (2) maps maximal τ -plays to maximal σ -plays. Moreover, according to (14), 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 (4), $y \neq \epsilon$ thus we get property (60).

This terminates the proof of Lemma 3. \square

Lemma 4 *Let σ be a distributed strategy of a (N, \preceq) -broadcast game. Assume that for some $\mathbb{Q} \subseteq \mathbb{P}$, σ has a \mathbb{Q} -thread of length more than $K_{\mathbb{Q}}$. Then there is a useless thread in σ .*

Proof. Without losing generality, we can choose \mathbb{Q}_{\preceq} minimal for the inclusion so that every \mathbb{Q}'_{\preceq} -thread in σ has length less than $K_{\mathbb{Q}'_{\preceq}}$ whenever $\mathbb{Q}'_{\preceq} \subsetneq \mathbb{Q}_{\preceq}$.

By hypothesis there exists $\mathbb{Q} \subseteq A$ and a σ -play uv such that $\text{dom}(v) \subseteq \mathbb{Q}$ and $|v| \geq K_{\mathbb{Q}}$. Without loss of generality we can assume $\mathbb{Q}_{\preceq} = \mathbb{Q}$. And w.l.o.g., since there are at most $|\mathbb{Q}|$ maximal events in a trace, we can assume that uv is prime and $|v| \geq \frac{K_{\mathbb{Q}}}{|\mathbb{Q}|}$.

We consider the complete graph G_c with vertices $1, \dots, |v|$ and the label of the edge $\{i < j\}$ is defined as follows. We denote $v[i, j]$ the subword of v between positions i and j (both included) and $A[i, j] = \text{Alph}(v[i, j]) \subseteq \text{Alph}(v)$ and $P[i, j] = \text{dom}(v[i, j])_{\preceq} \subseteq \mathbb{Q}_{\preceq}$. Then for $1 \leq i_0 \leq i_1 \leq i_2 \leq |v|$ we set

$$E[i_0, i_1, i_2] = (A[i_0, i_2], a_{i_1}, Q_{i_1}, L[i_1]) \text{ with}$$

- a_{i_1} the letter at position i_1 in v ,
- Q_{i_1} the states of the processes in $uv[1, i_1]$.
- $L[i_1] = \{w \in A^* \mid (\text{dom}(w) \subseteq \mathbb{Q}_{\preceq}) \wedge (w \mathbb{I} \max_{\preceq} \text{dom}(v))\}$.

Let N_L be the number of possible values of $E[i_0, i_1, i_2]$ and $H = N * N_L$. Notice that every trace in $L[i_1]$ is a \mathbb{Q}' -thread with $\mathbb{Q}' = (\mathbb{Q} \setminus \max_{\preceq} \text{dom}(v))$. Since $\mathbb{Q}'_{\preceq} \subsetneq \mathbb{Q}_{\preceq}$ then we can apply the inductive hypothesis: every trace in $L[i_1]$ has length at most $|A|^{K'_{\mathbb{Q}'}}$. Thus

$$H = N * N_L \leq N * 2^{|A|} |A|^{|\mathbb{P}|} 2^{|A|} |A|^{K'_{\mathbb{Q}'}} .$$

By definition of $K_{\mathbb{Q}}$, $|v| \geq R(2^{\mathbb{Q}}, H)$. Hence by definition of Ramsey numbers, G_c contains a monochromatic clique $i_1 < i_2 < \dots < i_H$. Let $v = v_0 v_1 v_2 \dots v_H v_{H+1}$ be the corresponding factorization of v , such that $\text{dom}(v_1) = \dots = \text{dom}(v_H)$.

For every $1 \leq i \leq H$, every process in $\text{dom}(v_1)$ has played at least N times in $v'_i = v_{i*N+1} \dots v_{i*N+N}$. Thus, since the game is a broadcast game, for every $1 \leq i \leq N_L$, there exists a prefix $v''_i \preceq v'_i$ such that

$$uv_0 \dots v_{i*N} v''_i \text{ is a well-ordered } \text{dom}(v_1)_{\preceq}\text{-broadcast.}$$

By definition of N_L we can find two indices $1 \leq i < j \leq N_L$ such that $E(i, |v'_i|, i + |v_i|) = E(j, |v'_j|, j + |v_j|)$. Let $B = \text{Alph}(v'_i) = \text{Alph}(v'_j)$ and $\mathbb{Q}' = \text{dom}(v_1) = \text{dom}(v'_i)_{\preceq} = \text{dom}(v'_j)_{\preceq}$ and $x = v_0 \dots v_{i-1} v'_i$ and $y = v_0 \dots v_{j-1} v'_j$. Then both ux and uy are \mathbb{Q}' -broadcasts, end up with the same letter, processes have the same state and local strategies coincide ($L[i] = L[j]$). Thus all conditions are met and σ has a useless thread. \square

8 Broadcast games

Proposition 2 *It is decidable whether a game G is a broadcast game. In case it is then G is an N -broadcast-game for some $N \leq \Pi_{p \in \mathbb{P}} |Q_p|$.*

Proof. Let $M = \Pi_{p \in \mathbb{P}} |Q_p|$. Let \preceq be an inductive decomposition of A . A standard pumping argument of automata theory shows that the conditions in Definition 13 are satisfied for some N is and only if they are satisfied when $N = M$ and both traces u and v have length less than M .

Let $b \in A$ and $\mathbb{Q} \subseteq \mathbb{P}$. Let $I_b = \{a \in A \mid a \parallel b\}$ and $I_{\mathbb{Q}} = \{a \in A \mid a \parallel \mathbb{Q}\}$. Then a b -primary play u is a \mathbb{Q} -broadcast if and only if there does not exist a prime trace $v \in A^*$ such that uv is a play and

$$(v \notin I_b^*) \tag{62}$$

$$\wedge (\text{dom}(v) \not\subseteq \mathbb{Q}) \tag{63}$$

$$\wedge (v \notin I_{\mathbb{Q}}^*) . \tag{64}$$

The conjunction of these three conditions is indeed equivalent to the opposite of the definition fo broadcasts given in Proposition 1. Indeed, according to Proposition 3,

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

Again, a standard pumping argument shows that if there exists $v \in A^*$ which satisfies (62), (63) and (64) 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 $\mathbb{Q} \subseteq \mathbb{P}$ and for each of those, enumerating all u, v of length less than M , and for those which satisfies the conditions in Definition 13, enumerating all prefixes $v' \sqsubseteq v$ and check whether (62), (63) and (64) are satisfied with u replaced by uv' . If a witness is found then G is not a broadcast game, otherwise G is a broadcast game. \square